UCLA/Arisaka-GG-A10-Cosmos
v23.19 — September 11, 2026 — *Preprint*

The Standard Cosmology as Theorem of GG-Theory

The Cosmic Inventory as Observations of One Six-Dimensional Geometry

Katsushi Arisaka

Department of Physics and Astronomy, Electrical and Computer Engineering, University of California, Los Angeles, Los Angeles, CA 90095, USA

Contact: arisaka@physics.ucla.edu

Abstract

For forty years the standard model of cosmology, ΛCDM, has described the universe with breathtaking precision and explained almost none of it: a handful of numbers, read off the sky, reproduce everything, yet the theory cannot say why those numbers take the values they do.

This paper argues that they are not free. The same six-dimensional geometry that fixes the elementary particles, with three declared inputs and no dial adjusted to what it predicts, also forces the present-epoch *shape* of the cosmos: the matter fraction $1/\pi$, the dark-energy fraction $1 - 1/\pi$, the ordinary-matter fraction $1/(2\pi^2)$, and the geometry's switch at the epoch of the first atoms at redshift e^7, the exponential of a single whole number. Inflation, today's acceleration ($w = -1$), and a dark-matter particle near $8.1\,\mu$eV, the Kaxion, follow from the same shape, with the tilt and tensor ratio computed rather than fitted ($n_s = 0.963$, $r = 0.004$).

The two great puzzles of present-day cosmology appear in a new light: the Hubble tension as one derived expansion rate with the distance ladder as the outlier the geometry predicts, and the early, massive galaxies seen by JWST read structurally. As in its companion on particle physics, the paper states plainly which numbers are fully derived, which rest on a single named step, and which remain open, tuning nothing after the fact. If the argument holds, the cosmic inventory stops being a list of accidents: its form is forced, the 10^{120} vacuum fine-tuning dissolves, and what remains is one number read from the sky, *when* we look.

Executive Summary

The problem. ΛCDM — cold dark matter with a cosmological constant — is the rock-solid empirical baseline of contemporary cosmology, the most successful predictive framework after the Standard Model of particle physics. Six numbers, read off the data, compress every microwave-background anisotropy, every baryon-acoustic oscillation, every supernova and weak-lensing map into one grammar. But the framework is silent about *why* those numbers hold. Why is the universe roughly five percent ordinary matter, twenty-seven percent dark matter, sixty-eight percent dark energy? Why do the early- and late-universe measurements of the expansion rate disagree at more than six sigma — the *Hubble tension* — and why has the James Webb Space Telescope (JWST) found massive galaxies in place a few hundred million years after the Big Bang? Seventeen recognized problems, set out in Part I, sit outside the six-parameter language. This paper does not abandon ΛCDM; it asks whether its numbers can be *derived*. Beneath the seventeen sits a question the field has circled for half a century without naming, the *second horizon problem* (§72): inflation explains why disconnected regions share the same *conditions*, one temperature to a part in 10^5, but not why they hold the same *things* — why an electron across the universe is identical to one here, though no signal ever compared them. Part X poses it, and answers it.

The single input. The derivation rests on one object: a six-dimensional geometry, fixed once and for all, with no continuous free parameter. Its five integers $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ — the ledger this program calls OGN-5 — are not adjusted but forced, the unique set that survives the anomaly-free closure of the six-dimensional bulk (Arisaka, 2026, A3-BI6; Arisaka, 2026, A4-Euclid); $L = 5$ is the unification block, $N = 3$ the three colors. The same geometry admits exactly two transcendental constants, e for the exponential redshift along the radial fifth dimension and π for the angle of the compact sixth, so every exponential in the inventory is a scale factor and every π a holonomy or a volume. Beyond that ledger the cosmology declares three inputs, none a dial: an overall amplitude, the epoch at which the present snapshot is read, and one calibrated scale, fixed by the dark-matter abundance in the companion (Arisaka, 2026, C1-Kaxion). This paper reads that one geometry at the largest scale: our four-dimensional universe is its boundary, and the contents of the cosmos are what the geometry casts onto that boundary — observations of an open system, not free inputs (§6). The cosmos and the electron are read from one ledger (Arisaka, 2026, A9-SM), one story apart.

The headline results. The composition of the universe takes a forced shape. Read at the present epoch, the matter fraction is $\Omega_m = 1/\pi$, the dark-energy fraction $\Omega_\Lambda = 1 - 1/\pi$, and the ordinary-matter fraction $\Omega_b = 1/(2\pi^2)$, landing on the measured values to within a few percent, the baryon fraction the farthest at 2.3σ (Figure 1); their ratio $1/(2\pi)$ and the vacuum scale are fixed for all time, and only the present-day split is read from the sky. The switch that governs recombination descends from one whole number in the ledger, $z_* = e^7 \approx 1097$, within 0.62% of the measured decoupling redshift. Inflation is a derived plateau one e-fold below the grand-unification edge, with the tilt computed rather than fitted: $n_s = 0.9634$, within 0.4σ of *Planck* and 3.2σ below

the 2025 value that includes baryon-acoustic data — a tension the paper states rather than softens. Dark energy is the residual $1 - 1/\pi$ of the same inventory with $w = -1$ exactly, a two-sided bet that Euclid and Rubin will settle this decade. Dark matter is the Kaxion, a light particle placed near $8.1\,\mu$eV in a narrow laboratory window by the one calibration (Arisaka, 2026, C1-Kaxion). The Hubble tension is read as one derived expansion rate near 67, the value the microwave background gives, with the distance ladder as the outlier the geometry predicts (Arisaka, 2026, C2-Hubble); the age of the universe comes out at 13.9 billion years. The primordial gravitational-wave amplitude is predicted at $r = 0.0038$, a factor nine below the present bound and within reach of the Simons Observatory and LiteBIRD. And the three chapters — inflation, dark matter, dark energy — are one descent read three times (Figure 3); they cannot be tuned separately, and a failure in any one reverberates through all.

How confident. Twenty-one cosmological observables are classified by provenance: one derived outright, nineteen derived up to a named conjecture each (two for the dark-matter row), and one — the amplitude — declared (two calibrations, the observation epoch and the Kaxion's scale, enter beside it and are not codebook rows). No number is tuned to the data after the fact. The title's word *theorem* means exactly this conditional ledger, stated in the codebook of Part IX, and nothing beyond it. The zero-parameter temperature spectrum is reported at $\chi^2/\mathrm{dof} \approx 2.3$ — not a fit; the whole excess is a 0.3% deficit in the acoustic angle, and the shape is already right — and the tilt is recorded as the framework's one currently failing test (Figure 6). The matter–antimatter asymmetry is predicted from the baryon fraction ($\eta_{10} = 6.18$ against the measured 6.12); the leptogenesis mechanism meant to produce it falls short by a factor of three to ten, depending on how the calculation is done, and the paper says so rather than closing the gap. Deuterium and helium come out standard, and the lithium-7 deficit is read as a $1/\pi$ port on the $A = 7$ channel.

How it can die. A theory with no dials earns credibility only by exposure, here deliberate, two-sided and dated: a gravitational-wave siren reading of the expansion rate at the ladder's 73 rather than ≈ 67; a tensor amplitude off $r = 0.0038$ in either direction — beyond the launch's own spread, 0.0037–0.0040, and in any case outside $[10^{-3}, 10^{-2}]$; w driven off -1 at Euclid and Rubin precision; the 1.59–2.13 GHz Kaxion window swept at DFSZ depth and silent; a firm primordial lithium-7 abundance away from SBBN/π. Nothing can be renegotiated afterwards.

Against ΛCDM and the multiverse. The two standard accounts of the cosmic numbers do not derive them. ΛCDM measures its six parameters, a fitting language of exquisite accuracy. The inflationary multiverse proposes that the constants vary across an unobservable ensemble and that we inhabit one that permits observers — a selection that accommodates almost any value and therefore predicts none. One fixed geometry, with no ensemble to choose from and no dial to turn, is *obliged* to produce a definite form for each cosmic number; that is why it derives the coupled π-shape of the inventory rather than measuring it, and why a single wrong number can kill it.

Where this paper sits. This is the cosmological member of a twenty-two-paper program that begins with a constitution (Arisaka, 2026, A1-Constitution), builds the six-dimensional geometry

(Arisaka, 2026, A2-GG6; Arisaka, 2026, A3-BI6), and reads the Standard Model from it (Arisaka, 2026, A9-SM). The Kaxion's laboratory window is placed in Arisaka (2026, C1-Kaxion), the TeV-scale collider signature in Arisaka (2026, P1-KK), and the same geometry is carried toward life and mind in Arisaka (2026, A11-Origin). A reader needs none of these to follow the argument; the falsifiable core stands on its own.

The identity of the cosmos. Part X turns from the inventory's values to its *sameness*. Why is dark energy the same on every sightline, every Kaxion identical, the vacuum one vacuum? The identity theorem of the particle paper (Arisaka, 2026, A9-SM) lifts one story: one deepest basin selects the vacuum as it selects the electron, and one angle, held across the sky during inflation — a conjecture this paper names and does not prove — selects the Kaxion; two routes to identity, a settled minimum and a frozen phase. The classical alternatives, design, a multiverse, brute fact, are dissolved rather than refuted: a forced geometry leaves no ensemble to select from and no dial for a designer to have set.

A philosophical conclusion. If the argument holds, the largest questions change character. The mixture of matter, dark matter and dark energy, and the near-coincidence of their densities in our own epoch, are not numbers the universe happened to draw; they are two faces of one uniqueness statement on a single geometry — the statement that, in the companion paper, makes every electron identical to every other. This is not the claim that the cosmos was designed for us, nor that its history was fixed in every detail. It is the more restrained and, we think, more striking claim that *given one geometry*, a universe with this inventory is a requirement rather than an accident, and that the observers who measure it belong to the same structure that produced it: the layered form is forced, the history within the layers is not — freedom sheltered, not preprogrammed. The two oldest questions, why these particles and why this universe, turn out to be one question asked twice; and a universe that can be derived is a universe that can be understood, which is the only promise physics has ever made. Where physics ends and interpretation begins, the paper marks the line and leaves the final step to the reader.

Keywords: GG-Theory program; cosmological-boundary closure; π-locked inventory ($\Omega_m = 1/\pi$); corridor-switch redshift $z_* = e^7$; one derived Hubble background; Kaxion dark matter; baryogenesis and BBN; primordial gravitational waves; MSMF/OGN cost-minimum principle; codebook provenance triage; second horizon problem; universal-cosmic-identity principle.

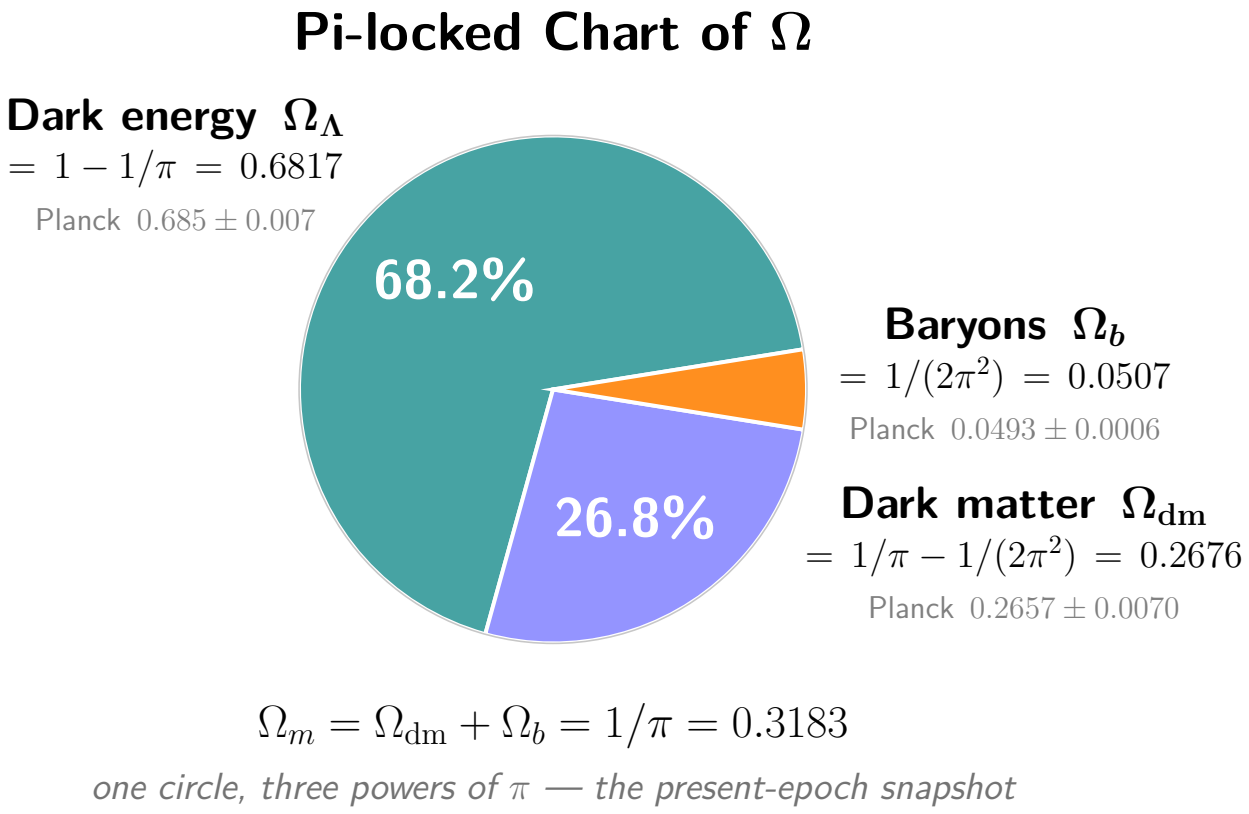

Figure 1: Pi-locked Chart of Ω — the composition of the universe in one circle. Every slice of the present-epoch cosmic energy budget is a single closed-form fraction of π, the three forced together as one coupled set: dark energy $\Omega_\Lambda = 1 - 1/\pi$, dark matter $\Omega_{\rm dm} = 1/\pi - 1/(2\pi^2)$, and baryons $\Omega_b = 1/(2\pi^2)$, with the two matter slices summing to $\Omega_m = 1/\pi$. Written just outside each slice is the prediction and, beneath it, the Planck 2018 (Planck Collaboration, 2018) measured value; the two agree to better than 3%. The coupled π-form is the headline structural result, derived (conditional on Conjecture VIII) as Theorem 74.1 in Part X; the matter-composition ratio $1/(2\pi)$ and the vacuum scale hold at every epoch, while these absolute present-day fractions are the one snapshot calibrated to the sky (§32.11).

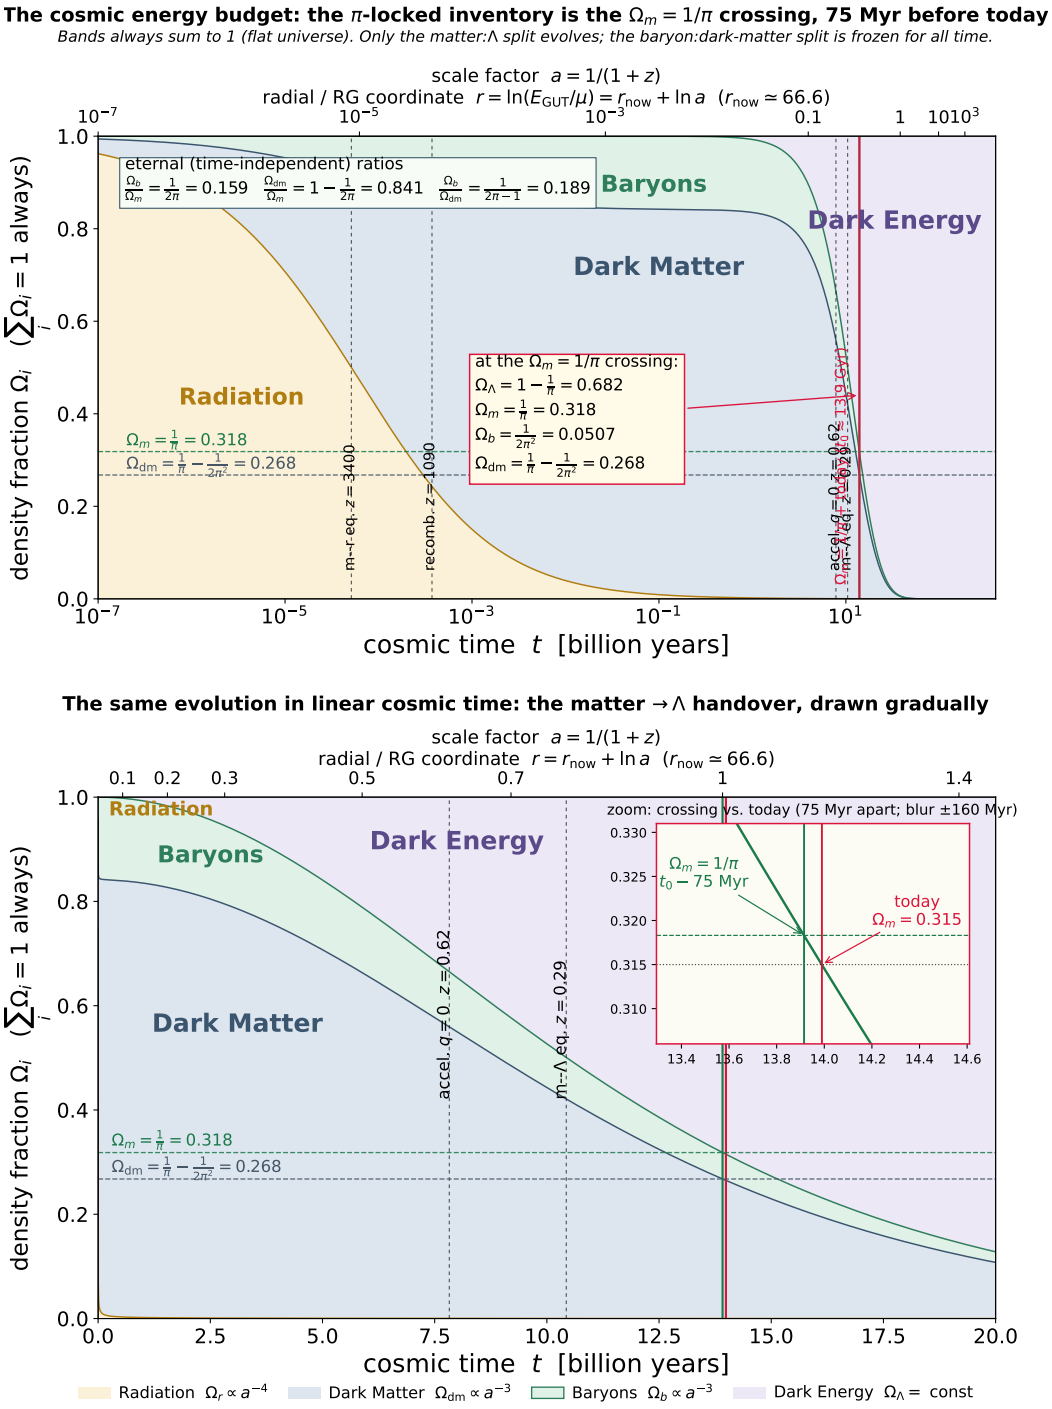

Figure 2: The cosmic energy budget through history, shown twice — the story of this paper in one image. Only the matter:vacuum seam moves; the π-locked inventory is the snapshot at the $\Omega_m = 1/\pi$ crossing. Derived, and defended, in Part IV (Figure 22).

The Entire History of the Universe, Computed — One Two-Basin Descent, from the GUT Edge to the de Sitter Future

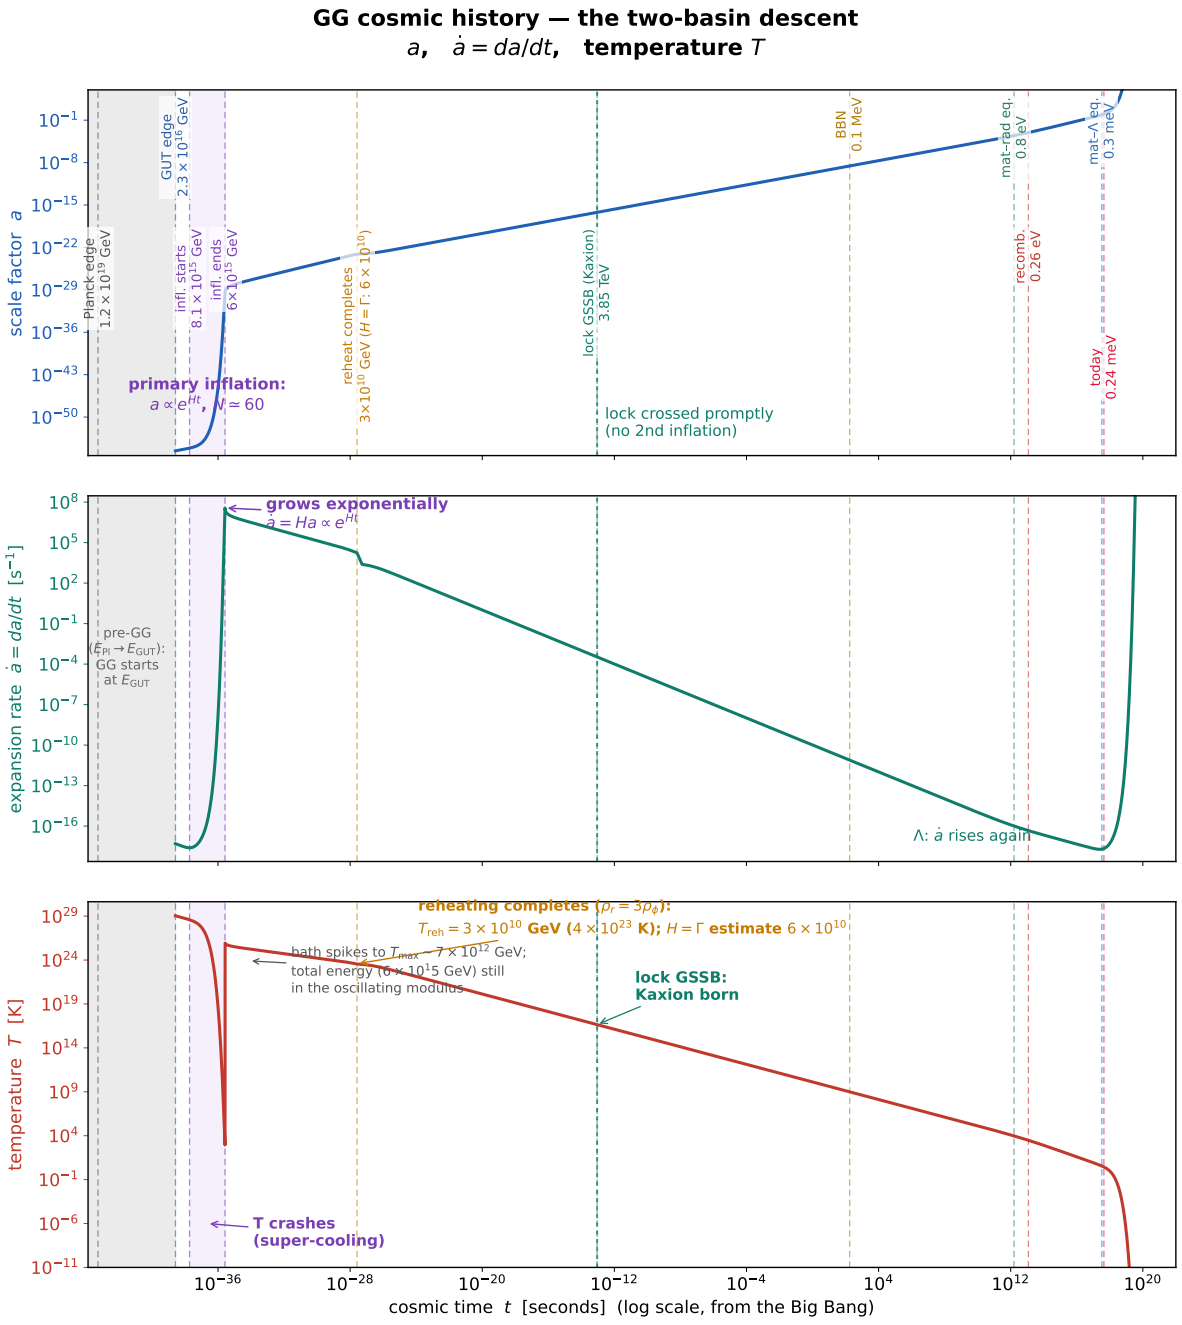

Figure 3: The two-basin descent, numerically integrated (Friedmann + inflaton + reheating fluid, smooth by construction): the scale factor a, the expansion rate $\dot{a}$, and the temperature T from the Planck edge to 10^{22} s, every vertical epoch line carrying its energy scale, the two de Sitter epochs bracketing the whole history. Derived, and fully annotated, in Part IV (Figure 19).

The Same History in the Descent's Own Variables — Scale Dimension, Energy Scale, and the Cosmic Energy Budget

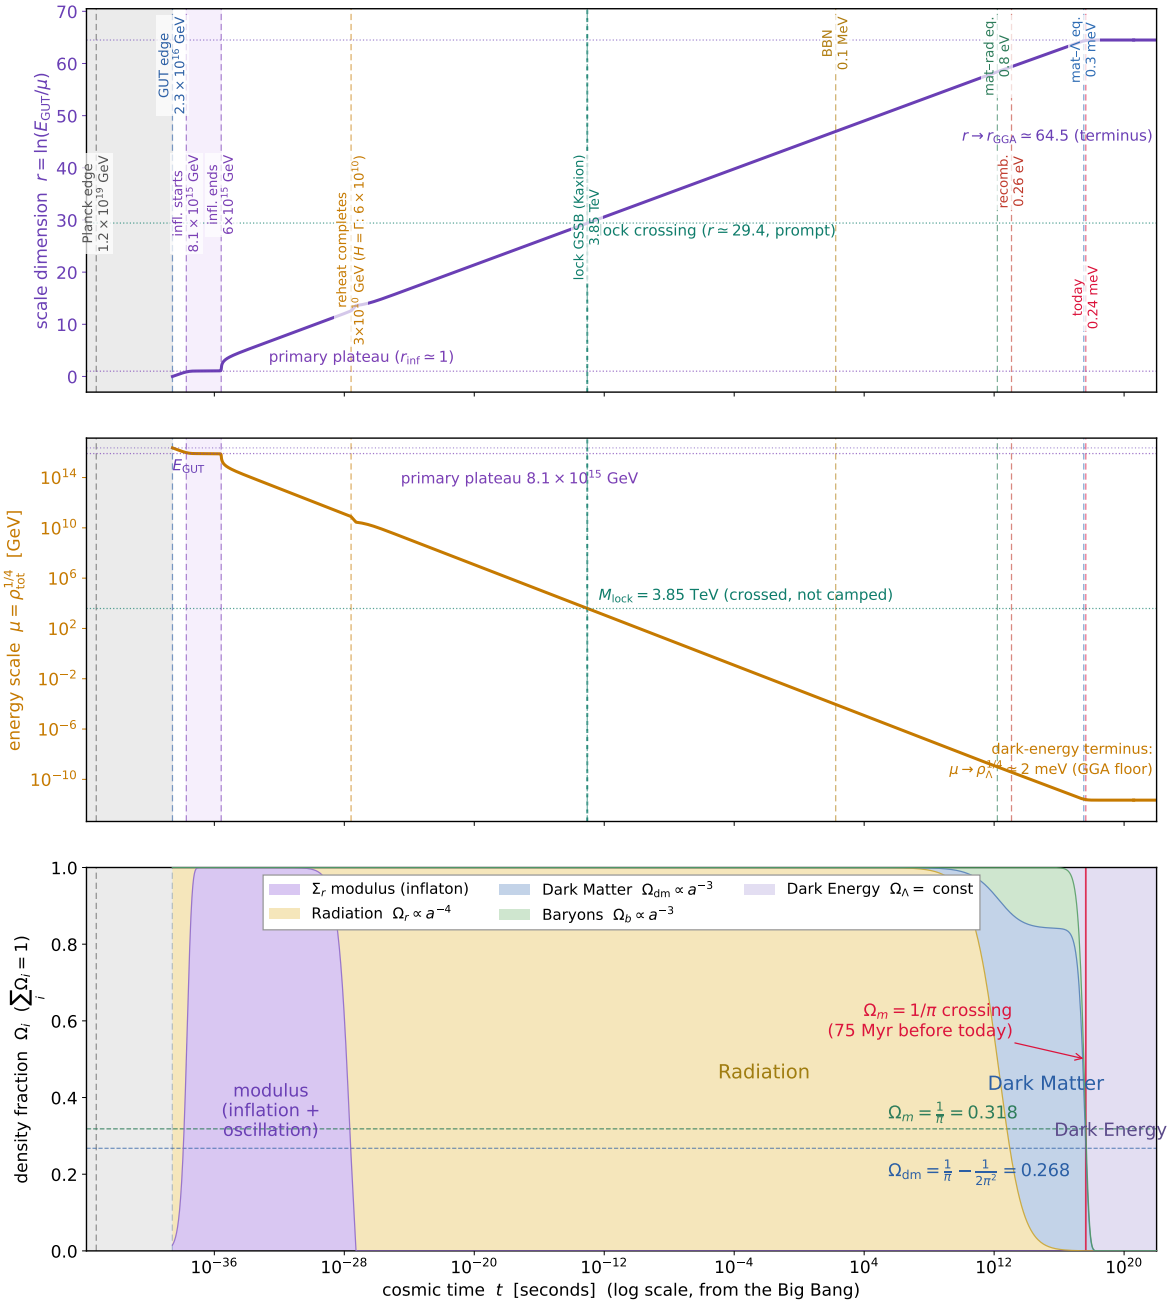

Figure 4: The descent's own variables: the scale dimension $r = \ln(E_{\rm GUT}/\mu)$, descending monotonically from the primary plateau to the terminus; the energy scale $\mu = \rho_{\rm tot}^{1/4}$, from 8.1×10^{15} GeV to the ≈ 2 meV dark-energy floor; and the cosmic energy budget $\Omega_i(t)$, with the $\Omega_m = 1/\pi$ crossing marked in red. Fully annotated in Part IV (Figure 20).

Figure 5: The cosmic GRIP descent, drawn to scale on $r = \ln(E_{\mathrm{GUT}}/\mu)$: one open-system descent from the GUT edge ($r = 0$) to the de Sitter future, read simultaneously in energy μ, time t, and the scale factor a. The key epochs are drawn to scale and colored; the late universe is magnified below the break so that recombination, structure, and the present epoch stay legible. Every value is the standard cosmological reading of the master chronology (Table 36); nothing on the axis is fitted.

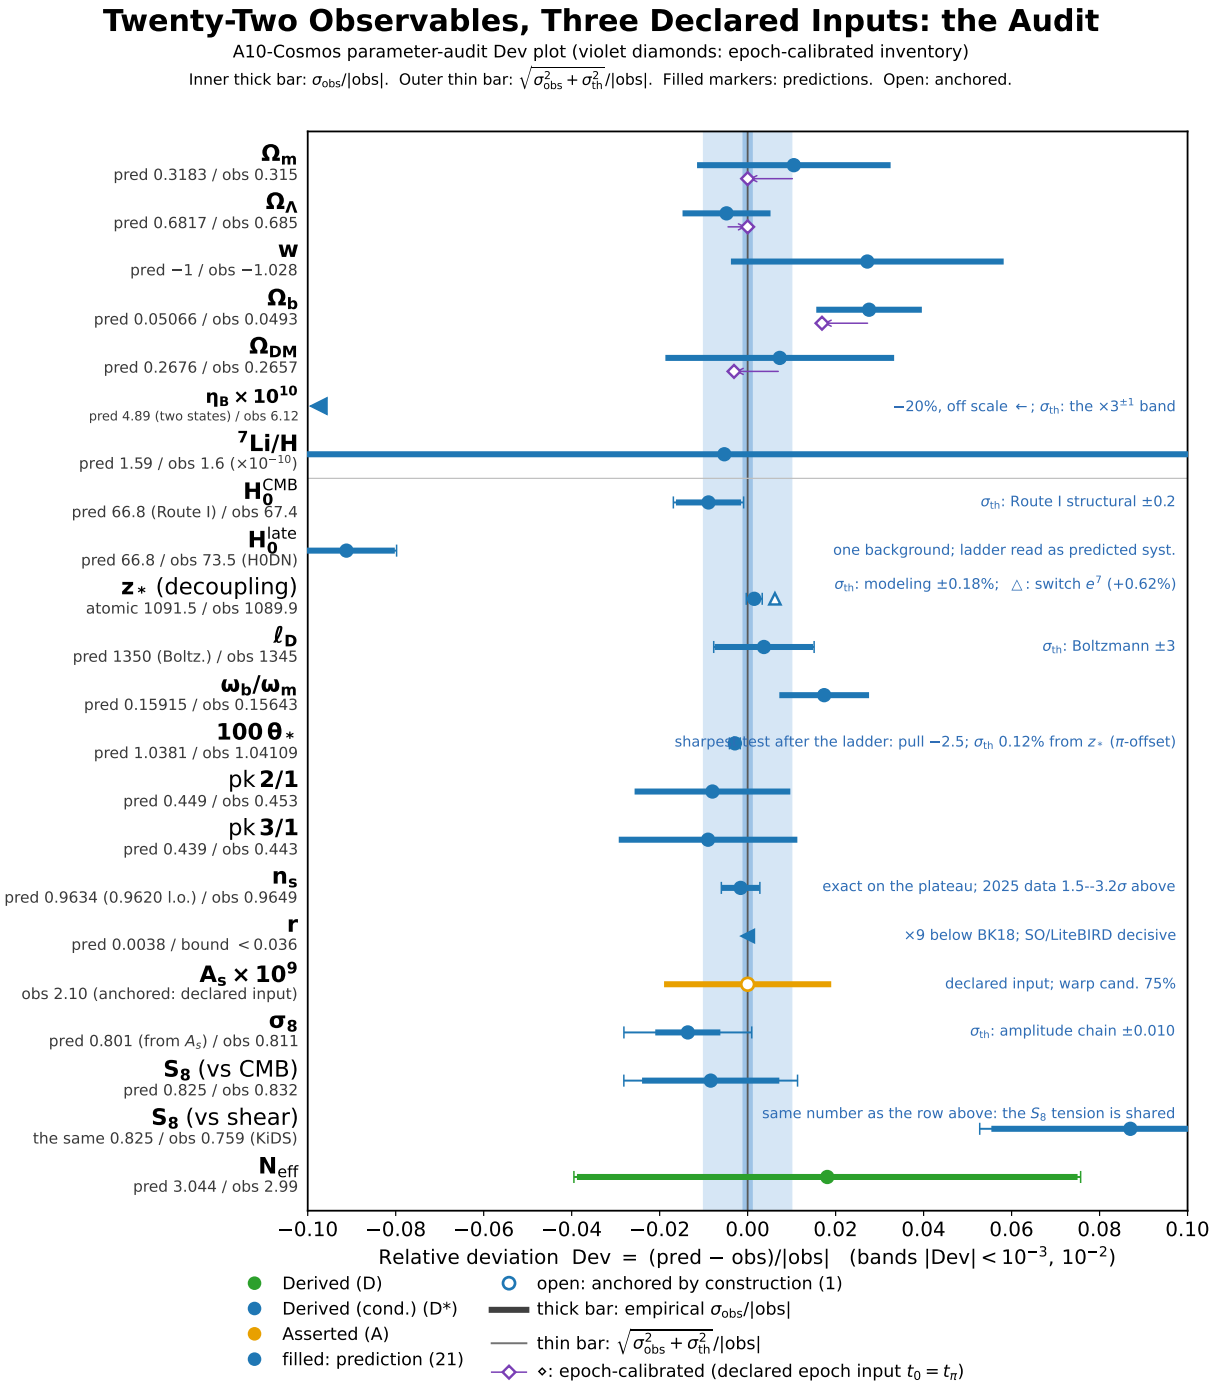

Figure 6: The bottom line, in one chart: the fractional deviation (pred − obs)/|obs| for the twenty-two rows carrying an empirical comparison — twenty-one predictions (filled), one anchored by construction (open) — with nested precision bands at $|\mathrm{Dev}| < 10^{-3}, 10^{-2}$. The z_* entry scores the framework's atomic-chain decoupling prediction against the measured value, the corridor switch e^7 annotated as a distinct topological marker (§42.1). Presented with its data table in Part I (Figure 10) and reproduced, audited row by row, as Figure 31 in the codebook of Part IX.

A Six-Dimensional Bulk Whose Fifth Dimension Is Scale and Whose Sixth Is Phase

The whole program in one diagram, with the scale axis carrying numbers. The bulk $\mathcal{M}_4 \times I_r \times S^1_\theta$ is a fiber bundle: over every point of the Lorentzian stage sit the two *non-spatial* fibers, scale and phase. AX-2 (GGEP) acts *upward* and builds them; AX-3 (OGGEP) acts *downward*, and registration is the projection that erases (r, θ) — the origin of quantum behavior. The vertical axis is calibrated: $E_{\rm GUT} = E_{\rm Pl} e^{-2\pi}$ at $r = 0$, the OGN-5 depth $r = K = 35$ at the infrared boundary, and the lock at $r = 29.41$, which is $K - 2\pi + \ln 2$ exactly.

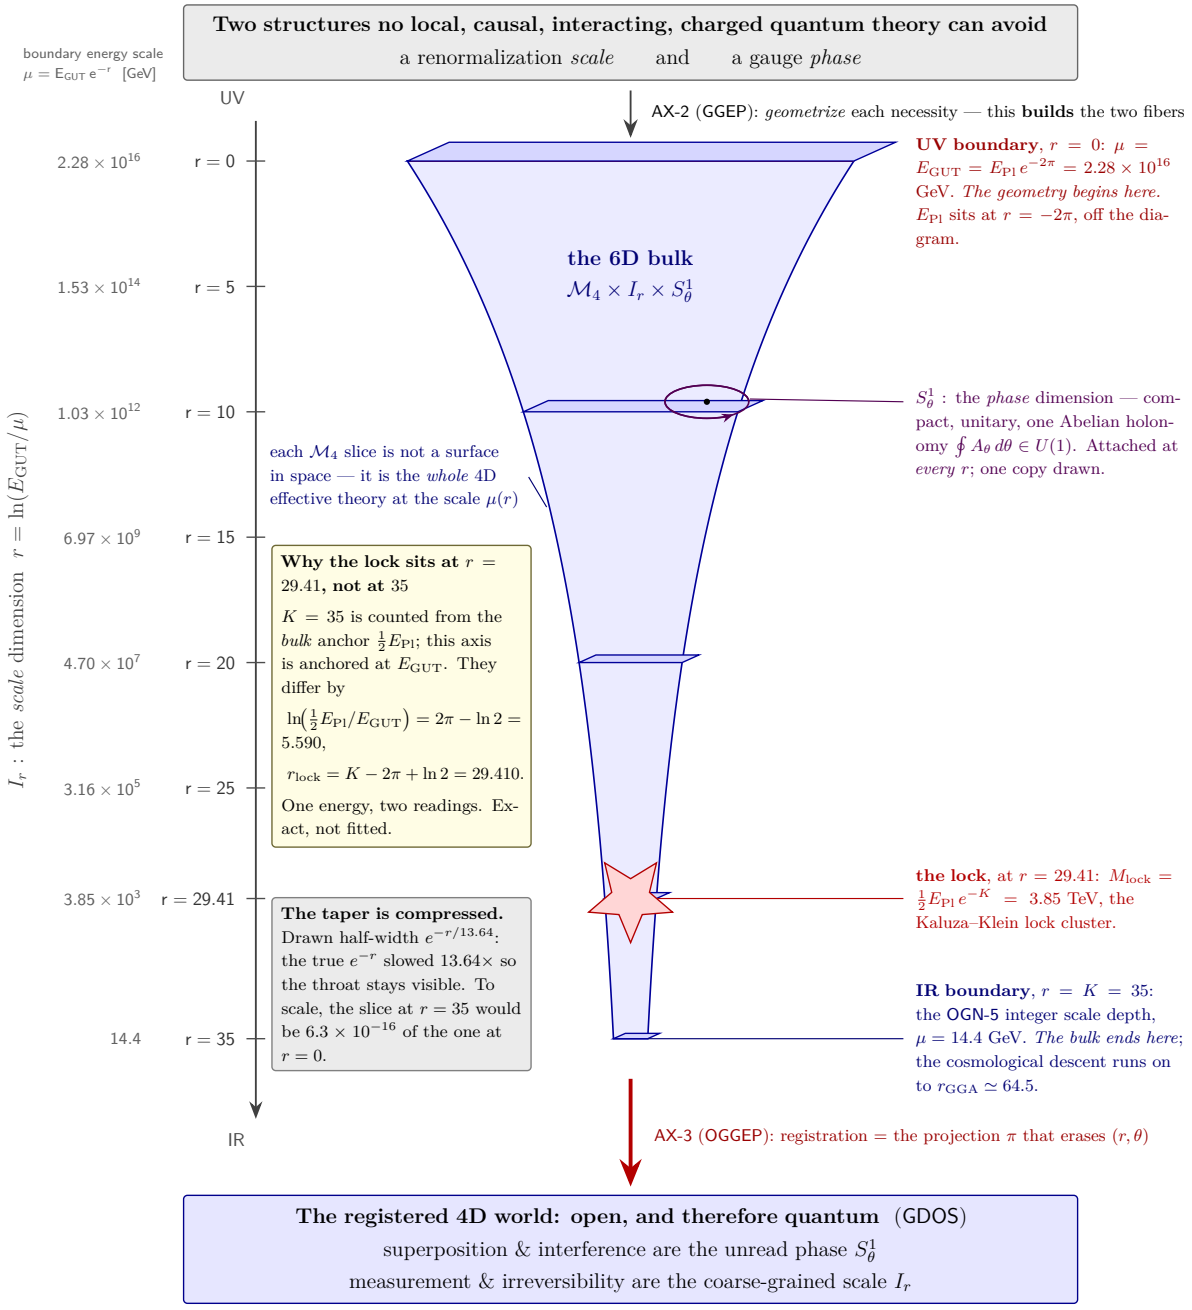

What is actually claimed, in one page

The input. One six-dimensional geometry, inherited unchanged from the particle-physics companion GG-A9-SM, with zero continuous parameters adjusted to the observables it predicts, plus the proved OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, fixed upstream in the particle sector. Three inputs are declared, each named where it is used: the primordial amplitude A_s, the timing of the present epoch, and the bulk Green–Schwarz normalization $f_{\rm bulk}$, calibrated by the dark-matter abundance in GG-C1-Kaxion. Said plainly, that calibration means the dark-matter row $\Omega_{\rm DM}$ is the row $f_{\rm bulk}$ was fixed on, and is a consistency statement rather than a prediction, while the Kaxion mass and decay constant that descend from the same scale are predictions conditional on that one calibration (row D14 is declared, not derived).

The output. The present-epoch inventory in closed form — $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $w = -1$ exactly — plus the corridor-switch redshift $z_* = e^7$ (within 0.6% of the measured decoupling), the background $H_0 = 66.8 \pm 0.2$ km/s/Mpc (66.6 on the two-road reading), the tilt $n_s = 0.9634$ exact ($1 - 2/N^* = 0.9620$ at leading order), the tensor ratio $r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order), and a dark-matter particle at 8.1 μeV: twenty-one codebook entries (twenty-two comparisons in the Dev audit: three unmeasured rows out, four acoustic images of the same fractions in).

The status, exactly. One entry Derived (a null); nineteen Derived-conditional, each resting on a named, falsifiable conjecture (ten in all, and the dark-matter row on two: its completeness lemma and the inflation-era pinning C-Kaxion-d, which the isocurvature bound makes necessary); one Asserted (A_s). Nothing is tuned after the fact; the correlated +2.7% cluster is adopted as a declared second-order precision floor.

What is not claimed. Not a likelihood-level fit (the zero-parameter TT curve gives $\chi^2/\text{dof} \approx 2.3$ — the 0.3% acoustic-angle deficit, reported, not negotiated); not the measured recombination redshift (z_* is a topological landmark); not a settled resolution of the Hubble tension (one background is derived and the ladder is read as the predicted outlier; the joint Einstein–Boltzmann fit is named as open); and not the closure of the Hopf-fiber lemma VIII-b″, which is owed, not delivered.

How it dies. A gravitational-wave siren reading of H_0 at the ladder's 73 rather than ≈ 67; a tensor measurement inconsistent with $r = 0.0038$ beyond the launch's spread 0.0037–0.0040, in either direction; a tilt measurement excluding the derived $n_s = 0.9634$ in either direction at three standard deviations — a criterion already met on one of the two 2025 combinations: the CMB-plus-BAO value sits 3.2σ above it, the CMB-only value 1.5σ, and the paper records the tilt as its one currently failing test (§28.10); w driven off -1 at Euclid/Rubin precision; the Kaxion absent at DFSZ depth across 1.59–2.13 GHz. Each is two-sided, dated, and printed in the falsification ledger: Table 50 states the full set in "how the theory can lose" language, and Table 51 attaches to each one the instrument class and the decision horizon.

Why it is not numerology. The integer ledger is pinned upstream, not searched (unique among the ten closed OGN-5 ledgers under $N = 3$, $L = 5$); the look-elsewhere exposure is answered by a rule, not a number: every form is assigned by the stratum it inhabits, stated before the value is compared (§29.8 is the worked example, with its enumeration released). An earlier version quoted a global trials factor $p \sim 10^{-3}$–10^{-4} and promised a Monte-Carlo to back it; the Monte-Carlo has since been run, and on an *unrestricted* alphabet it returns order unity, so the bound is withdrawn and the discipline stated in its place.

Overview of Contents

Part I. Introduction 33
Part II. The inherited framework: GG-Theory in one Part 59
Part III. The cosmological stage 73
Part IV. Inflation and dark energy: the cosmic GRIP descent 97
Part V. Origin of matter: Kaxion dark matter, baryogenesis, and BBN 159
Part VI. Growth of structure and the present-day tensions 197
Part VII. Black holes and primordial relics 248
Part VIII. Evolutionary chronology of the universe 280
Part IX. Cosmological codebook: Derived / Derived (cond.) / Asserted 315
Part X. Topological completion: *why the universe has the inventory it does* 346
Part XI. The assembled ΛCDM: the standard cosmological model as a constitutional export 376
Part XII. Summary and cross-paper coherence 382
Part XIII. Discussions 401
Part XIV. Conclusions: the summit view 423
Part XV. Technical Appendices 433
Part XVI. Common Objections and Responses 622
Part XVII. References 667

List of Figures

1 Pi-locked Chart of Ω: the cosmic inventory as three powers of π (front-matter showcase). . . . 5
2 The cosmic energy budget through history (preview; derived in Part IV). 6
3 Cosmic history, numerically integrated (showcase; full version in Part IV). 7
4 Cosmic history in descent variables (showcase; full version in Part IV). 8
5 The cosmic GRIP descent on the radial axis (showcase; master chronology in Part VIII). . . . 9
6 The parameter-audit Dev plot (showcase; audited row by row in Part IX). 10
7 A century of cosmology, from a static universe to precision ΛCDM. 34
8 Why ΛCDM is believed: four independent probes, one six-parameter model. 36
9 The central claim of GG-Theory, as one deductive funnel. 44
10 Preview: the parameter-audit Dev plot (full version in Part IX). 49
11 The twenty-two-paper cohort map. 55
12 Roadmap of the seventeen Parts of GG-A10-Cosmos. 57
13 6D bulk and its OGGEP projection to the 4D cosmological boundary. 61
14 Zero-parameter derivation tree of GG-A10-Cosmos. 66
15 The GG-Theory derivation ladder, from the constitution to ΛCDM. 74
16 The beginning: the Planck-scale start versus the GG-6 GUT-edge start. 76
17 The whole history of time, on one line. 94
18 The derived launch: potential and predictions. 114
19 Cosmic history, numerically integrated: scale factor, expansion rate, temperature. 115
20 Cosmic history in descent variables: scale dimension, energy scale, energy budget. 116
21 The trio on the two fibers: three dark episodes and their joint terminus 141
22 Stacked-band evolution of the cosmic energy budget (log and linear time). 143
23 The matter–vacuum ratio versus the radial coordinate. 145
24 The H_0 consistency web: independent CMB roads converge on ≈ 67. 202
25 The microwave sky from the derived parameters. 205
26 The derived Hubble parameter $H(t)$ of the GG-6 background, of which H_0 is one point. . . . 223
27 One background, three ways of reading it. 225
28 Cosmic-history timeline of GG-A10-Cosmos: aligned t, T, z axes. 282

29 The cosmic GRIP descent: every key epoch on the radial axis. . . . 299
30 The radial ladder of cosmic history. . . . 308
31 Parameter-audit deviation plot for the cosmological codebook. . . . 323
32 Provenance map of the cosmological codebook. . . . 324
33 The Hopf fibration behind the π-locked inventory. . . . 363
34 Pi-locked Chart of Ω: the cosmic inventory as three powers of π. . . . 364
35 The twenty-two-paper cohort map of GG-Theory (Summary recap). . . . 387
36 Falsification / kill-test timeline. . . . 420
37 Hierarchy of program-defined acronyms in the GG-Theory program. . . . 439
38 Comoving Hubble radius $(aH)^{-1}$ across cosmic history. . . . 446
39 Slow-roll inflaton potential and the GG normalization gap . . . 484
40 BBN abundances vs the baryon-to-photon ratio . . . 500
41 Linear growth factor $D(a)$ and structure freeze-out. . . . 526
42 CMB acoustic peaks and the computed damping scale. . . . 550
43 The retired two-stage corridor (record). . . . 568
44 Kaxion dark matter from vacuum misalignment . . . 586

List of Tables

1 The open problems of contemporary cosmology (2026), in five clusters . . . 38
2 Preview: the numerical bottom line of GG-A10-Cosmos. . . . 48
3 The twenty-two-paper GG-Theory cohort and its relevance to this cosmology. . . . 51
4 Map from the seventeen open problems to the inherited cohort components . . . 71
5 The four scales that organize the cosmological history . . . 78
6 LCDM as a theorem: the six concordance parameters and their closed forms . . . 85
7 Background scaling laws used in Part II . . . 91
8 Ordinary high-scale inflation against the minimal GG module . . . 112
9 The stratum cluster for the vacuum ratio . . . 128
10 The minimal late-time GG inventory . . . 136
11 The three dark episodes on the two fibers . . . 140
12 The cosmic energy budget through history in the pi-locked inventory . . . 147
13 Inherited particle-sector inputs from GG-A9-SM . . . 160
14 The single physical Kaxion benchmark, as GG-C1-Kaxion states it . . . 164
15 Signal logic of the Kaxion window . . . 171
16 The resonant-leptogenesis chain used in Part V . . . 182
17 Light-element audit used in Part V . . . 188
18 Bridge from baryogenesis language to the late-time inventory . . . 191
19 Chronological ordering of the Part V mechanisms . . . 193
20 Compact parameter ledger for the Part V matter sector . . . 194
21 Compressed summary of the Part V outputs . . . 195
22 Headline falsifiable numbers of Part VI. . . . 199
23 Core CMB observables forced by the π-locked chart (CLASS-verified). . . . 203
24 Inherited background inventory for the CMB and growth audit . . . 212
25 Condensed CMB and primordial-sector audit . . . 216
26 Status of each element of the Hubble reading . . . 227
27 One-background decomposition of the Hubble reading . . . 228
28 Sign structure of the Part VI growth channels . . . 233
29 Mechanism ladder for the JWST-era early-galaxy problem . . . 242
30 What counts as a resolution of the information paradox . . . 258
31 The island/QES program beside the GG open-system reformulation . . . 259
32 Compact observational dictionary for the black-hole block . . . 262
33 Illustrative PBH mass windows relevant to GG-A10-Cosmos . . . 271
34 Integrated summary of Part VII . . . 277
35 Coordinate dictionary used throughout Part VIII . . . 286
36 Master chronology of the universe in GG-A10-Cosmos . . . 300
38 Radial bookkeeping for the main segments of cosmic history . . . 310
39 The four milestones and the present epoch, as a five-clock ladder . . . 312
40 Cosmological codebook: twenty-one observables by provenance . . . 317
41 Data table for the parameter-audit Dev plot. . . . 322

42 Closure economy: how the codebook evolves as each lemma is proved . . . 338
43 The electron, the Kaxion, and dark energy as one GDOS + GRIP construction. . . . 349
44 The silent structural axioms behind textbook LCDM, and their replacements . . . 377
45 Compressed codebook: twenty-one observables and their status . . . 389
46 Compressed audit roadmap: the ten named conjectures and what each closes . . . 390
47 GG-distinct comparison: free parameters and closure mode . . . 398
48 How a few inherited modules address several problems at once . . . 404
49 GG cosmology against several leading theoretical habits . . . 406
50 Practical falsification ledger for the cosmological branch . . . 417
51 The falsification decade: each kill criterion, instrument and horizon . . . 419
52 A practical forward program suggested by GG-A10-Cosmos . . . 422
53 Summary of abbreviations used in the present paper . . . 435
54 Glossary of key symbols used in the present paper . . . 440
55 Compact codebook glossary with predicted and empirical values . . . 443
56 The regulator landscape at the lock crossing . . . 516
57 Validation of the Boltzmann integration . . . 521
58 Record: kill conditions for the retired Part VI growth block . . . 543
59 Record: sign structure of the retired Part VI growth corrections . . . 572
60 The pi-locked fractions against Planck . . . 582
61 Internal source ledger for GG-A10-Cosmos . . . 614
62 Compact cross-paper inheritance: what A10 draws from each paper . . . 615
63 Consolidated benchmark audit for GG-A10-Cosmos . . . 615
64 Named-statement deck of GG-A10-Cosmos . . . 617
65 Named-lemma promotion roadmap for the MSMF/OGN audit . . . 619
66 Falsification-test catalog for GG-A10-Cosmos . . . 621

Contents

Executive Summary . . . 2
List of Figures . . . 13
List of Tables . . . 15
Preface . . . 30

Part I. Introduction **33**
1 The oldest question, asked precisely . . . 33
2 A century of cosmology: how the Standard Model of the universe was built . . . 34
3 ΛCDM: the concordance model, and why it works . . . 35
4 What the standard model of cosmology measures but cannot explain . . . 37
4.1 The seventeen open problems, one by one . . . 39
5 A radically new idea: the universe as the boundary of a geometry . . . 43
6 The central claim: the cosmic inventory consists of open-system observables . . . 44
7 What this paper claims . . . 46
8 Scope, the cohort, and how to read this paper . . . 50
9 The cohort inheritance and the roadmap of the seventeen Parts . . . 51
9.1 The cohort inheritance . . . 51
9.2 Roadmap of the seventeen Parts . . . 54
10 Part I in plain English: what the standard model of the universe does and does not explain . . . 58

Part II. The inherited framework: GG-Theory in one Part **59**
11 One geometry, no dials: the six-dimensional bulk . . . 59
12 The boundary projection: how four-dimensional physics arises . . . 60
12.1 Which particles? The observer-indexed vacuum, and horizon thermality as openness . . . 62
13 Zero free parameters: BI-6 and the integer ledger . . . 63
14 The runtime: how cosmic branches are selected . . . 64
14.1 Why GRIP and Wilson-loop GSSB now control the cosmology paper . . . 67
15 What this buys cosmology . . . 70
16 Part II in plain English: what this paper borrows, and at what strength . . . 72

Part III. The cosmological stage **73**
17 The thesis: ΛCDM is a theorem of GG-Theory . . . 73
18 Why the universe began at the GUT edge, not the Planck scale . . . 74
19 The scale chain: $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock}$ . . . 76
19.1 The inherited scales . . . 77
19.2 What each inherited scale actually controls . . . 77
20 The inherited foundation: the first three rungs . . . 78
20.1 Rung 1 — the constitution fixes what may count as a cosmology (GG-A1-Constitution) . . . 79
20.2 Rung 2 — the bulk and its anomaly lock (GG-A2-GG6, GG-A3-BI6) . . . 79
20.3 Rung 3 — zero free parameters (GG-A4-Euclid) . . . 79
20.4 Rung 3, restated for cosmology . . . 79
21 General relativity as a theorem of GG-6/GDOS . . . 80
21.1 The starting point: curvature in six dimensions . . . 80
21.2 Integrating out the throat: the scale-profile reduction . . . 80
21.3 Reading off the four-dimensional theory . . . 81

21.4 Conservation, the Newton constant, and the vacuum scale . . . 81
22 The Friedmann equations as a theorem . . . 82
22.1 The maximally symmetric metric . . . 82
22.2 The curvature of an expanding universe . . . 83
22.3 The source: a perfect fluid from the inherited ledger . . . 83
22.4 The Friedmann equations . . . 83
22.5 What is a theorem, and what is inherited . . . 84
23 ΛCDM as a theorem: the six parameters . . . 84
24 Expansion, cooling, and horizons . . . 86
24.1 Exact scaling law for a constant equation-of-state component . . . 86
24.2 Cooling law from comoving entropy conservation . . . 87
24.3 Why the expansion cools every sector differently . . . 88
24.4 Why the universe expands at all . . . 88
24.5 The comoving Hubble radius and the meaning of horizon growth . . . 89
24.6 Particle horizon, event horizon, and causal size . . . 89
24.7 Freeze-out, decoupling, and the universal competition Γ/H . . . 90
24.8 A compact table of expansion and cooling laws . . . 91
25 The radial chronology and the arrow of time . . . 91
25.1 Reinterpreting the source sectors in GG-6/GDOS . . . 91
25.2 Expansion, cooling, and the arrow of time in the GDOS language . . . 92
25.3 The fiber without holonomy: why the chronology cannot run backward . . . 93
25.4 Bottom line: the stage is set . . . 95
26 Part III in plain English: why the universe starts at the GUT edge and not at the Planck scale . . . 96

Part IV. Inflation and dark energy: the cosmic GRIP descent 97
27 One descent, from the electron to the cosmos . . . 97
28 The launch: inflation as the opening of the descent . . . 98
28.1 Why the universe needs a launch . . . 98
28.2 Why a free inflaton is forbidden by the inherited program . . . 98
28.3 The minimal effective action for the inflationary modulus . . . 99
28.4 The two basins: a GUT-adjacent plateau, and a lock that is crossed rather than camped . . . 100
28.5 The first e-fold: the cold start and the forced throw . . . 102
28.6 The entrance: what the geometry must still supply . . . 102
28.7 Slow-roll equations and the meaning of "slow-smooth" . . . 105
28.8 Scalar tilt, amplitude, and the tensor target: two predictions and one input 106
28.9 How the descent does the cosmological work . . . 107
28.10 The tilt and the tensor ratio evaluated exactly on the plateau . . . 109
28.11 Graceful exit as a geometric spontaneous symmetry breaking . . . 109
28.12 Reheating and the return to the hot Big Bang . . . 111
28.13 A compact comparison and the bottom line of the launch . . . 112
28.14 The launch, graded . . . 113
28.15 Inflation on the two fibers: the interval's coordinate, rolling once . . . 117
29 The terminus: dark energy as the Ground Geometric Attractor . . . 119
29.1 What has to be explained at the end of the descent . . . 119
29.2 Residual vacuum energy as a GFF/CEP boundary functional . . . 120
29.3 A natural meV scale from the lock-scale seesaw . . . 120
29.4 Why the equation of state is exactly $w = -1$: the transition is complete . . . 121
29.5 From flatness to the late-time density fractions . . . 124
29.6 The cosmic coincidence as an attractor-slice . . . 125
29.7 From the residual vacuum density to the present Hubble scale . . . 126

29.8 The derivation of H_0: the two-route anchor of the present expansion rate . 127
29.9 The de Sitter future and the entropy of the terminus 133
29.10 Dark energy on the two fibers: the circle's share, the interval's scale 134
29.11 Bottom line of the terminus 135
30 Two crossings, three legacies 136
31 The same principle, from the electron to dark energy: the vacuum generations . . . 137
31.1 The vacuum generations 138
31.2 The cosmological triptych: rolling, ringing, rest 139
32 The cosmic coincidence, bounded: why the inventory is π-locked *now* 139
32.1 The exact background and the flatness sum rule 142
32.2 What is eternal and what is a snapshot: the decomposition 142
32.3 The full evolution and the four transition epochs 146
32.4 The true age of the universe 148
32.5 The future: the de Sitter fate and the exact $a(t)$ 149
32.6 The coincidence quantified: the radial observation window 150
32.7 Physics intuition: why these magnitudes, and why this window 150
32.8 The growth-suppression integral: structure formation is freezing out now . 151
32.9 The anthropic weight, quantified 152
32.10 The resolution in three movements 152
32.11 What is predicted, and what is not: form, scale, and epoch 153
32.12 The honest residue: GG-Theory fixes the form and scale, not yet the epoch 154
32.13 Falsifiable corollaries 155
33 Falsification and bottom line 156
33.1 The significance: two unknowns, one hierarchy 156
34 Part IV in plain English: the two ends of a single descent 158

Part V. Origin of matter: Kaxion dark matter, baryogenesis, and BBN 159
35 Kaxion dark matter 161
35.1 Historical background 161
35.2 Why dark matter should be sought first in the compact-angle sector 161
35.3 From BI-6 and the compact angle to an axion-like 4D field 162
35.4 Mass and decay-constant benchmark: one calibration, then a ladder 164
35.5 Relic abundance from misalignment and the π-locked inventory 165
35.5.1 Oscillation onset, free streaming, and why the Kaxion behaves as cold dark matter 166
35.5.2 The history of the angle: released at the lock crossing, single-patch by conjecture 167
35.5.3 Why the Kaxion is not merely a rebranded Peccei–Quinn axion 169
35.5.4 What the direct-search program would actually look for 169
35.6 Couplings, lifetime, and the single-line haloscope signature 170
35.7 Why the Kaxion is the first dark-matter candidate of GG-A10-Cosmos . . . 172
35.8 The Kaxion on the two fibers: the anomaly ledger's own phase, ringing still 172
36 Baryogenesis via the inherited neutrino and CP ledger 173
36.1 Historical background 173
36.2 Why the matter–antimatter problem becomes sharper after GG-A9-SM . . 174
36.3 Sakharov conditions in GG-6/GDOS language 174
36.4 The inherited light-neutrino block and the heavy completion 175
36.5 Seesaw matching: the geometry fixes the Yukawas 177
36.6 Resonant CP asymmetry from the inherited phases 177
36.7 Washout and the strong-washout efficiency 178
36.8 Assembling η_B: the mechanism behind $e^{-K/2}$ 179
36.9 Boltzmann transport and why the geometry compresses the calculation . . 180

36.9.1 A numerical audit table for the baryogenesis chain . . . 182
36.10 From baryon asymmetry to the present baryon density fraction . . . 182
36.11 Why exact proton stability and successful baryogenesis are compatible . . 183
36.12 What the baryogenesis block predicts and how it can fail . . . 183
36.13 Exact neutrality: why the books balance and the electrons match the protons 184
37 Big Bang nucleosynthesis and the lithium-7 audit . . . 185
37.1 Historical background . . . 185
37.2 Why BBN is the first precision audit of Part V . . . 185
37.3 Standard BBN baseline: freeze-out, bottleneck, and synthesis . . . 185
37.3.1 Freeze-out and bottleneck in explicit numbers . . . 186
37.3.2 Why deuterium is such a powerful baryometer . . . 186
37.4 The lithium-7 problem and the A=7 channel . . . 186
37.5 The A=7 port and the $1/\pi$ prediction . . . 187
37.6 Why deuterium and helium are preserved . . . 188
37.7 What the lithium audit predicts and how it can fail . . . 189
37.8 A consistency chain from η_B to Ω_b and the baryon-to-dark-matter ratio . . 189
37.9 Chronological ordering inside Part V . . . 191
37.10 A compact parameter ledger for the matter sector . . . 193
37.11 Bottom line of Part V . . . 194
38 Part V in plain English: where the matter came from, and what is assumed . . . 195

Part VI. Growth of structure and the present-day tensions 197
39 One switch, two faces: the geometric mechanism behind the present-day tensions . 197
40 Determinism from geometry: the microwave background as an over-determined test 199
40.1 The master identity: fractions, densities, and several roads to one rate . . 200
40.2 The core observables, computed and confronted . . . 201
40.3 The full spectrum from the derived parameters . . . 202
40.4 The π-offset: one named systematic, visible everywhere . . . 204
40.5 What is claimed, and what is not . . . 206
41 CMB baseline, primordial perturbations, and the scalar/tensor audit . . . 206
41.1 Historical background . . . 206
41.2 What this section inherits from the previous papers . . . 207
41.3 Standard perturbation baseline: what remains unchanged . . . 208
41.4 Recombination as a GDOS transition: from DGI-dominated plasma to geometric attractor . . . 209
41.5 The π-locked background inventory and baryon loading . . . 211
41.6 Acoustic ruler, peak sequence, and the fixed ratio r_s/D_A . . . 212
41.7 Silk damping, BFFT, and the ultraviolet knee of the primary CMB . . . 213
41.8 Scalar tilt, polarization, and the tensor kill corridor . . . 214
41.9 Lensing amplitude and the recombination-to-structure bridge (the early residue withdrawn) . . . 215
41.10 Bottom line of the CMB audit . . . 216
42 Hubble tension . . . 217
42.1 The switch and the first atoms: two redshifts, one descent . . . 218
42.2 Why the Hubble tension is really an inference-chain tension . . . 220
42.3 $H(t)$: the derived expansion history, of which H_0 is one point . . . 221
42.4 One Friedmann background, and how each probe reads it . . . 222
42.5 The absolute anchor and the local-ladder excess . . . 225
42.6 The distance ladder as the predicted outlier . . . 227
42.7 Success criteria and kill switches for the Hubble module . . . 229
43 The S_8/σ_8 matter-clustering tension . . . 229
43.1 Definition of the observable and why it is difficult to move cleanly . . . 230

43.2 Standard linear growth equation and the scale on which S_8 lives 230
43.3 The amplitude computed: σ_8 and S_8 from the declared amplitude on the π-locked background 231
43.4 Growth on one background: what moves S_8 and what does not 232
43.5 Why the S_8 statement does not contradict the early-galaxy statement . . . 232
43.6 A compact consistency ladder for the growth block 233
43.7 Kill criteria for the growth module 234
44 Early galaxy formation in the JWST era 234
44.1 Why the JWST problem is not just "more S_8" 234
44.2 Press–Schechter baseline and the reason tiny changes can matter enormously 235
44.3 Small-scale boost from the early-dark-energy residue — and its withdrawal 235
44.4 The baryon–dark-matter streaming velocity: a lever with no mediator . . . 238
44.5 Collapse-threshold reduction and the high-mass halo tail 239
44.6 Direct-collapse seeds and the early supermassive black-hole problem 240
44.7 The chronology shift: why $\Delta z \sim 2$–3 is enough to matter 240
44.8 Kill criteria for the JWST module — and the verdict, now that they have been run 242
45 Integrated summary of the growth-and-tension block 243
45.1 One ledger, four manifestations 243
45.2 What Part VI does and does not claim 244
45.3 The real significance of the open-system viewpoint 245
45.4 An invitation: revisiting ΛCDM in the light of a derived $H(t)$ 245
45.5 Bridge to the next part 246
46 Part VI in plain English: one switch, and what it does and does not fix 246

Part VII. Black holes and primordial relics 248
47 Black holes as GRIP-In way-stations: the organizing idea 248
48 The black-hole information paradox 249
48.1 Why black holes are the sharpest information-processing objects in the paper 250
48.2 Closed-system versus operational formulations 250
48.3 The black hole as a System Instance Packet 252
48.4 GFF and the correct entropy ledger 253
48.5 DGI and the factorization puzzle 254
48.6 The Page curve as a GSSB saddle switch 254
48.7 OGGEP and boundary-accessible information 255
48.8 Why the firewall is not forced in the GDOS formulation 256
48.9 What would count as a genuine resolution 257
48.10 The clean GDOS resolution in one sentence 258
48.11 How the GDOS resolution compares with the island/QES program 258
48.12 Scrambling, Hayden–Preskill logic, and operational information return . . 260
48.13 Why remnants, baby universes, and final bursts are not needed in the minimal module 261
48.14 A compact falsification dictionary for future black-hole observations 261
48.15 Why the black-hole block matters for the rest of cosmology 262
48.16 What the present framework claims, and what it does not 264
48.17 Interim conclusion of the information-paradox section 265
49 Primordial black holes 265
49.1 Historical background 265
49.2 Formation windows 266
49.3 A minimal quantitative estimate 267
49.4 Formation windows tied to the inherited thermal history 268
49.5 The integer-ledger mass ladder 268

49.6 Connection to early background dynamics . . . 269
49.7 PBHs versus Kaxion dark matter . . . 270
49.8 Observational searches . . . 270
49.9 Bottom line of the PBH section . . . 271
50 Primordial gravitational waves . . . 272
50.1 Historical background . . . 272
50.2 Tensor modes from the derived plateau . . . 272
50.3 Why the primordial signal is exactly one signal: vacuum tensors, and silence from the transitions . . . 273
50.4 B-mode polarization . . . 274
50.5 Interferometer and CMB cross-checks . . . 275
50.6 What the primordial-gravitational-wave section contributes to the full paper 275
51 Integrated summary of the black-hole and primordial-signature block . . . 276
51.1 Cross-consistency and falsifiability inside the extreme-gravity block . . . 277
51.2 Bridge to the chronology of the universe . . . 278
52 Part VII in plain English: what is resolved here, and what is only reformulated . . . 279

Part VIII. Evolutionary chronology of the universe 280
53 The chronology as the GRIP descent . . . 280
54 Why a chronology block is required . . . 280
55 The chronology dictionary: translating between t, μ, r, a, and z . . . 283
55.1 Five variables, one caveat . . . 283
55.2 Exact algebraic relations . . . 283
55.3 Radiation-era time-temperature law . . . 284
55.4 Matter-era and vacuum-era clocks . . . 285
55.5 A compact dictionary table . . . 285
56 The first microsecond: from the GUT boundary to the electroweak seam . . . 286
56.1 Why the first microsecond already includes several logically distinct stages 286
56.2 Ultraviolet initialization and the descent to the plateau . . . 287
56.3 Slow-smooth inflation on the GUT-adjacent plateau . . . 287
56.4 Graceful exit, reheating, and the two crossings that stock the sectors . . . 288
56.5 TeV leptogenesis and the electroweak corridor . . . 288
57 The first three minutes: nucleosynthesis and the fixation of ordinary matter . . . 289
57.1 From the GeV corridor to the MeV corridor . . . 289
57.2 Hadronization and the operational meaning of $U(1)_B$ . . . 289
57.3 Neutrino decoupling, freeze-out, and the first-second milestone . . . 290
57.4 The deuterium bottleneck and the ignition of BBN . . . 290
57.5 The three-minute milestone . . . 291
57.6 The lithium window inside the chronology . . . 291
57.7 What is fixed by the end of the first three minutes . . . 292
58 The first 3×10^5 years: equality, recombination, and the release of the CMB . . . 292
58.1 The long radiation-dominated bridge . . . 292
58.2 Matter-radiation equality and why it is a pivotal checkpoint . . . 292
58.3 Recombination as a GSSB event . . . 293
58.4 The recombination milestone . . . 293
58.5 The CMB as fossil record and chronological anchor . . . 293
59 The first 10^8 years: dark ages, first stars, and early galaxy formation . . . 294
59.1 Dark ages and residual baryon-dark-matter streaming . . . 294
59.2 The one-hundred-million-year milestone . . . 294
59.3 Population III stars and the end of the dark ages . . . 294
59.4 The first galaxies and the JWST corridor . . . 295
59.5 Reionization as the first large-scale feedback era . . . 295

60 From the first gigayear to the present attractor . . . 296
60.1 From young galaxies to the mature cosmic web . . . 296
60.2 Why the late universe is narrow in r . . . 296
60.3 Peak star formation and mature structure . . . 296
60.4 Matter-dark-energy equality and the end of deceleration . . . 296
60.5 The present epoch . . . 297
61 Beyond the present: the asymptotic future . . . 297
62 Master chronology table . . . 298
63 Worked conversion examples . . . 305
63.1 Example A: the one-microsecond milestone . . . 305
63.2 Example B: the deuterium bottleneck (the first three minutes) . . . 305
63.3 Example C: the recombination milestone . . . 306
63.4 Example D: the hundred-million-year milestone . . . 306
63.5 Example E: the present epoch . . . 307
63.6 Example F: matter-dark-energy equality . . . 307
64 Radial bookkeeping: how much history is packed into each segment . . . 308
64.1 The universe is front-loaded in energy and back-loaded in clock time . . . 308
64.2 A second summary table: intervals rather than milestones . . . 309
64.3 Why this bookkeeping is scientifically useful . . . 310
64.4 The chronology in one radial sentence . . . 311
65 The milestone ladder in one glance . . . 311
66 Chronological synthesis . . . 313
67 Part VIII in plain English: why a chronology block earns its place . . . 314

Part IX. Cosmological codebook: Derived / Derived (cond.) / Asserted 315
68 Codebook overview and the MSMF/OGN principle on the cosmological-boundary stratum . . . 315
68.1 The cosmological codebook . . . 315
68.2 Per-parameter MSMF/OGN audit . . . 326
68.2.1 Statement of the MSMF/OGN principle on the cosmological-boundary stratum . . . 327
68.2.2 Audit of the one Derived (D) entry, and of the tensor row it formerly held . . . 328
68.2.3 Audit of the nineteen Derived (cond.) (D*) entries — nine named lemmas . . . 329
68.2.4 Audit of the inflationary sector: two promoted entries and the one remaining Asserted (A) entry — two named lemmas . . . 336
68.2.5 Closure summary and the promotion roadmap . . . 337
68.3 The Hubble two-stage corridor: canonical theorem statement . . . 339
68.3.1 Empirical motivation: the H_0 tension as of 2026 . . . 339
68.3.2 Setup and hypotheses . . . 339
68.3.3 Statement of the theorem . . . 340
68.3.4 Proof outline . . . 342
68.3.5 Corollaries . . . 343
68.3.6 Falsifiable predictions and observational tests . . . 343
69 Part IX in plain English: how to read the codebook, and why the grades are uneven 345

Part X. Topological completion: *why the universe has the inventory it does* 346
70 Overview of the topological and deep-geometric origin Part . . . 346
71 The geometric origin of the cosmic inventory: identity, one floor up . . . 347
71.1 The ground-state refinement: from GA to GGA . . . 347

71.2 GDOS and GRIP, three objects side by side . . . 348
71.3 Three cosmic questions, finally answerable . . . 350
71.4 The same principle, layer by layer, from the electron to the cosmos . . . 351
71.5 The vacuum cascade as the cosmic relaxation channel . . . 353
72 The second horizon problem: sameness of things, not sameness of temperature . . . 354
72.1 The problem the textbooks leave unnamed . . . 354
72.2 Fifty years of near misses: what was posed, and what was not . . . 356
72.3 The first half of the answer: one mold . . . 358
72.4 The second half of the answer: the ordering . . . 358
72.5 The theorem, assembled — and its exposure . . . 359
73 The third horizon problem: sameness of choices . . . 360
74 Why the inventory is π-locked: the closed-packet derivation . . . 361
74.1 Why $\Omega_m = 1/\pi$ and not some other value? . . . 361
74.2 Why is the cosmic coincidence not a coincidence? . . . 364
74.3 The closed-packet four-form integration: $\Omega_m = 1/\pi$ as a topological theorem 365
75 Two fundamental questions, one geometric answer . . . 371
75.1 The structural symmetry of the two cosmological questions . . . 372
75.2 Dissolving the three classical alternatives . . . 372
75.3 The self-referential closure . . . 373
75.4 A careful word about determinism and freedom . . . 374
75.5 The closing punch line . . . 374
76 Part X in plain English: the deepest claim in the paper, and the conjecture it rests on 375

Part XI. The assembled ΛCDM: the standard cosmological model as a constitutional export 376
77 The model the reader was taught to fit . . . 376
78 The zeroth and the six silent axioms of ΛCDM . . . 376
79 The export chain in eight stages . . . 377
79.1 C-1 — the bulk action and the ledger . . . 378
79.2 C-2 — the boundary Einstein equation . . . 378
79.3 C-3 — the FRW symmetry and the Friedmann equations . . . 378
79.4 C-4 — the components as geometry . . . 378
79.5 C-5 — the fractions . . . 379
79.6 C-6 — the initial conditions . . . 379
79.7 C-7 — the expansion history . . . 379
79.8 C-8 — what ΛCDM does not contain . . . 379
80 The assembled model . . . 380
81 What the assembly does and does not claim . . . 380
82 Part XI in plain English: the standard model of cosmology, reassembled from the geometry . . . 381

Part XII. Summary and cross-paper coherence 382
83 Overview of the Summary Part . . . 382
84 Cross-paper coherence map: the eight upstream papers of A10-Cosmos . . . 383
84.1 What each upstream paper contributes to A10-Cosmos . . . 383
84.2 Cross-paper coherence map: outward and inward . . . 385
84.3 Compressed codebook summary . . . 388
84.4 Compressed MSMF/OGN audit roadmap . . . 390
84.5 The second-order residual: the +2.7% common offset, honestly tiered . . . 390
85 Post-hoc-fit defense: why the π-locked inventory is not numerology . . . 393

85.1 Forcing versus searching: the discrete look-elsewhere effect, and the Euclidean reading . . . 394
86 GG-distinct comparison: what makes GG-Theory cosmology structurally different . 398
87 Part XII in plain English: how this paper is held to the rest of the cohort . . . 399

Part XIII. Discussions **401**
88 What this Part does . . . 401
89 Significance and uniqueness . . . 401
89.1 The essence of GG-Theory: OGGEP → GDOS → GRIP → GGA . . . 402
89.2 It begins from a constitutional layer . . . 402
89.3 It geometrizes gauge structure at E_{GUT}, not gravity at E_{Pl} . . . 402
89.4 Dark matter, inflation exit, and dark energy are linked, not merely coexisting 403
89.5 Observation itself is treated as open-system projection . . . 403
89.6 It is austere in a correlated way . . . 403
90 Comparison with the leading frameworks . . . 404
91 The Ten Structural Innovations of the Cosmological Closure . . . 407
92 Accomplishments: the seventeen problems, resolved, conditional, or frontier . . . 409
93 The honest boundary: what is proved, what is conditional, what is open . . . 411
93.1 What is proved, what is suggested, what remains open . . . 411
93.2 What GG-Theory computes, and what it does not: the cosmology tier ladder 412
93.3 Dependencies . . . 414
93.4 The problem tally . . . 414
94 The cost of being right: zero parameters, maximal exposure . . . 414
95 Kill criteria and the honest test ledger . . . 415
95.1 The decade timeline: what 2030 and 2035 decide . . . 418
95.2 Hostile audit: what if every test fails? . . . 419
96 The bottom line . . . 421
97 Part XIII in plain English: what is proved, what is conditional, and what is open . 422

Part XIV. Conclusions: the summit view **423**
98 The summit thesis . . . 423
99 The composition of the universe is a number . . . 423
100 One origin, from the first instant to today . . . 424
101 One universe, one story above the Standard Model . . . 425
102 The fourth revolution, written across the sky . . . 425
103 A universe made answerable . . . 426
104 One staircase, from the Big Bang to living molecules . . . 427
105 The trio: from the electron, through the cosmos, to the question of us . . . 428
106 Fine-tuning, answered without luck and without a multiverse . . . 429
107 A personal word from the cosmological floor . . . 430
108 Closing remarks: why are we here? . . . 430
Acknowledgments . . . 432

Part XV. Technical Appendices **433**
A Summary of Abbreviations . . . 435
B Glossary of Key Symbols . . . 440
C Codebook glossary (compact) . . . 442
D Global naming convention for cross-cohort citation . . . 444
E Direct derivation of the 4D Einstein equation, the Friedmann equations, and ΛCDM 445
E.1 The four debts of the textbook derivation . . . 447

E.2 Step 1 — the six-dimensional stage and its ledger . . . 448
E.3 Step 1 — the bulk action, read term by term . . . 449
E.4 Step 1 — the Einstein–Hilbert term descends to the boundary . . . 449
E.5 Step 1 — Newton's constant computed, and the scale ladder . . . 450
E.6 Step 1 — the boundary Einstein equation, earned . . . 451
E.7 Step 2 — the FRW metric and its conventions . . . 453
E.8 Step 2 — Christoffel symbols . . . 454
E.9 Step 2 — the Ricci tensor and Ricci scalar . . . 455
E.10 Step 2 — the Einstein tensor: components . . . 456
E.11 Step 2 — the perfect-fluid stress-energy tensor . . . 457
E.12 Step 2 — the two Friedmann equations, the acceleration equation, and the continuity equation . . . 457
E.13 Step 2 — spatial curvature, the critical density, and the density parameters 459
E.14 Step 2 — worked numerical example: the age of the universe and the critical density in flat ΛCDM . . . 460
E.15 Step 3 — the source ledger: the components from geometry . . . 461
E.16 Step 3 — the fractions, and the epoch identity . . . 462
E.17 Step 3 — the initial conditions (paying debt 4) . . . 463
E.18 Step 3 — the assembled expansion law . . . 463
E.19 What this appendix establishes . . . 464
F Cosmic expansion, scaling laws, and thermal history . . . 469
F.1 The continuity equation and the scaling law $\rho \propto a^{-3(1+w)}$ . . . 470
F.2 Solving for $a(t)$: radiation, matter, and Λ domination . . . 471
F.3 Redshift $1+z = 1/a$ and the redshifting of momenta and temperature . . . 472
F.4 Entropy conservation, $g_{*s}T^3a^3 = \text{const}$, and reheating of the photons . . . 473
F.5 Freeze-out: the Γ/H competition and how decoupling works . . . 474
F.6 The thermal milestones: equality, recombination, BBN . . . 475
F.7 The three horizons: comoving Hubble radius, particle horizon, event horizon 476
F.8 Worked numerical examples . . . 477
F.9 The five clocks of cosmology: t, a, z, T, and r . . . 477
F.10 What the standard thermal history assumes — and cannot explain . . . 478
F.11 The geometric derivation: the cosmic clock as a GDOS descent . . . 478
F.12 Status . . . 481
G Slow-roll inflation from first principles . . . 483
G.1 The three problems of the hot Big Bang . . . 485
G.2 Counting e-folds: how much inflation is enough . . . 487
G.3 A scalar field in FRW: action, stress tensor, equation of motion . . . 487
G.4 The slow-roll approximation and the potential slow-roll parameters . . . 488
G.5 The number of e-folds from the potential . . . 489
G.6 Quantum fluctuations and the scalar power spectrum . . . 489
G.7 The tensor spectrum, r, and the consistency relation . . . 490
G.8 Reheating: from cold inflaton to hot Big Bang . . . 491
G.9 Worked example (STANDARD TEXTBOOK PEDAGOGY): quadratic and Starobinsky potentials . . . 492
G.10 What standard inflation explains — and what it leaves open . . . 493
G.11 The geometric derivation: the two-basin descent, taught slowly . . . 494
H Big-Bang nucleosynthesis and the lithium-7 problem . . . 498
H.1 The neutron-to-proton ratio and weak freeze-out . . . 499
H.2 The deuterium bottleneck . . . 502
H.3 The reaction network: D, He-3, He-4, and Li-7/Be-7 . . . 503
H.4 Dependence on the baryon-to-photon ratio . . . 504
H.5 Observational status: triumph and the lithium-7 problem . . . 504
H.6 Worked example: n/p at freeze-out and the helium fraction . . . 505

H.7 The geometric derivation: the lithium-7 deficit as a closed-packet $1/\pi$ port 506
I Resonant leptogenesis and the lithium-7 port . . . 509
I.1 Type-I seesaw matching to the GG-A9-SM Weinberg operator . . . 510
I.2 Resonant CP asymmetry on the inherited phases . . . 511
I.3 The partial-resonance mechanism: the flavor-measurement regulator . . . 513
I.4 Washout and the strong-washout efficiency . . . 518
I.5 Assembling η_B . . . 519
I.6 The flavored Boltzmann integration: the window deficit measured . . . 519
I.7 The open-system A=7 port and the $1/\pi$ survival fraction . . . 522
J Linear perturbations and the growth of structure . . . 525
J.1 The perturbed FRW metric in conformal-Newtonian gauge . . . 527
J.2 The perturbed fluid: continuity, Euler, Poisson . . . 528
J.3 The sub-horizon growth equation . . . 530
J.4 The growth factor $D(a)$ and its exact matter-dominated solutions . . . 531
J.5 The Meszaros effect: suppression during radiation domination . . . 532
J.6 The growth index: $f \approx \Omega_m^{0.55}$ . . . 533
J.7 Transfer function and the matter power spectrum $P(k)$ . . . 533
J.8 From $P(k)$ to σ_8 and S_8 . . . 534
J.9 What the standard growth theory leaves open: the S_8 tension . . . 535
J.10 The retired geometric derivation: the two-stage corridor and the sign of S_8 (record) . . . 536
J.11 Legacy growth response, cross-consistency Jacobian, and worked derivations (Part VI) . . . 538
J.11.1 Legacy S_8 growth-suppression linear-response . . . 538
J.12 Cross-consistency and falsifiability inside the growth block . . . 539
J.12.1 The minimal Part VI parameter vector . . . 539
J.12.2 Sign matrix and one-ledger discipline . . . 540
J.12.3 Consistency inequalities that must hold simultaneously . . . 541
J.12.4 A worked one-line closure check . . . 542
J.12.5 What a full numerical pipeline would have to do next (record) . . . 543
J.12.6 The kill table for Part VI as a whole (record) . . . 543
J.12.7 Why this section matters scientifically . . . 544
J.13 Worked derivation notebook for the Part VI observables . . . 544
J.13.1 Derivation notebook A: how a narrow early residue shortens the sound horizon (a lever this paper no longer carries) . . . 544
J.13.2 Derivation notebook B: why a five-percent broad growth shift is natural (a lever this paper no longer carries) . . . 545
J.13.3 Derivation notebook C: exponential amplification in the halo tail 545
J.13.4 Derivation notebook D: how a streaming damping would lower the filtering mass (a lever this paper no longer carries) . . . 546
J.13.5 Derivation notebook E: freeze-out of the acoustic phase at recombination . . . 546
J.13.6 Derivation notebook F: the tensor corridor is small but nonzero 546
J.13.7 Derivation notebook G: the difference between broad S_8 scales and first-light scales . . . 547
J.13.8 Final notebook takeaway . . . 547
K The cosmic microwave background – acoustic peaks and the angular power spectrum 549
K.1 The tightly-coupled photon-baryon fluid and its acoustic oscillations . . . 551
K.2 The sound speed and the sound horizon . . . 552
K.3 The angular-diameter distance to last scattering and the acoustic angular scale . . . 553
K.4 Peak positions: the harmonic series and the first peak near $\ell_1 \approx 220$ . . . 554
K.5 Baryon loading and the odd/even peak-height alternation . . . 555

K.6 Silk damping and the damping tail . . . 555
K.7 The temperature angular power spectrum: C_ℓ and D_ℓ . . . 556
K.8 How the six ΛCDM parameters map onto peak features . . . 557
K.9 The geometric derivation: what the cohort computes, and what it does not 558
K.10 Legacy early-dark-energy linear-response: a non-unique estimate of the corridor gap . . . 562
K.10.1 Deriving the early-universe shift from the inherited source ledger 562
K.10.2 The local lift: the effective void correction, with the gauged-baryon route excluded . . . 563
K.10.3 Compact numerical closure of the legacy estimate . . . 564
K.10.4 The retired two-stage corridor, as the body stated it before its retirement (record) . . . 565
K.10.5 What the pair actually delivers, computed (record) . . . 566
K.10.6 The two-stage corridor: one Friedmann structure, two projections (record) . . . 567
K.10.7 The local branch and the inherited baryonic sector (record) . . . 568
K.10.8 The GDOS interpretation: a mis-specified dictionary problem rather than a contradiction of laws (record) . . . 569
K.10.9 A compact derivation of the numerical corridor (record) . . . 569
K.10.10 The retired growth block on the early residue, as the body stated it before its retirement (record) . . . 569
K.10.11 Why the same early correction that helps H_0 usually hurts growth (record) . . . 569
K.10.12 The role of the Kaxion and the inherited neutrino ledger (record) 570
K.10.13 Why the S_8 solution does not automatically contradict the early-galaxy solution (record) . . . 570
K.10.14 A compact consistency ladder for the growth block (record) . . . 571
K.10.15 Kill criteria for the growth module (record) . . . 572
L Topological background and the pi-locked inventory . . . 573
L.1 Differential forms and the wedge product . . . 574
L.2 The field strength two-form F and the four-form $\mathrm{tr}(F \wedge F)$ . . . 575
L.3 The (first) Pontryagin density and Pontryagin / Chern numbers . . . 575
L.4 Characteristic classes, and why their integrals are integers or rationals in the standard setting . . . 576
L.5 The OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ . . . 577
L.6 What standard cosmology cannot explain: the coincidence problem . . . 577
L.7 The idea of a "unit-trace closure" . . . 578
L.8 The $\Omega_m = 1/\pi$ argument as the paper states it (Steps 1–5) . . . 578
L.9 The residual-Pontryagin-normalization reading, stated honestly . . . 580
L.10 The fiber-measure resolution: $\Omega_m = 1/\pi$ as the Hopf ratio . . . 580
L.11 Why an irrational number cannot be a topological invariant in the standard sense . . . 581
L.12 Worked numerical example: the inventory against Planck . . . 581
L.13 Status . . . 583
M Cold dark matter from vacuum misalignment – the Kaxion . . . 585
M.1 A light pseudo-scalar as a misalignment condensate . . . 587
M.2 Equation of motion in an expanding universe . . . 588
M.3 Frozen by Hubble friction, then released . . . 588
M.4 Redshifting as cold, pressureless matter . . . 589
M.5 The relic-abundance formula . . . 590
M.6 The QCD relation $m_a f_a = \sqrt{\chi_{\mathrm{QCD}}}$ . . . 592
M.7 What the standard misalignment story leaves free . . . 592
M.8 The geometric reading: one calibration, then an inversion . . . 592

M.9 Worked numerical example . . . 595
N Black-hole entropy A/4 from induced gravity . . . 598
N.1 The four laws of black-hole mechanics and the thermodynamic analogy . . 599
N.2 The Hawking temperature $T_H = \hbar\kappa/2\pi$ and the surface gravity κ . . . 600
N.3 The Bekenstein–Hawking entropy $S = A/4G_N\hbar$ and the first law $dM = T\,dS$ 601
N.4 Sakharov's induced gravity: the Einstein–Hilbert action is generated by integrating out fluctuations . . . 602
N.5 The heat-kernel expansion: one coefficient sets both $1/G_N$ and the entanglement entropy . . . 603
N.6 The conical-replica (deficit-angle) method: the coefficient $1/4 = 4\pi/16\pi$, regulator-independently . . . 604
N.7 What standard physics leaves open: the origin of $A/4$ . . . 605
N.8 The geometric derivation: $A/4$ from the canonical bulk (Arisaka (2026, C3-BH)) . . . 605
N.9 Information and the Page curve: conceptually closed, quantitatively open . 606
N.10 A worked numerical example: a one-solar-mass Schwarzschild black hole . 607
N.11 Worked derivation notebook for the Part VII observables . . . 609
N.11.1 Derivation notebook A: Schwarzschild area, entropy, and temperature . . . 609
N.11.2 Derivation notebook B: evaporation law and Page-time estimate 610
N.11.3 Derivation notebook C: the operational map and why pure-to-mixed is not paradoxical . . . 610
N.11.4 Derivation notebook D: horizon mass as a timestamp for PBH formation . . . 611
N.11.5 Derivation notebook E: why the minimal PBH abundance is exponentially tiny . . . 611
N.11.6 Derivation notebook F: evaporation survival threshold . . . 612
N.11.7 Derivation notebook G: tensor amplitude from the plateau height 612
N.11.8 Derivation notebook H: today's stochastic background . . . 613
N.11.9 Final notebook takeaway . . . 613
O Internal source ledger for GG-A10-Cosmos . . . 613
P Cross-paper inheritance table (compact) . . . 614
Q Consolidated benchmark audit . . . 615
R Named-theorem-and-conjecture deck . . . 617
S Named-lemma promotion roadmap . . . 619
T Falsification-test catalog . . . 620

Part XVI. Common Objections and Responses 622
O1 For the curious reader: the big questions, plainly answered . . . 623
O2 The ΛCDM and CMB-precision critic . . . 632
O3 The inflation and dark-energy critic . . . 636
O4 The Hubble-tension and large-scale-structure critic . . . 639
O5 The dark-matter, baryogenesis, and BBN critic . . . 644
O6 The gravitation, black-hole, and primordial-gravitational-wave critic . . . 648
O7 The quantum-field-theory and string-theory critic . . . 652
O8 The methodology, honesty, and numerology critic . . . 656
O9 The cohort-coherence and philosophy-of-cosmology critic . . . 659
O10 The curious reader returns: the multiverse and the three horizon problems . . . 664

Part XVII. References 667

Preface

This is a long paper about a large subject — where the universe came from, what it is made of, and why — and it is written so that a reader who is not a professional cosmologist can follow the argument and judge it. The mathematics is here for those who want it, but the central ideas are meant to be understood in plain language first. Before the formal matter begins, a word of welcome, and an explanation of why this paper is addressed to a wider circle than cosmologists.

Who this paper is for. It is written with a particular reader in mind: the graduate student — in physics, certainly, but just as much in chemistry, biology, engineering, computer science, mathematics, or philosophy — who suspects that the largest questions can still be asked seriously, and wants to see one attempt worked out in full, with its evidence and its weaknesses in the same place. You were probably taught that the composition of the universe is a list of measured numbers: so much matter, so much dark energy, an expansion rate, a handful of primordial parameters, all read off the sky and none explained. This paper argues that those numbers may instead be consequences — outputs of one geometry rather than inputs to a fit. You do not need a cosmology background to weigh that claim. Part I poses every problem from scratch, in words; the appendices of Part XV are a self-contained graduate course — the Friedmann equations, inflation, nucleosynthesis, the microwave background, each built line by line for a first-year student meeting it for the first time; and the Questions-and-Answers Part answers, in plain language, the questions people actually ask: about the multiverse, fine-tuning, free will, and what came before the Big Bang.

What this paper claims. The claim is easy to state and hard to earn: that the large-scale contents of the universe are not free numbers to be measured, but consequences of a single six-dimensional geometry — the same geometry that, in a companion paper, fixes the masses of the elementary particles. The fraction of the cosmos in matter, in dark energy, and in ordinary atoms; the moment the first atoms formed; the identity of the dark matter; the cause of inflation — all are offered as quantities the geometry forces, with no adjustable dials.

An honest accounting. A paper that claims this much owes the reader a candid ledger. Of twenty-one cosmological quantities, one is fully derived; nineteen are derived up to a named conjecture — for the inventory, the closed-packet conjecture, whose once-open mathematical step has been traced to one declared input, the timing of the present epoch — and one, the amplitude, is declared; the paper marks each one, and tunes nothing to the data after the fact. The test of an honest framework is not that it answers everything on the first pass, but that it always knows which questions it has answered, which it has half-answered, and which it has not yet touched.

The decade of the verdict. Here is why this paper is addressed to graduate students in particular, and why now. Every headline number in this book is exposed to instruments that will run in the coming decade. Gravitational-wave standard sirens will read the cosmic expansion rate directly, and this framework needs them to say ≈ 67, not the distance ladder's 73. The Simons Observatory and LiteBIRD will hunt the primordial gravitational-wave background, and this framework has pinned

it, two-sidedly, at $r = 0.0038$ — too high fails, too low fails. Euclid and the Rubin Observatory will weigh whether dark energy drifts, and this framework says it cannot: $w = -1$ exactly, forever. Haloscopes will sweep the microwave band where the Kaxion is placed, near 1.95 GHz, in a window narrow enough to search honestly. By the mid-2030s the sky will have voted on all of it. And the people who will build those instruments, calibrate them, write their analysis pipelines, and catch their systematic errors are, for the most part, the people who are in graduate school today. The verdict on this framework belongs to your generation. That is not a rhetorical flourish; it is the schedule.

An invitation to break it. So do not read this paper deferentially. It was written to be attacked, and it tries to make the attack easy: every claim is graded in a published codebook — Derived, Derived-conditional, or Asserted — so you can see at a glance which numbers are earned and which are still owed; the retired claims are kept visible together with what refuted them; and the Objections Part collects the ninety-seven hardest questions we could pose against ourselves, because the ones we missed are the ones that matter. If you are looking for a problem worth a thesis chapter, the named open lemmas are stated precisely enough to be settled either way — above all the Hopf-fiber identification on which the matter fraction $\Omega_m = 1/\pi$ now visibly rests, and the trials-factor question, which the Monte-Carlo has answered only for an unrestricted alphabet: a restricted, stratum-assigned enumeration is still open. Break one, and you will have done physics a service and this program a favor; prove one, and you will have promoted an entry of the codebook. Either outcome is progress, and both are open to you.

The larger arc. This paper is the middle panel of a deliberate arc. Its companion GG-A9-SM derives the particle world — the masses, the charges, the identity of every electron — from the same geometry; the capstone GG-A11-Origin carries the arc onward toward life and mind. Whatever your own field, the arc touches it at one point: *identity*. Why is every copy of your object — every electron, every carbon atom, every ATP molecule, every ribosome — the same as every other copy, everywhere, including in places that could never have compared notes? Part X of this paper asks that question about the universe itself and proposes that physics, chemistry, and biology inherit their sameness from a single geometric principle. That is why a chemist, a biologist, or a computer scientist may find that the deepest Part of a cosmology paper speaks, unexpectedly, to their own bench.

How to read this paper. No reader needs all the body Parts at once, and there are three honest ways in. *Twenty minutes:* the one-page statement of what is actually claimed (just after the opening figures), then the Executive Summary, then the decade-timeline of kill tests in Part XIII — enough to decide whether the argument deserves your evening. *One evening:* add Part I for the problems, Part IV for the descent that launches inflation and leaves the dark sector behind, and Part X for the identity of the cosmos and the view from the top. *A semester:* the whole book, in order, with the Part XV appendices as the accompanying course and the Objections Part as the exam. And a note for the reader whose interest is strictly the physics: the falsifiable core stands entirely on its own — a reader may stop at the end of Part XIII and lose no testable content, taking

the Cosmos-to-observer arc of Part X as optional (§105). Throughout, the body keeps the standard observations — Planck, the supernovae, the weak-lensing surveys — in plain view, and every result is stated so that it can be checked, and challenged, on its own. However you enter: welcome. Bring your skepticism; the paper was written in its honor. And if, ten years from now, the sirens say 73, the tensor sky is empty at $r = 0.004$, or the Kaxion window has been swept clean and silent — then this book will have failed honorably, in public, the way a scientific claim should. That outcome, too, would be worth your generation's while to deliver.

Part I
Introduction

The oldest questions about the universe — and a new geometric answer

1 The oldest question, asked precisely

Every culture that has ever looked up has asked the same three questions. Where did the universe come from? What is it made of? Where is it going? For most of human history these were questions for myth. In the last hundred years they became, for the first time, questions for measurement — and that transformation is one of the great achievements of the scientific age.

The arc is short and astonishing. Copernicus moved the Earth from the center; Newton gave the heavens and the falling apple one law; Einstein made space and time dynamical (Einstein, 1917); Hubble, in 1929, found that the galaxies are flying apart (Hubble, 1929) and so handed us an expanding universe with a beginning (Friedmann, 1922; Lemaître, 1927). The hot Big Bang followed, confirmed in 1965 when two engineers found its afterglow as a faint microwave hiss in every direction (Penzias and Wilson, 1965). By the turn of this century these threads had been woven into a single quantitative model — ΛCDM, the concordance cosmology (Planck Collaboration, 2018) — that reproduces the cosmic microwave background, the abundances of the light elements, the large-scale distribution of galaxies, and the accelerating expansion (Perlmutter et al., 1999; Riess et al., 1998), with a precision that would have seemed miraculous a lifetime ago.

And yet. For all its precision, the standard model of cosmology is, at its foundation, a model that *measures* the universe without *explaining* it. It tells us, to a fraction of a percent, that the cosmos is about five percent ordinary matter, twenty-seven percent dark matter, and sixty-eight percent dark energy — and it cannot tell us why. It describes an early burst of inflation without saying what drove it. It carries a cosmological constant whose value is the most catastrophic mismatch between theory and experiment in the history of physics. It is the most precisely measured account of the universe ever achieved, and the one most riddled with numbers it can record but never derive.

This paper, GG-A10-Cosmos, is about those numbers. It proposes that they are not free, not fortuitous, and not forever beyond explanation — that the composition and history of the universe are written into the shape of a single geometry, and can be read out rather than fitted. It is, in other words, the introduction whose questions the Conclusions of this paper (§XIV) return to answer.

2 A century of cosmology: how the Standard Model of the universe was built

The concordance model did not arrive in one piece, and no one designed it. It is the residue of a century of measurement, in which each new instrument forced a revision the previous picture could not absorb. It is worth retracing that arc, because the model's authority — and the precise shape of its unfinished business — comes from how it was built; the milestones are gathered in Figure 7.

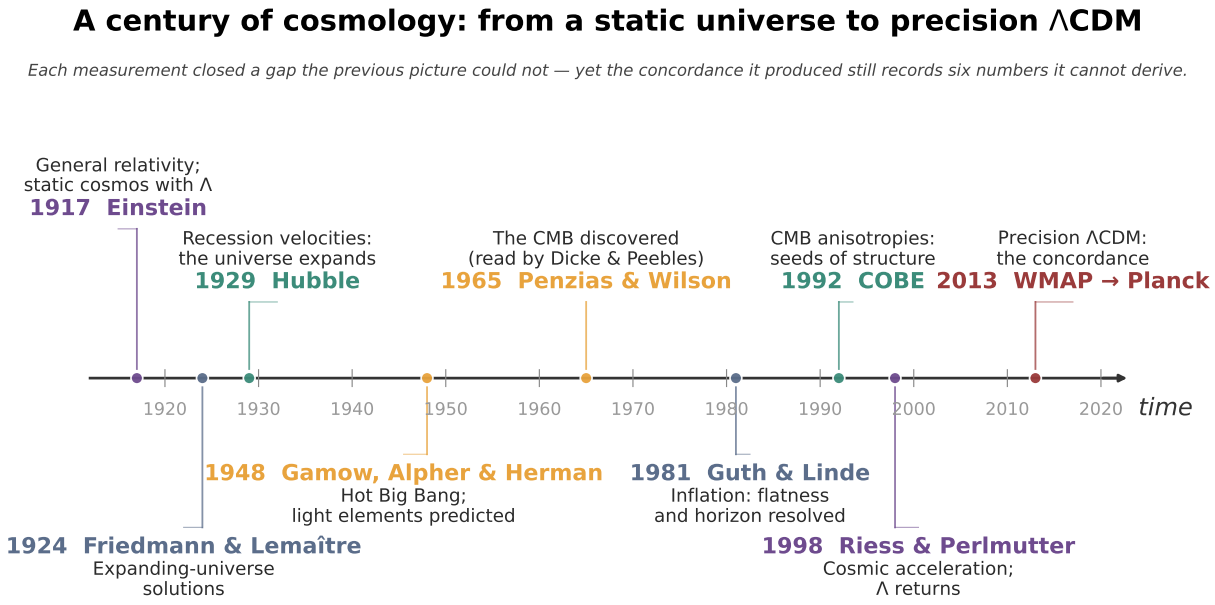

Figure 7: A century of cosmology. The discoveries that assembled the Standard Model of the universe, each forced by a measurement the previous picture could not accommodate: the dynamical universe of Einstein, Friedmann, and Lemaitre; Hubble's expansion; the hot Big Bang and its predicted relic radiation; the CMB; inflation; the CMB anisotropies; cosmic acceleration; and the precision concordance of WMAP and Planck. The model was not designed — it was forced, step by step — and it nonetheless ends with six numbers it can record but not derive.

The modern story opens in 1915–1917, when Einstein completed general relativity and applied it to the universe as a whole (Einstein, 1917). Expecting a static cosmos, he found that gravity would not permit one: a universe of matter must collapse. To hold it still he introduced a new term into his equations, the cosmological constant Λ, tuned to balance the pull of matter. Within a decade that fix was overtaken. Friedmann in 1922 (Friedmann, 1922) and, independently, Lemaitre in 1927 (Lemaître, 1927) showed that the natural solutions of general relativity are not static at all but dynamic — expanding or contracting — and Lemaitre connected the expansion to the redshifts of distant nebulae and to the idea of a dense beginning.

The decisive measurement came in 1929. Hubble (Hubble, 1929), calibrating distances to galaxies and reading their redshifts, found that recession velocity grows in proportion to distance: the universe is expanding. Einstein abandoned Λ, and the expanding universe became the new

baseline. Its immediate implication is a beginning: run the expansion backward and the cosmos was once hotter and denser than any furnace. In the 1940s Gamow and his collaborators Alpher and Herman pursued that implication quantitatively (Alpher et al., 1948), working out how the lightest nuclei would have been forged in the first minutes and predicting that the early fireball should have left a faint relic glow, cooled by expansion to a few degrees above absolute zero.

That prediction sat unconfirmed for two decades, until in 1965 Penzias and Wilson (Penzias and Wilson, 1965), troubleshooting a radio antenna, found an inexplicable microwave hiss coming from every direction — a signal that Dicke and Peebles immediately recognized as the predicted afterglow (Dicke et al., 1965). The cosmic microwave background (CMB) turned the hot Big Bang from a hypothesis into the standard framework. Two structural puzzles remained, however. Why is the universe so very flat, when general relativity drives curvature away from flatness over time? And why is the CMB at the same temperature in regions that, in the standard expansion, were never in causal contact? In the early 1980s inflation supplied a single answer to both: Guth (Guth, 1981), and then Linde (Linde, 1982, 1983) together with Albrecht and Steinhardt (Albrecht and Steinhardt, 1982) in the form that survives — with an independent precursor in Starobinsky's curvature-driven model (Starobinsky, 1980) — proposed a brief epoch of accelerated expansion in the first instant, which flattens space and stretches one causal patch across the whole sky, and, as a bonus, converts quantum fluctuations into the seeds of galaxies (Mukhanov & Chibisov, 1981; Mukhanov *et al.*, 1992; Starobinsky, 1979).

The seeds were found in 1992, when the COBE satellite (Smoot *et al.*, 1992) detected temperature ripples in the CMB at the level of one part in a hundred thousand — the large-angle imprint, through the Sachs–Wolfe effect (Sachs & Wolfe, 1967), of the density variations from which all structure grew. The last major ingredient arrived at the end of the decade. In 1998–1999 two teams measuring distant Type-Ia supernovae, led by Riess (Riess et al., 1998) and Perlmutter (Perlmutter et al., 1999), found that the expansion is not slowing under gravity but *accelerating*. Einstein's discarded Λ returned, no longer as a balancing trick but as the dominant component of the universe. Through the 2000s and 2010s the WMAP and Planck satellites (Planck Collaboration, 2018) mapped the CMB to sub-percent precision, pinning every ingredient and welding the threads into a single quantitative model (Weinberg, 2008). That model is ΛCDM — and the century of measurement that produced it also handed it the unfinished business that the rest of this Part confronts.

3 ΛCDM: the concordance model, and why it works

Before asking what ΛCDM cannot do, one must see clearly how much it does, because a new idea earns nothing by underselling the incumbent it hopes to deepen. ΛCDM is general relativity applied to a universe that is, on the largest scales, homogeneous and isotropic — the same everywhere and in every direction — and expanding according to the Friedmann equations (Misner *et al.*, 1973; Wald, 1984; Weinberg, 1972). Into that geometry it places a short list of ingredients: radiation (the photons of the CMB and the cosmic neutrinos), ordinary "baryonic" matter (the atoms of stars, gas, and us), cold dark matter (an unseen, slowly moving substance that outweighs the atoms five to one), and the cosmological constant Λ (the dark energy driving acceleration). Space is taken

to be flat, and the whole structure is seeded by a nearly scale-invariant spectrum of primordial fluctuations of the kind inflation predicts (Lyth & Riotto, 1999; Mukhanov *et al.*, 1992).

What makes the model so powerful is its economy. The entire expansion history and the full statistical pattern of the CMB and the galaxy distribution are fixed by just *six* numbers: the physical densities of baryons and of cold dark matter, $\omega_b \equiv \Omega_b h^2$ and $\omega_c \equiv \Omega_c h^2$; the angular size of the sound horizon θ_* that sets the expansion rate; the optical depth τ to the epoch when the first stars reionized the gas; and the amplitude A_s and tilt n_s of the primordial spectrum. From these six, everything else is derived — the Hubble constant H_0, the matter and dark-energy fractions Ω_m and Ω_Λ, the clustering amplitude σ_8, and the 13.8-billion-year age of the universe. A model of the entire cosmos runs on six adjustable dials.

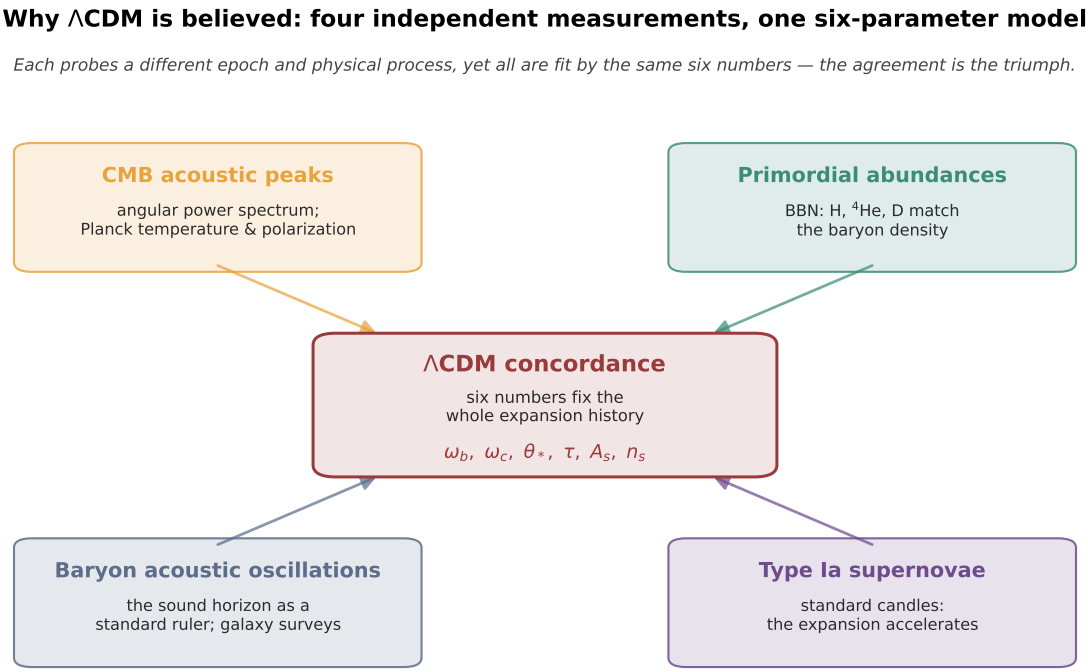

Figure 8: Why ΛCDM is believed. Four independent classes of measurement — the CMB acoustic peaks, the primordial light-element abundances, the baryon acoustic oscillation scale in large-scale structure, and the acceleration read from Type-Ia supernovae — each probe a different cosmic epoch and a different physical process, yet all are fit by the same six parameters (ω_b, ω_c, θ_*, τ, A_s, n_s). The agreement among independent probes, not any single one, is the triumph that makes ΛCDM the Standard Model of cosmology.

The reason cosmologists trust it is not any single measurement but the way *independent* measurements, probing different epochs and different physics, converge on the same six numbers (Figure 8). Four pillars carry the weight. The first is the CMB itself: before the universe became transparent, photons and baryons formed a single fluid in which sound waves rang, and those waves froze into a precise pattern of hot and cold spots whose statistics — the heights and spacings of the "acoustic peaks" — are predicted in detail and measured by Planck to exquisite accuracy (Hu & Sugiyama, 1995; Peebles & Yu, 1970; Planck Collaboration, 2018; Sachs & Wolfe, 1967; Silk, 1968;

Sunyaev & Zeldovich, 1970). The second is primordial nucleosynthesis: the same baryon density that fits the CMB peaks independently predicts the amounts of hydrogen, helium, and deuterium forged in the first three minutes, and the prediction matches observation across nine orders of magnitude in abundance (Cyburt *et al.*, 2016; Fields, 2011; Pitrou *et al.*, 2018; Steigman, 2007; Wagoner *et al.*, 1967). The third is large-scale structure: the same sound horizon that sets the CMB peaks leaves a standard ruler — the baryon acoustic oscillation scale — imprinted in the clustering of galaxies (Cole *et al.*, 2005; Eisenstein *et al.*, 2005), while the growth of that clustering from the tiny COBE-level seeds, by gravitational instability, follows the rate ΛCDM predicts (Bardeen, 1980; Bardeen *et al.*, 1983, 1986; Kodama & Sasaki, 1984; Lifshitz, 1946; Peebles, 1993; Press & Schechter, 1974). The fourth is the accelerating expansion read directly from Type-Ia supernovae (Perlmutter et al., 1999; Riess et al., 1998).

These four probes could easily have disagreed. The CMB is a snapshot from redshift $z \approx 1100$; nucleosynthesis is a clock running at $z \approx 10^9$; the supernovae and galaxy surveys report from the recent universe at $z \lesssim 2$. They involve different forces, different particles, and different decades of instrument-building. That they are all fit by one set of six numbers, with no adjustment between them, is the concordance — and the concordance, not any lone result, is why ΛCDM is the Standard Model of cosmology. The standard graduate treatments develop each pillar in full (Dodelson, 2003; Kolb & Turner, 1990; Lyth & Riotto, 1999; Mukhanov, 2005; Peebles, 1993; Weinberg, 2008).

And yet the triumph is one of *description*, not explanation. ΛCDM tells us, to a fraction of a percent, that the universe is about five percent atoms, twenty-seven percent dark matter, and sixty-eight percent dark energy — and it cannot say why those fractions, and not others. It measures the value of Λ without accounting for why it is so extraordinarily small yet not zero (Carroll, 2001; Copeland *et al.*, 2006; Peebles & Ratra, 2003; Weinberg, 1989); it names dark matter without identifying it (Bertone & Hooper, 2018; Marsh, 2016); it uses inflation without specifying what drove it. The six numbers are the model's inputs, read from the sky and inserted by hand. The success of ΛCDM is precisely what sharpens the question this paper asks: granted that these six numbers describe the universe so well, where do the numbers themselves come from?

4 What the standard model of cosmology measures but cannot explain

Let us be clear, first, about how good the incumbent is, because a new idea earns nothing by underselling what it must respect. ΛCDM — four-dimensional general relativity, a homogeneous expanding metric, cold dark matter, a vacuum-like dark energy, and a nearly scale-invariant spectrum of primordial ripples — is a genuine triumph. At the background level it gives the redshift–distance relation, the expansion history, and the timing of the radiation, matter, and acceleration eras; at the perturbative level it gives the acoustic peaks of the microwave background, the baryon-acoustic scale (Eisenstein *et al.*, 2005), the growth of structure, and a long ledger of survey data (Planck Collaboration, 2018); the standard treatments are the textbooks of Dodelson (2003); Mukhanov (2005); Peebles (1993). Any deeper theory must reproduce this success before it is allowed to differ from it. GG-A10-Cosmos keeps every one of these equations.

The trouble is not the equations. It is that the model runs on roughly half a dozen numbers — the matter density, the baryon density, the dark-energy fraction, the amplitude and tilt of the primordial spectrum — that are measured from the sky and inserted by hand. The model fits the data beautifully; it does not say where the data come from. Modern cosmology, in this precise sense, resembles the Standard Model of particle physics before a derivational program: the formalism works, but several of its source terms are placeholders, not closed explanations.

Sharpened into a finite audit, the open frontier of contemporary cosmology is the set of seventeen problems in Table 1, grouped into five clusters: the two flagship tensions; the dark Universe and its fate; the beginning and the primordial signals; structure, cosmic dawn, and the JWST surprises; and the cracks in the foundations.

Table 1: The open problems of contemporary cosmology (2026), grouped into five clusters, with the reading GG-A10-Cosmos offers for each. The Hubble tension and the JWST early-galaxy results are widely regarded as the two most exciting current problems. Entries marked *frontier* are ones the framework touches but does not yet close.

#	Open problem (current status)	The reading **GG-A10-Cosmos** offers
I. The two flagship tensions		
1	Hubble tension: local $H_0 \approx 73$ vs. early ≈ 67, now $\sim 6\sigma$; JWST sharpened it	One derived background; the ladder the predicted outlier; switch at $z_* = e^7$
2	S_8/σ_8 growth tension (easing toward Planck in 2024–25)	S_8 computed from the declared amplitude; the tension shared, not resolved
II. The dark Universe and its fate		
3	The cosmological-constant problem (Λ's tiny value)	$\Omega_\Lambda = 1 - 1/\pi$ as a residual share
4	Constant or evolving dark energy? (DESI $w_0 w_a$ hint)	Predicts $w = -1$ exactly — a live, falsifiable stand
5	Cosmic coincidence / "why now"	The π-locked inventory $\Omega_m = 1/\pi$
6	The nature of dark matter (WIMP nulls, axions, PBH windows)	The Kaxion, 8.1 μeV, in a narrow window
III. The beginning and the primordial signals		
7	Origin / initial singularity / quantum gravity	A finite, lawful edge at E_{GUT}
8	The cause of inflation	A derived plateau at the top of the geometric descent
9	Primordial gravitational waves (no B-modes; $r < 0.036$)	$r = 0.0038$ — a decisive forward prediction
10	Matter–antimatter asymmetry (baryogenesis)	Leptogenesis on the inherited neutrino ledger
11	Big-Bang nucleosynthesis and the lithium-7 problem (factor ~ 3)	An $A = 7$ open-system port
IV. Structure, cosmic dawn, and the JWST surprises		
12	JWST's too-early, too-massive galaxies ($z \gtrsim 10$)	Early-structure amplification, on the baryonic levers (the corridor route and the streaming damping withdrawn)
13	"Little red dots" / overmassive early black holes	*Frontier*: open-system horizon reading

#	Open problem (current status)	The reading GG-A10-Cosmos offers
14	Cosmic dawn / reionization / 21 cm	*Frontier*: inherited early-structure timing
15	The black-hole information paradox	Horizon as an OGGEP port; exact $A/4$
V. Cracks in the foundations		
16	Cosmic isotropy / dipole and CMB large-angle anomalies	*Frontier*: boundary-projection statistics
17	Neutrino masses and $N_{\rm eff}$ (DESI low Σm_ν)	Inherited neutrino block; $\Delta N_{\rm eff} \leq 0.03$, a derived null

Two of these now dominate the field's attention. The first is the *Hubble tension* (Abdalla *et al.*, 2022; Di Valentino *et al.*, 2021; Riess et al., 2022) — a roughly six-sigma disagreement between the locally measured expansion rate ($H_0 \approx 73$ km s^{-1}Mpc^{-1}) and the value inferred from the early universe (≈ 67), which the James Webb Space Telescope has sharpened rather than removed. The second is JWST's discovery of *surprisingly luminous, massive galaxies already in place less than three hundred million years after the Big Bang.* Honesty requires adding at once that the prevailing reading of the JWST result is that it stresses the astrophysics of how galaxies and their black holes form more than the cosmological framework itself; the stubborn residual is the excess of bright galaxies at the very highest redshifts. GG-A10-Cosmos takes both seriously, and offers a geometric reading of each.

4.1 The seventeen open problems, one by one

The table above is a census; a broader-audience reader deserves to know what each entry actually *means*. We therefore walk the seventeen problems in turn — what each one is, why the standard model of cosmology cannot close it, and, in a closing phrase, the reading this paper will offer. The cluster headings are those of Table 1.

I. The two flagship tensions.

1. The Hubble tension. The expansion rate of the universe today, the Hubble constant H_0, can be measured two ways: directly, by watching nearby supernovae and Cepheid stars recede, which gives $H_0 \approx 73$ km s^{-1}Mpc^{-1} (Riess et al., 2022); and indirectly, by reading the sound waves frozen into the microwave background of the infant universe and evolving them forward with the standard model, which gives ≈ 67. These two numbers should agree, and they do not: the gap is now near six standard deviations (Abdalla *et al.*, 2022; Di Valentino *et al.*, 2021; Verde *et al.*, 2019), far too large to dismiss as measurement error, and the James Webb Space Telescope has tightened rather than relaxed it (Casertano et al., 2026). Either one measurement hides an unknown systematic, or the standard model is missing a physical ingredient between the early and the late universe. GG-A10-Cosmos reads the two values as one derived background rate, 66.8 ± 0.2, which the microwave background reads at 67.4 through the inherited sound horizon and every ladder-free probe is predicted to read directly; the ladder's 73.5 is the one measurement displaced from it, by a calibration systematic the framework predicts (Part VI).

2. The S_8–σ_8 growth tension. A second, quieter disagreement concerns not how fast the

universe expands but how strongly matter has clumped into the cosmic web. Surveys of weak gravitational lensing and galaxy clustering have tended to find the present universe slightly *smoother* — a lower clustering amplitude S_8 (Asgari et al., 2021) — than the microwave background, evolved forward, predicts. The discrepancy is milder than the Hubble tension and has eased toward agreement in 2024–25 data, so its very reality is still debated. If genuine, it would signal that the growth of structure departs subtly from the standard prescription. GG-A10-Cosmos computes the amplitude from the declared A_s on the same background, $S_8 = 0.825 \pm 0.010$, and shares the tension with cosmic shear rather than resolving it (Part VI).

II. The dark Universe and its fate.

3. The cosmological-constant problem. The accelerating expansion of the universe is driven by a smoothly spread energy of empty space. If that energy is the quantum vacuum, the naive theoretical estimate exceeds the measured value by some 120 orders of magnitude (Weinberg, 1989) — routinely called the worst quantitative prediction in the history of physics. Why the vacuum energy is so extraordinarily small, yet not exactly zero, is the cosmological-constant problem (Carroll, 2001), and it lies on the unhealed seam between quantum theory and gravity. GG-A10-Cosmos does not tune the number: it reads the dark-energy share as the residual $\Omega_\Lambda = 1 - 1/\pi$ of a fixed geometric inventory, with the absolute vacuum scale a meV-level seesaw residue rather than a free input (Parts IV and X).

4. Constant or evolving dark energy? Even granting a small vacuum energy, is it truly *constant* — a genuine cosmological constant with equation of state $w = -1$ — or does it drift slowly with cosmic time (Copeland *et al.*, 2006; Peebles & Ratra, 2003)? Recent data from the DESI survey have hinted at a possible evolution, a w_0–w_a departure from -1, though the evidence is not yet decisive. The answer decides whether dark energy is a property of spacetime itself or a new dynamical field. GG-A10-Cosmos takes a sharp, falsifiable stand: the late-time vacuum is a completed geometric transition, so $w = -1$ *exactly*, with no appreciable clustering — a clean bet that a secure detection of $w \neq -1$ would lose (Part IV).

5. The cosmic coincidence, or "why now." Matter dilutes as the universe expands while the vacuum energy does not, so the two are comparable only during a brief window of cosmic history — and we happen to be observing inside it. Why should our epoch coincide with the rare moment when matter and dark energy are of the same order? In the standard model this is an unexplained accident stacked on top of the tiny value of Λ. GG-A10-Cosmos reads the present budget as a π-locked inventory with $\Omega_m = 1/\pi$ (Theorem 74.1); crucially, it does not pretend the coincidence vanishes, but shows it demoted from a 10^{120} fine-tuning to a mild, few-percent matter of *timing*, with the form forced and only the observation epoch read from the sky (Part IV, §32; the form/scale/epoch split is §32.11).

6. The nature of dark matter. About a quarter of the cosmic energy budget — five times the ordinary matter — is some substance that bends light and binds galaxies but neither emits nor absorbs it. Decades of searches for the leading candidate, weakly interacting massive particles, have returned nulls (Bertone & Hooper, 2018), redirecting attention to ultralight axion-like particles (Marsh, 2016; Peccei and Quinn, 1977) and to primordial black holes. The standard model of cosmology names dark matter as a slot in the budget but supplies no particle to fill it. GG-A10-

Cosmos identifies it with the Kaxion, a light pseudoscalar near 8.1 μeV placed in a narrow laboratory window (companion paper C1 (Arisaka, 2026, C1-Kaxion); Part V).

III. The beginning and the primordial signals.

7. The origin and the initial singularity. Run the expansion backward and the standard equations end in a singularity — an instant of infinite density where the theory itself breaks down. This is less a prediction than a confession that general relativity cannot describe its own beginning (Misner *et al.*, 1973; Wald, 1984), which is widely believed to require a quantum theory of gravity that does not yet exist. GG-A10-Cosmos replaces the singular beginning with a finite, lawful edge at the grand-unification scale E_{GUT}, the Planck scale becoming a derived rather than a primary quantity (Parts II and III).

8. The cause of inflation. To explain why the universe is so uniform and so flat, cosmology posits an early burst of accelerated expansion called inflation (Albrecht and Steinhardt, 1982; Guth, 1981; Linde, 1982; Starobinsky, 1980). The mechanism works beautifully as a description, but the standard model does not say *what* drove it: the "inflaton" field and its potential are added by hand, chosen from a vast menu to fit the data. GG-A10-Cosmos derives the driver: the scale modulus of the inherited geometry, on a plateau whose *class* is forced by the boundary law and whose height sits one e-fold below the GUT edge — no field added by hand, no potential chosen from a menu; the lock scale is crossed on the way down, not camped on (Part IV). And the derived placement pays an unexpected dividend: because the launch sits at the top of the descent, inflation happened *before any matter existed* — which turns out to answer a question the horizon-problem literature circled for fifty years but never posed in its own right, the *second* horizon problem of Part X: why the universe is everywhere made of the same things, not merely everywhere the same temperature (§72).

9. Primordial gravitational waves. Inflation should imprint a faint background of primordial gravitational waves on the polarization of the microwave background, in a characteristic "B-mode" pattern whose strength is the tensor-to-scalar ratio r (Mukhanov & Chibisov, 1981; Starobinsky, 1979). None has been seen; the present bound is $r < 0.036$. A detection would pin the energy scale of inflation; continued non-detection presses on high-scale models. GG-A10-Cosmos makes this a forward prediction with a number on it: the derived plateau fixes $r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order), a factor nine below the current bound and squarely within Simons-Observatory/LiteBIRD reach — a two-sided target that a measurement inconsistent with it in either direction would break (Part IV).

10. The matter–antimatter asymmetry. The visible universe is built almost entirely of matter, with only a trace of antimatter, in the small but decisive proportion of about six baryons per ten billion photons. Producing this imbalance from a symmetric beginning requires conditions — baryon-number violation, the right kind of symmetry violation, and a departure from equilibrium (Sakharov, 1967) — that the standard model cannot supply in sufficient measure. GG-A10-Cosmos generates it by leptogenesis (Fukugita & Yanagida, 1986; Kuzmin *et al.*, 1985) on the neutrino-and-CP ledger inherited from the particle sector, rather than as a new free input (Part V).

11. Big-Bang nucleosynthesis and the lithium-7 problem. The light elements forged in the first minutes — hydrogen, helium, and traces of lithium (Alpher et al., 1948; Wagoner *et al.*, 1967) — match prediction superbly, with one stubborn exception: the observed primordial lithium-7

is about three times smaller than the standard calculation gives (Cyburt *et al.*, 2016; Fields, 2011; Pitrou *et al.*, 2018). Either the stellar measurements or the early-universe nuclear network hides an error, or new physics depletes lithium. GG-A10-Cosmos reads the deficit as the action of a geometric $A = 7$ "port" that suppresses lithium-7 by the same closed-packet factor that appears in the inventory, leaving the successful hydrogen and helium yields untouched (Part V).

IV. Structure, cosmic dawn, and the JWST surprises.

12. JWST's too-early, too-massive galaxies. The James Webb Space Telescope has found surprisingly luminous, massive galaxies already assembled less than a few hundred million years after the Big Bang — earlier and heavier than the standard structure-formation timeline comfortably allows. Whether this stresses the cosmology or only the astrophysics of early star formation is still debated, but the excess of the brightest high-redshift galaxies is stubborn. GG-A10-Cosmos reads it as a mild, inherited amplification of early structure, carried by the baryonic levers of the ledger — the collapse threshold in the baryonic environment, the one lever left standing — and *not*, as earlier versions claimed, by a small-scale power boost from the corridor's own residue or by a gauged-baryon damping of the streaming velocity, both of which are withdrawn (Part VI, §44.3, §44.4).

13. "Little red dots" and overmassive early black holes. JWST has also revealed a population of compact red objects — and black holes that appear too massive for their host galaxies and their early epoch — straining the picture of how black holes grow so fast. This is a young and rapidly moving observational frontier. GG-A10-Cosmos offers only a tentative, open-system reading of the horizon physics here, and marks the item honestly as *frontier*: touched, not closed (Part VII).

14. Cosmic dawn, reionization, and the 21 cm signal. The epoch when the first stars lit up and reionized the intergalactic hydrogen — "cosmic dawn" — is only now becoming accessible through the redshifted 21-centimeter line of neutral hydrogen, and early claims of a detection remain contested. It is a frontier of measurement as much as of theory. GG-A10-Cosmos expects the timing to follow from its inherited early-structure history, but marks the item *frontier*: a target, not yet a closed prediction (Part VI).

15. The black-hole information paradox. When a black hole forms and slowly evaporates, the standard combination of quantum theory and general relativity appears to destroy information, in contradiction with the unitarity that quantum mechanics demands (Bekenstein, 1973; Hawking, 1975). Reconciling the two has driven a generation of work on horizons, entropy, and holography (Maldacena, 1998; Page, 1993). GG-A10-Cosmos reads the horizon as an open-system projection port in the sense of OGGEP, recovering the exact Bekenstein–Hawking entropy $A/4$ and treating evaporation as a lawful boundary process rather than a loss (Part VII).

V. Cracks in the foundations.

16. The large-angle anomalies and cosmic isotropy. The microwave background is very nearly the same in every direction, yet a handful of large-angle features — curious alignments and a possibly anomalous dipole — sit oddly with the assumption of perfect statistical isotropy. Each is individually marginal, and whether they are real or flukes is unsettled. GG-A10-Cosmos suggests they could reflect the statistics of the boundary projection, but marks the item *frontier*: a hint, not a claim (Part VI).

17. Neutrino masses and the radiation content. Neutrino oscillations prove that neutrinos have mass, but cosmology constrains only their *sum* — and recent DESI data push that sum to surprisingly low values, in some tension with particle-physics expectations — while the effective number of relativistic species $N_{\rm eff}$ measures the radiation budget of the early universe (Steigman, 2007). GG-A10-Cosmos inherits the neutrino block from the particle sector and adds nothing to it: the radiation content is the Standard Model's, $N_{\rm eff} = 3.044$, $\Delta N_{\rm eff} \leq 0.03$ — a derived null to within the freeze-in ceiling, and a two-sided one (Part IX).

Two features of this list matter for what follows. First, most of these are not failures of calculation but gaps in explanation: the Friedmann equations do not tell us what inflation *is*, what the vacuum term *means*, or why the present energy fractions take the values they do. Second, the problems are coupled — inflation sets the primordial spectrum, which sets the microwave background and the timing of galaxies; the baryon content inferred from the CMB constrains nucleosynthesis; the dark sector governs the growth of structure. A genuine explanation should therefore not arrive as seventeen separate patches. It should arrive as one constrained ledger that the open problems audit at once (that ledger is the cosmological codebook of §68.1, Table 40).

It is worth saying plainly what does *not* count as such an explanation. To answer "why is the dark-energy density so small?" by replying that in a vast multiverse only the rare regions with a small value can host observers is to offer a selection effect, not a derivation — it explains the specific number we measure no more than "the survivors survived" explains who lives. A real answer must produce the value, and must risk being wrong. That is the standard GG-A10-Cosmos sets for itself.

5 A radically new idea: the universe as the boundary of a geometry

Suppose you have never encountered this framework before. Here it is in one sentence. GG-Theory (Geometry–Gauge Theory) is a program that derives the laws of physics, the contents of the cosmos, and — in its later chapters — the conditions for life and mind, from a *single six-dimensional geometry with no continuous free parameters to adjust*. There are no dials. Once the geometry is fixed, everything else is supposed to follow as a consequence.

The central move is easy to picture even without the mathematics. Our familiar four-dimensional world — three of space, one of time — is not, in this picture, the whole of reality. It is the *boundary* of a richer six-dimensional bulk: a throat with one extra radial direction, encoding energy scale, and one extra compact angle, encoding the locked gauge structure. What a laboratory or a telescope can register is the shadow this bulk casts on its four-dimensional edge — a boundary/bulk relation in the spirit of horizon thermodynamics and the thermodynamic reading of gravity (Gibbons & Hawking, 1977; Jacobson, 1995). The companion papers of the program show that the same geometry, read at the scale of a single particle, fixes the mass and charge of the electron (Arisaka, 2026, P3-Electron). GG-A10-Cosmos is what happens when you read that very same geometry at the largest scale of all. The cosmos and the electron come from one ledger, one story apart.

If that is right, then the deepest question in cosmology changes character. It is no longer "what values shall we measure and accept?" but something far sharper:

Are the numbers that define our universe — the dark-energy density, the matter fraction,

the tilt of the primordial spectrum — free inputs to be measured, or are they theorems of the geometry, fixed by π and a short ledger of integers?

This is the cosmological form of an ideal that physics has carried, and quietly abandoned, since Euclid: the ideal of a description with nothing left over to choose. GG-A10-Cosmos's wager is that for cosmology the ideal is recoverable — that the cosmic numbers are not free, not postulated, and not fitted; they are derived.

6 The central claim: the cosmic inventory consists of open-system observables

The question just posed has a precise answer, and it is worth stating in one place, without hedging, before the specific numbers appear. The claim is not that this paper fits the cosmic inventory better than ΛCDM. It is a claim about *what the inventory is*. Figure 9 traces the whole argument in one view.

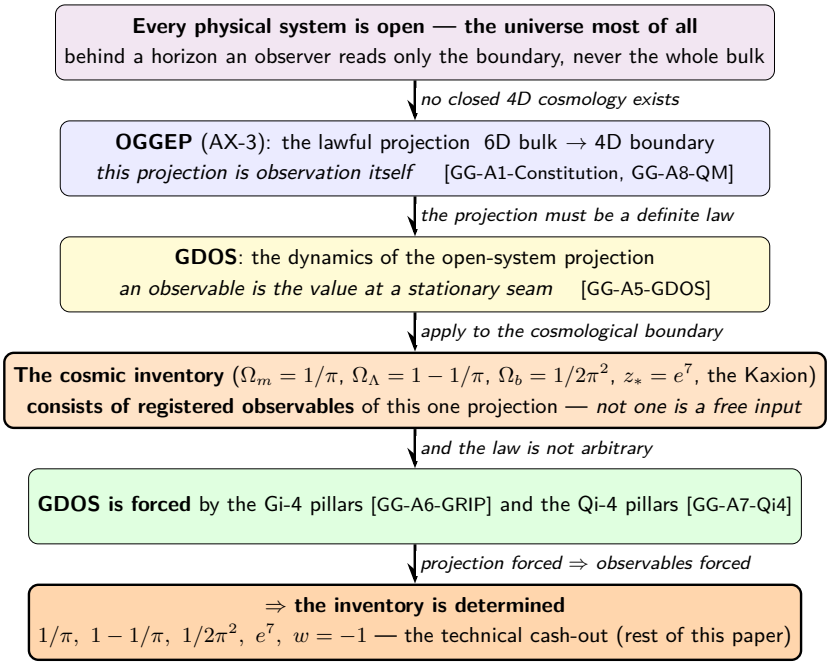

Figure 9: The central claim, as one deductive funnel. Because no physical system is closed — and the universe, seen from behind a horizon, least of all — the four-dimensional cosmos is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection is observation itself; its dynamics is the open-system law GDOS. The cosmological inventory is therefore a set of *registered observables* of the projection, not free inputs — and because GDOS is forced by the Gi-4 and Qi-4 pillars, those observables are determined. The closed forms derived in the rest of the paper (the matter fraction $1/\pi$, the dark-energy residue $1 - 1/\pi$, the baryon fraction $1/2\pi^2$, the recombination redshift e^7,[a] and $w = -1$) are the technical cash-out of this single recognition. The funnel is identical to the one that opens the companion particle-physics paper GG-A9-SM; only the payload differs, one story up.

[a] Prose that calls z_* "the recombination redshift" does so in the loose epochal sense — the epoch of the first atoms. The precise object is always the topological corridor-switch redshift $z_* = e^7$, distinct from the measured decoupling redshift $z_{\rm dec} = 1089.92 \pm 0.25$, with which it coincides to 0.6% (Part VI).

Every physical system is open — the universe most of all. No system in nature is truly closed. A particle is defined by its couplings to an environment that cannot, even in principle, be switched off; and the universe is the extreme case, not the exception. An observer inside a cosmological horizon can never access the whole of the bulk — only the part that has been projected onto the four-dimensional boundary they inhabit. What a telescope registers is, by its nature, a boundary readout of a larger geometry, never the geometry entire. The closed four-dimensional cosmology in which Ω_m, Ω_Λ, and the rest appear as free inputs is therefore an idealization — and it is precisely the idealization that makes those numbers look like quantities one is obliged to measure.

Openness forces a projection, and the projection is observation. If there is no closed four-dimensional theory, then the four-dimensional world an observer reports is the image of a deeper bulk under a lawful projection. This is the content of the third axiom of the GG-Constitution, the open-system geometric–gauge equivalence principle (OGGEP, AX-3; Arisaka, 2026, A1-Constitution): every admissible theory must supply a map from the six-dimensional bulk to the registered observables on its four-dimensional boundary. The decisive recognition — absent from a century of quantum mechanics — is that this bulk-to-boundary projection *is* observation itself. Measurement is not a postulate appended to a theory; it is what the OGGEP projection does, with the recovery of quantum mechanics carried out in GG-A8-QM (Arisaka, 2026, A8-QM) and its information-theoretic content in GG-A7-Qi4 (Arisaka, 2026, A7-Qi4).

The projection is a definite law: GDOS. OGGEP requires that a projection exist; GG-A5-GDOS makes it dynamical and unique. The geometric dynamics of open systems (GDOS; Arisaka, 2026, A5-GDOS) is the open-system law whose stationary, attractor configurations are the values an observer reads off. An observable, in this framework, is not a coefficient in a cosmological action; it is the value assigned at a stationary seam of a lawful open-system flow.

Therefore the cosmic inventory consists of observables — without exception. It follows that every entry of the cosmological inventory — the matter fraction Ω_m, the dark-energy fraction Ω_Λ, the baryon fraction Ω_b, the recombination redshift z_*, the dark-energy equation of state w, and the Kaxion that fills the dark-matter slot — is a *registered observable* of the same OGGEP/GDOS projection. Not one is a free input. The six-number "parameter space" of ΛCDM is not a space of freedoms at all; it is a list of measurement outcomes of one fixed geometry. This is an *ontological* statement, and it holds for the whole inventory regardless of how completely the present paper succeeds in *computing* each entry; the honest accounting of that progress is the provenance triage of §68, in which the entries not yet closed are observables whose determining projection has not been worked out in detail, not exceptions to the claim.

And the law itself is forced. Finally, GDOS is not an arbitrary open-system equation. It is the one consistent with the two four-pillar structures the constitution erects over it: the geometric pillars Gi-4 (TGG → CCI → GFF → GRIP; Arisaka, 2026, A6-GRIP) and the quantum-information pillars Qi-4 (Arisaka, 2026, A7-Qi4). Because the projection that produces the observables is itself forced, the observables it produces are determined.

Figure 9 draws this chain as the single deductive funnel it is, and its conclusion is what makes the rest of the paper possible. For the forty years that ΛCDM has measured the universe — and the century that quantum mechanics has described measurement without deriving it — the cosmic numbers have been treated as inputs to a closed four-dimensional theory, and no one asked whether they are instead observables of an open-system projection from a higher-dimensional bulk. That question, and its affirmative answer, is the move on which everything below depends: once it is granted, the matter fraction $1/\pi$, the dark-energy residue $1 - 1/\pi$, and the integer-ledger redshift e^7 cease to be separate mysteries and become readings of one ledger. The next section states those readings; the body of the paper earns them. The funnel is deliberately the same one that opens the companion particle-physics paper GG-A9-SM — the cosmos and the electron are read from one ledger, one story apart, and the whole GG-Theory cohort rests on this single recognition.

7 What this paper claims

The claim is bold, and it is meant to be, because the whole value of a zero-parameter theory is that it can be wrong in public. Stated plainly, GG-A10-Cosmos proposes the following.

The composition of the universe is a geometric figure. The matter fraction is the reciprocal of π; the dark-energy fraction is what remains; the baryon fraction is one part in $2\pi^2$:

$$\Omega_m = \frac{1}{\pi} \approx 0.318, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi} \approx 0.682, \qquad \Omega_b = \frac{1}{2\pi^2} \approx 0.051,$$

landing on the measured concordance values. The three great chapters of cosmic history share one origin: inflation is the first stretch of a geometric descent — a derived plateau one e-fold below the GUT edge, with the lock scale $M_{\rm lock} \approx 3.85\,\text{TeV}$ crossed promptly on the way down, the second break that stocks the dark sector; dark energy is the residual $1 - 1/\pi$ of the same closed inventory, with an equation of state $w = -1$; and dark matter is a specific particle, the Kaxion, with a definite mass, $m_{\rm Kaxion} \simeq 8.1\ \mu\text{eV}$, ringing in a narrow laboratory window near two gigahertz — one scale calibrated to the observed abundance and declared as such, the mass its lattice-QCD image. The two warring values of the Hubble constant are read as one derived rate and one predicted calibration systematic; the switch that separates the two dictionaries through which the rate is read is fixed at the integer-ledger redshift $z_{\rm sw} = e^7 \approx 1097$, with the surface of last scattering derived from the framework's own atomic chain 0.4% below it, on the measured value (§42.1).

None of this is offered as a story to be admired. Every one of these claims is staked on an experiment that can refute it within the decade — a primordial gravitational-wave measurement against the staked $r = 0.0038$, a dark-matter haloscope scan, a precision measurement of the dark-energy equation of state, where — against the recent DESI hint that dark energy may be evolving — GG-A10-Cosmos predicts $w = -1$ exactly, and stands or falls on it. GG-A10-Cosmos keeps ΛCDM as its boundary baseline, retains its equations and its observables, and aims to *derive the numbers* ΛCDM *can only measure*. Exactly how much of that program is proved today, how much rests on a stated conjecture, and how much remains open, is set out honestly in the codebook of Part IX and the critical discussion of Part XIII; the present Part states the claims, and the body earns them.

Before the body unpacks the apparatus, it is worth seeing the result itself. Figure 10 and Table 2 set the whole of GG-A10-Cosmos's output — every cosmological number the framework puts on the table — beside the measured universe, some two hundred pages ahead of where the full audit is developed. Each entry is shown as its fractional deviation from observation: filled markers are genuine zero-parameter predictions, open markers are quantities the framework currently inherits rather than derives, and the color records how much of each has been earned — Derived, Derived-conditional, or Asserted.

Reading the table. The twenty-two rows are not twenty-two *independent* numbers: the three inventory fractions are one Hopf-fibration form and sum to one, $\Omega_{\rm DM}$ is their difference, S_8 is σ_8 times a fraction, the ledger road's $\eta_B = 6.18 \times 10^{-10}$ is $\Omega_b h^2$ in other units (the row scores the mechanism's 4.89 against the datum), and r and n_s share one count — the honest count of independent predictions is about eight, and the rest are the same geometry read in other units, a consistency check and not additional evidence. The columns are the prediction in closed form; the observed value with its 1σ uncertainty; the relative deviation Dev $= (\text{pred} - \text{obs})/|\text{obs}|$; the *pull* $(\text{pred} - \text{obs})/\sigma$ with $\sigma = \sqrt{\sigma_{\rm obs}^2 + \sigma_{\rm th}^2}$, the theory tolerance $\sigma_{\rm th}$ printed with its source in the pull cell of every row that carries one; and the provenance triage (**D** = Derived; **D*** = Derived-conditional, closed modulo a named lemma; **A** = Asserted). Hubble constants are in km/s/Mpc: the background 66.8 ± 0.2 is derived with zero sky input by the geometric Route I (Section 29.8), scored against the CMB reading 67.4 and against the H0DN ladder consensus 73.50 ± 0.81, whose -9.1% is the predicted ladder systematic (Section 42.6); the four acoustic rows are CLASS-verified images of the same fractions; the z_* row scores the atomic-chain decoupling prediction (1091.5, zero fitted parameters, at the framework's own h) against the measured value, with the topological switch e^7 annotated in the plot, not scored (§42.1). The 'cal.' entries are the *epoch calibration*: under Planck central values the exact crossing $\Omega_m = 1/\pi$ occurred 75 Myr before the present ($\pm$160 Myr; §32.11), and since the matter fractions fall with time while Ω_Λ grows, that one declared input takes the four inventory deviations to (input, 0.00, +1.69, −0.31)%, every remaining inventory residual an epoch-independent ratio within about 1σ; the rows outside the inventory carry their own pulls, the ladder's -8 and $100\,\theta_*$'s -2.4 among them, and are not smoothed by this sentence. Every entry is reproduced and audited row by row as Table 41 in the codebook of Part IX. This is the bottom line stated up front; the body earns it.

Table 2: The complete numerical bottom line of GG-A10-Cosmos. Twenty-two comparisons — twenty-one codebook entries less the three unmeasured, plus four acoustic images of the same fractions — against their measurements: prediction, observed value, relative deviation, pull (with the theory tolerance $\sigma_{\rm th}$ and its source where one applies), and provenance grade (D derived, D* derived-conditional, A asserted); conventions in the text above, the audited copy in Table 41.

Entry	GG prediction	Observed	Dev	Pull	Grade
The π-locked inventory					
Ω_m	$1/\pi = 0.3183$	0.315 ± 0.007	+1.05% (cal. input)	+0.5 (—)	D*
Ω_Λ	$1 - 1/\pi = 0.6817$	0.685 ± 0.007	−0.48% (cal. 0.00)	−0.5 (0.0)	D*
w	−1 (exact)	-1.028 ± 0.032	+2.72%	+0.9	D*
Ω_b	$1/(2\pi^2) = 0.05066$	0.0493 ± 0.0006	+2.76% (cal. +1.69)	+2.3 (+1.4)	D*
$\Omega_{\rm DM}$	$1/\pi - 1/(2\pi^2) = 0.2676$	0.2657 ± 0.0070	+0.73% (cal. −0.31)	+0.3 (−0.1)	D*
Baryogenesis and nucleosynthesis					
η_B	$\approx 5 \times 10^{-10}$ (central 4.89, closed form on two states, K=35; the priced 1.9 is 2.2σ short, the time-resolved 0.58 a factor ten, App. I.6)	$(6.12 \pm 0.04) \times 10^{-10}$	−20%	−0.3 ($\sigma_{\rm th}$ 4×10^{-10}, asserted; $\ln 3$ in the log, the $\times 3^{\pm 1}$ band, gives −0.2)	D*
$^7{\rm Li}/H$	SBBN/$\pi = 1.59 \times 10^{-10}$	$(1.6 \pm 0.3) \times 10^{-10}$	−0.6%	0.0	D*
The Hubble constant, recombination, the acoustic spectrum					
$H_0^{\rm CMB}$	66.8 ± 0.2 (Route I, zero sky input)	67.4 ± 0.5	−0.89%	−1.1 ($\sigma_{\rm th}$ 0.2, Route I; 2025: −1.1)	D*
$H_0^{\rm late}$	66.8 ± 0.2 (one background; the ladder's excess a predicted systematic)	73.50 ± 0.81	−9.12%	−8.0 ($\sigma_{\rm th}$ 0.2, same)	D*
z_*	1091.5 (atomic chain at the derived h; switch e^7 annotated)	1089.92 ± 0.25	+0.15%	+0.8 ($\sigma_{\rm th}$ 0.18%, modeling)	D*
ℓ_D	1350 ± 3 (Boltzmann; in CLASS's $k_D D_M$ convention 1956 vs. 1947 ± 5, +0.5%, +1.6)	$\approx 1345 \pm 15$ (convention spread)	+0.37%	+0.3 ($\sigma_{\rm th}$ 3, Boltzmann)	D*
ω_b/ω_m	$1/2\pi = 0.15915$ (h-free, epoch-free)	0.15643 ± 0.0016	+1.74%	+1.7	D*
$100\,\theta_*$	1.0381 (r_s/D_M at the carried rate)	1.04109 ± 0.00030	−0.29%	−2.4 ($\sigma_{\rm th}$ 0.12% via z_*; −10 on $\sigma_{\rm obs}$ alone)	D*
peak ratio 2/1	0.449 (image of $1/2\pi$)	0.453 ± 0.008	−0.8%	−0.5	D*
peak ratio 3/1	0.439 (image of $1/2\pi$)	0.443 ± 0.009	−0.9%	−0.4	D*
The inflationary spectrum (the derived plateau; prompt lock crossing, $N_2 = 0$)					
n_s	0.9634 exact (N^*=52.6; 0.9620 at leading order)	0.9649 ± 0.0042	−0.16%	−0.4 ($\sigma_{\rm th}$ 0.00034, half e-fold)	D*
r	0.0038 exact (0.0043 at leading order)	< 0.036 (95%, BK18, BICEP/Keck Collaboration, 2021)	×9 below	—	D*
A_s	anchored (the declared input; warp cand. 75%)	$(2.10 \pm 0.04) \times 10^{-9}$	0	0	A
Clustering and radiation					
σ_8	0.80 ± 0.01 (from A_s on the π-locked background)	0.811 ± 0.006	−1.4%	−0.9 ($\sigma_{\rm th}$ 0.010, amplitude)	D*
S_8 (vs. CMB)	0.825 ± 0.010	0.832 ± 0.013	−0.8%	−0.4 ($\sigma_{\rm th}$ 0.010, same)	D*
S_8 (vs. cosmic shear)	the same 0.825 ± 0.010	$0.759^{+0.024}_{-0.021}$	+8.7%	+2.5 ($\sigma_{\rm th}$ 0.010, same)	D*
$N_{\rm eff}$	3.044 ($\Delta N_{\rm eff} \leq 0.03$, derived null)	2.99 ± 0.17	+1.8%	+0.3 ($\sigma_{\rm th}$ +0.027, one-sided)	D

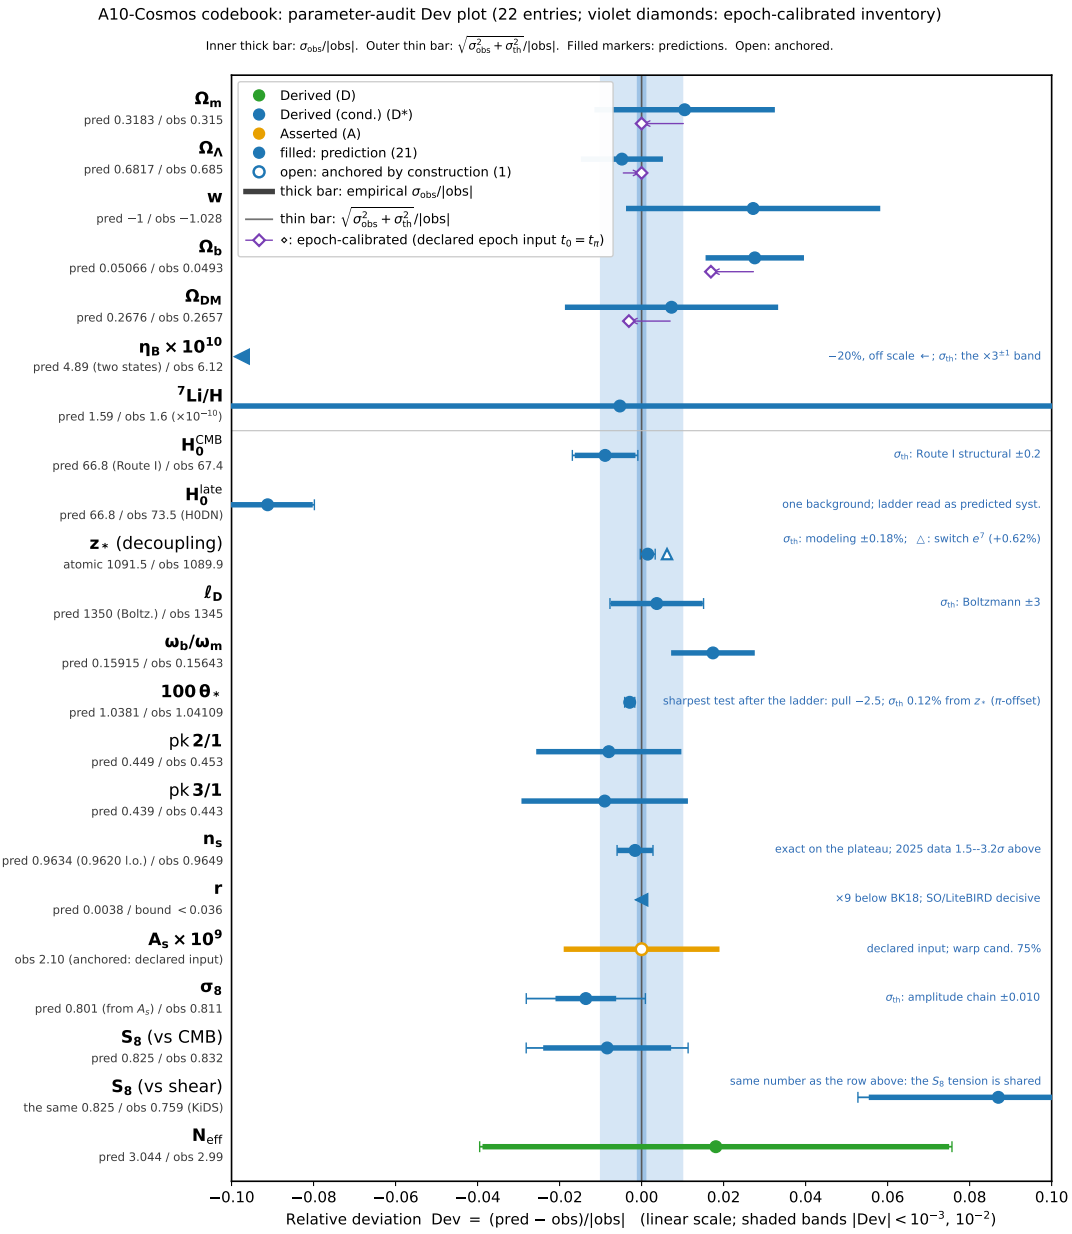

Figure 10: A10-Cosmos parameter-audit Dev plot. Fractional deviation $(\text{pred} - \text{obs})/|\text{obs}|$ for the twenty-two rows carrying an empirical comparison, with nested precision bands at $|\text{Dev}| < 10^{-3}$, 10^{-2}. *Filled* markers are predictions (twenty-one, of which about eight are independent); *open* markers are anchored by construction (Dev = 0). Color encodes status: green = Derived, blue = Derived (cond.), gold = Asserted. Inner thick bar: empirical $\sigma_{\text{obs}}/|\text{obs}|$; outer thin bar: combined $\sqrt{\sigma_{\text{obs}}^2 + \sigma_{\text{th}}^2}/|\text{obs}|$ — entries carrying a theory tolerance σ_{th} are annotated inside the panel. *Violet open diamonds*: epoch-calibrated positions of the four π-inventory rows (single declared input $t_0 = t_\pi$; see the table caption). The radiation entry is shown as the total N_{eff}. Full data, sources, and uncertainties: Table 2. Reproduced and audited entry by entry as Figure 31 in the codebook of Part IX.

A few entries show why this is more than a tidy refit of familiar numbers. ΛCDM *measures* the matter fraction Ω_m and the dark-energy fraction Ω_Λ: they are dialed until the model matches the sky, and the model is silent on why they take the values they do. Here the same two numbers take a *forced* coupled form under one closed geometry — $\Omega_m = 1/\pi$ and $\Omega_\Lambda = 1 - 1/\pi$ — landing within a fraction of a percent of Planck at the present epoch, so that the cosmic-coincidence problem ΛCDM files away is not patched but *demoted*: the geometry removes the 10^{120} vacuum-energy fine-tuning outright and leaves only a mild, few-percent matter of timing, honestly retained rather than declared an identity. The Hubble tension is the second example: where ΛCDM must pick a single H_0 and live with a six-sigma split between the early- and late-universe readings, the geometry *derives* the background rate outright — 66.8 ± 0.2 with zero sky input, 66.6 ± 0.2 from the π-locked inventory against the measured densities, coinciding with the CMB value — and reads the local-ladder excess whole as a calibration systematic the framework predicts and does not derive, so the "tension" is relocated from cosmology to metrology rather than patched with an extra parameter. And where ΛCDM leaves dark matter an unfilled slot and the dark-energy equation of state a fitted constant, GG-A10-Cosmos names the dark matter as a definite particle — the Kaxion, near 8.1 μeV — and predicts $w = -1$ exactly. That is the difference on display: not better fits to numbers already known, but numbers the geometry is *obliged* to produce — and that a single clean measurement can therefore prove wrong.

The reader who wants the derivations should turn to the codebook of Part IX, where this figure and table reappear (as Figure 31 and Table 41) with every entry audited row by row; the reader who wants the bottom line can read it here. After all, a zero-parameter theory is judged first by the numbers it dares to put on the table — so here they are.

8 Scope, the cohort, and how to read this paper

The plan of the body is plotted in Figure 12, in the roadmap subsection that closes this Part (§9.2).

A word on scope, on where this paper sits, and on how to read it.

GG-A10-Cosmos is a presentation of cosmology through one geometric lens; it is not a rejection of standard cosmology. It keeps the standard observables — $H(z)$, the CMB spectra C_ℓ, the matter power spectrum $P(k)$, the growth rate $f\sigma_8$, the density fractions, the light-element abundances — deliberately front and center, precisely so that it remains clear how the framework can fail. The novelty lies not in new equations but in a constrained source ledger and a single inherited interpretation of the sectors ΛCDM leaves open.

The paper is the cosmological member of a twenty-two-paper program. The reader does not need to have read the others: this Part and the body are self-contained, and each inherited result is named where it is used. In the program's architecture, GG-A10-Cosmos sits one story above the Standard Model of particle physics (the companion GG-A9-SM), applying the same constitution and the same six-dimensional geometry to the universe as a whole; the full cohort map is laid out in the section that closes this Part (Section 9.1). Different readers may travel differently: a working cosmologist may go straight to the codebook of Part IX and the falsification ledger of Part XIII; a

curious newcomer may follow the cosmic story from the beginning through Parts III–VIII; a skeptic may begin with the Common Objections. All paths meet at the Conclusions.

Throughout, the load-bearing items of the program carry cohort-global *Prefix-Suffix-N* identifiers — theorems T-Cosmos-N, innovations I-Cosmos-N, objections O-Cosmos-N, falsifiers F-Cosmos-N, and so on — so that any result can be cited from outside its home paper; the full scheme and the twenty-two-paper suffix list are tabulated in Appendix D.

9 The cohort inheritance and the roadmap of the seventeen Parts

Before the inherited framework is laid out (Part II), it is useful to lay out two reader-orientation maps on the same opening: a panorama of the twenty-two-paper cohort in which this paper stands (§9.1, Figure 11), and a panorama of the seventeen Parts of the present paper (§9.2, Figure 12). A reader of any later Part should be able to flip back to this Section and orient themselves in both directions: which cohort companion carries which imported result, and where the present argument stands in the seventeen-Part arc.

9.1 The cohort inheritance

GG-Theory is not a single paper but a connected program of twenty-two. They share one root, one geometry, and one descent: each paper hands the next a smaller, sharper set of conclusions, so that by the time cosmology is reached almost nothing is left to choose. The cleanest way to meet the program is to look at the whole of it at once. Figure 11, which closes this subsection, draws all twenty-two papers in one frame — and a reader new to GG-Theory can read its architecture off the figure from top to bottom — the constitutional root, the bulk trio, the open-system boundary law and its two pillars, the three textbook readouts, the capstone, the two phenomenology branches, and the Foundations Branch on the left flank — before following the thread of GG-A10-Cosmos itself.

Table 3: The twenty-two-paper GG-Theory cohort. Each paper's main achievement and its relevance to the present cosmology (GG-A10-Cosmos). Every paper name links to that paper's own page at preprints.arisaka-gg.org, where all twenty-two are collected.

Paper	Main achievement	Relevance to this cosmology (GG-A10-Cosmos)
GG-A1-Constitution	Four axioms and seven postulates fixing what counts as an admissible claim — the constitution of the whole program	Prevents cosmology from reintroducing hidden free sectors, undefined observables, or unconstrained late-time phenomenology
GG-A2-GG6	Derivation of the six-dimensional GG-6 bulk, BI-6, OGN-5, compact holonomy, and the unification of all six forces	Fixes the bulk geometry, anomaly structure, Wilson-line setting, and arithmetic ledger that constrain every cosmic sector

continued on next page

Table 3, continued.

Paper	Main achievement	Relevance to this cosmology (GG-A10-Cosmos)
GG-A3-BI6	Complete proof of BI-6 anomaly freedom, unified gauge consistency, and conservation of the dual current	Closes the ultraviolet legality question and certifies that the projected cosmic ledger may exist globally and locally
GG-A4-Euclid	The historical lineage from Euclid to GG-6, and the zero-parameter integer OGN-5 ledger behind the constants	Grounds the arithmetic — the endpoint $K = 35$ and the $1/\pi$ inventory — that this paper reads off without tuning
GG-A5-GDOS	OGGEP descent to boundary quantum mechanics, boundary general relativity, and formal 4D QFT: the GDOS boundary law	Establishes that four-dimensional GR is already derived, so FRW cosmology belongs inside the chain, not outside it
GG-A6-GRIP	The runtime constitution — GRIP, the force law $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$, and the system instance packet SIP	Provides the open-system runtime and phase-selection grammar for inflation, recombination, structure formation, and horizons
GG-A7-Qi4	The Qi-4 pillars — SeOS, QCI, the identity IH + OH = EIC — and the five-story informational pagoda	Provides the selector architecture, scale diagnostics, and hologram logic for the higher-floor reading of cosmic evolution
GG-A8-QM	Quantum mechanics derived from GG-Theory — the four-dimensional reformulation of measurement and state	Supplies the boundary-QM observation layer beneath horizon physics and the observer-projection reading of the cosmos
GG-A9-SM	The Standard Model as a theorem — the descent of all twenty-eight Standard-Model parameters	Fixes the particle ledger, electroweak sector, and neutrino block feeding dark matter, baryogenesis, and thermal history
GG-A10-Cosmos *(this paper)*	ΛCDM as the cosmological-boundary projection — the π-locked inventory $\Omega_m = 1/\pi,\ \Omega_\Lambda = 1 - 1/\pi$	The present paper: the cosmological-boundary member of the consequence trio, one story above GG-A9-SM
GG-A11-Origin	The closing capstone — the universal identity and the GRIP cascade across cosmos, life, and mind	Receives this paper's origin and why-now result and closes the A-series self-referential loop
GG-P1-KK	The Kaluza–Klein lock cluster near $M_{\text{lock}} \approx 3.85\,\text{TeV}$, with explicit LHC and future-collider falsifiers	Supplies the lock scale M_{lock} that enters this paper's H_0 anchor and the KK line of the codebook
GG-P2-N3	A three-leg derivation of $N = 3$ fermion generations from BI-6, MSMF, and the integer ledger	Fixes the generation count and neutrino content behind N_{eff} and the cosmic thermal history

continued on next page

Table 3, continued.

Paper	Main achievement	Relevance to this cosmology (GG-A10-Cosmos)
GG-P3-Electron	The electron as the ground state of nature — the six oldest electron questions, closed	Sibling consequence of the same ledger; anchors the charged-lepton stratum inherited here through GG-A9-SM
GG-P4-Real	The two real sources of the boundary's complex structure — positive frequency is holomorphy, from the compact phase circle and the causal clock	Foundations Branch; supplies the origin of the derived complex structure this paper leans on at the observer-indexed vacuum and at the dial-free door (T-Cosmos-18)
GG-P5-CPT	C, P, and T as orientation reversals of the three real bulk factors — why every particle has a twin, and why the universe is made of matter	Foundations Branch; the GG-A9-SM–GG-A10-Cosmos interface routing the inherited ledger phases into this paper's baryon-asymmetry chain (Part V, Appendix I)
GG-C1-Kaxion	Cold dark matter as the non-spatial compact-angle pseudoscalar Kaxion at 8.1 μeV	Identifies this paper's dark-matter component and its microelectronvolt mass on the same S^1_θ winding
GG-C2-Hubble	The Hubble tension read against one derived background	The dedicated observational study of the one background: the ladder read as one derived rate and one predicted calibration systematic, with the ladder-channel falsifiers
GG-C3-BH	Black-hole entropy $A/4$ recovered as a boundary theorem, Derived-conditional on three named premises, within GG-Theory's holographic ledger	Extends the same boundary law to horizon thermodynamics, consistent with this paper's holographic reading

The cohort falls into five strata, read top to bottom. At the head sits the constitutional root, GG-A1-Constitution (Arisaka, 2026, A1-Constitution), whose four axioms and seven postulates fix once and for all what counts as an admissible claim — the rule that forbids cosmology from quietly smuggling in a free sector later on. The root compiles into the *GG-6 trio*: GG-A2-GG6 (Arisaka, 2026, A2-GG6) derives the six-dimensional bulk, GG-A3-BI6 (Arisaka, 2026, A3-BI6) proves its BI-6 anomaly freedom, and the historical companion GG-A4-Euclid (Arisaka, 2026, A4-Euclid) supplies the integer ledger and lineage. From the trio descends the third stratum, the open-system boundary law GG-A5-GDOS (Arisaka, 2026, A5-GDOS), organized by its two pillar papers, GG-A6-GRIP (Arisaka, 2026, A6-GRIP) and GG-A7-Qi4 (Arisaka, 2026, A7-Qi4), which between them supply the runtime and selector architecture. The fourth stratum is the *consequence trio* that recovers textbook physics from this single source: quantum mechanics in GG-A8-QM (Arisaka, 2026, A8-QM), the Standard Model in GG-A9-SM (Arisaka, 2026, A9-SM), and the ΛCDM cosmology of the present paper, GG-A10-Cosmos. The capstone GG-A11-Origin (Arisaka, 2026, A11-Origin) closes the A-series, proving the universal identity and the GRIP cascade that the whole program uses. Hanging off

this spine are two phenomenology branches and one foundations column: the Predictions Branch — GG-P1-KK (Arisaka, 2026, P1-KK), GG-P2-N3 (Arisaka, 2026, P2-N3), GG-P3-Electron (Arisaka, 2026, P3-Electron) — and the Cosmology & Astrophysics Branch — GG-C1-Kaxion (Arisaka, 2026, C1-Kaxion), GG-C2-Hubble (Arisaka, 2026, C2-Hubble), GG-C3-BH (Arisaka, 2026, C3-BH); and on the spine's left flank sits the Foundations Branch — GG-P4-Real (Arisaka, 2026, P4-Real) and GG-P5-CPT (Arisaka, 2026, P5-CPT), read below.

This fixes where the present paper sits. GG-A10-Cosmos is the cosmological-boundary member of the consequence trio, one story above the Standard Model of GG-A9-SM: it reads ΛCDM as the cosmological-boundary projection of the very same six-dimensional bulk, inheriting the boundary law of GG-A5-GDOS and the GRIP cascade of GG-A6-GRIP. Cosmology, in other words, is not a fresh start. It is the top floor of a building whose foundation, frame, and lower stories are already standing.

The Foundations Branch. The map's left flank carries the cohort's third branch: the two member papers whose subject is the boundary's complex structure and its mirrors, a single column beside the spine. GG-P4-Real (Arisaka, 2026, P4-Real) is the origin-of-i member paper: it derives the boundary's complex structure from two real structures of the GG-A2-GG6 arena — the compact phase circle and the causal clock — and this paper's own derived reading of positive frequency at the observer-indexed vacuum (§12.1), curved-space residue and all, is that account read on the cosmological stage. GG-P5-CPT (Arisaka, 2026, P5-CPT) is the GG-A9-SM–GG-A10-Cosmos interface of the map: the ledger phases derived in GG-A9-SM — the PMNS Dirac phase and the Majorana parity $\alpha_{31} = \pi$ — reach this paper's baryon-asymmetry chain through it, and the η_B edge of the map is exactly the import that Part V and Appendix I consume, at the grade recorded there: derived-conditional on the named finite-temperature hypothesis of Appendix I. Both carry P-series designators; their place at the spine's flank, not the phenomenology rows, marks their foundational role.

The twenty-two-paper map, identical in every paper of the cohort, is Figure 11, which closes this subsection.

9.2 Roadmap of the seventeen Parts

The rest of the paper is organized into fifteen further Parts, and the walk through them is plotted in Figure 12. The present Part — the entrance — is nearly done: it has asked the oldest questions precisely, established the ΛCDM concordance baseline with its seventeen open problems, and laid out the two orientation maps of this closing section. Two Parts then put the machine on the table. Part II, the inherited framework, hands the reader the four ideas the cosmology leans on — the six-dimensional bulk, the boundary projection, the zero-parameter integer ledger, and the open-system runtime — each proved upstream and imported wholesale. Part III sets the cosmological stage: the lawful beginning at the GUT edge rather than the Planck scale, the inherited scale chain, general relativity and the Friedmann equations as theorems of the boundary law, and the radial chronology that orientates the arrow of cosmic time.

The two Parts that follow are the first half of the derivation core, and they are highlighted in the

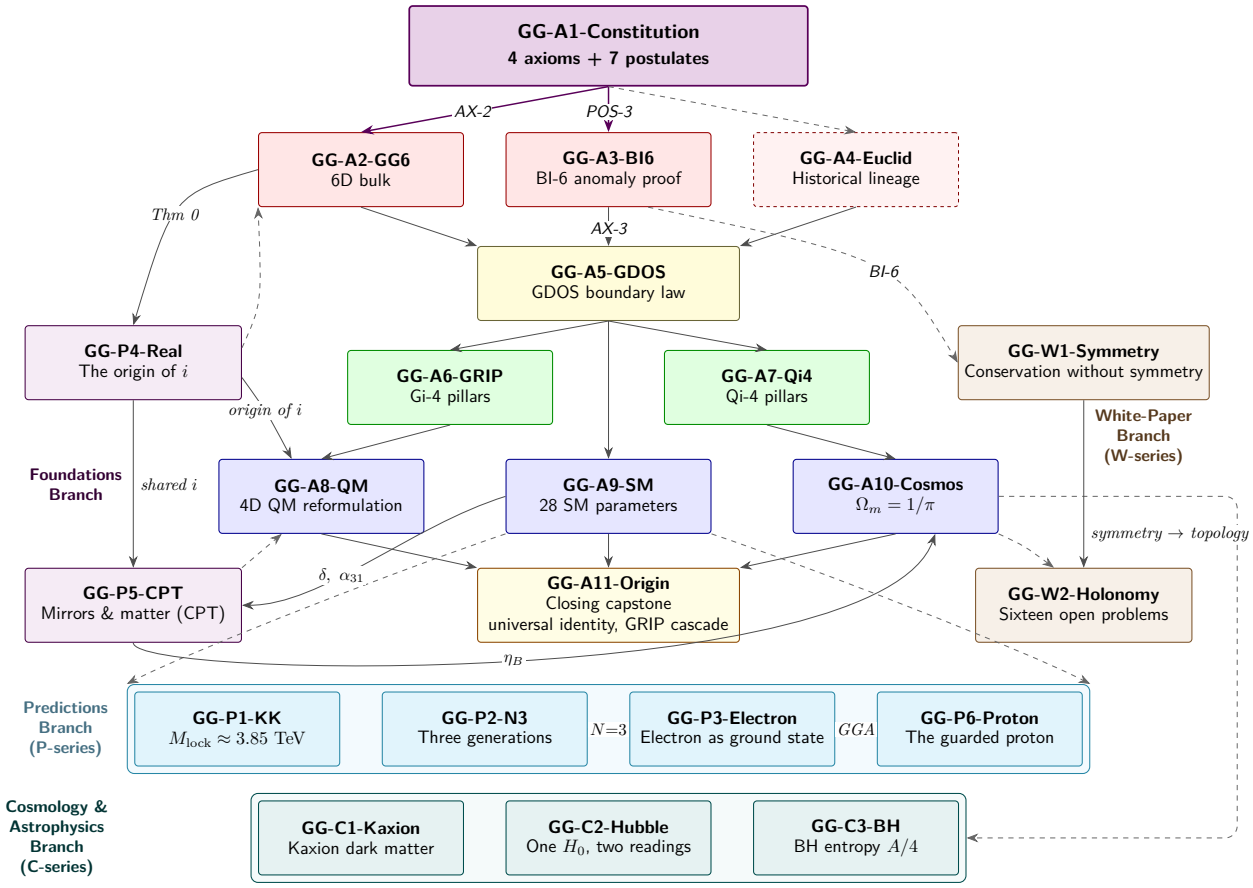

Figure 11: The twenty-two-paper cohort map of GG-Theory. The constitutional root GG-A1-Constitution (violet, top) supplies the four axioms and seven postulates from which the rest of the cohort descends: the GG6 Trio (GG-A2-GG6, GG-A3-BI6, GG-A4-Euclid, red) realizes the bulk geometry; GG-A5-GDOS (yellow) the boundary law; the GDOS Trio pillars (GG-A6-GRIP, GG-A7-Qi4, green) the geometric and quantum-information runtime; the 4D-Physics Trio (GG-A8-QM, GG-A9-SM, GG-A10-Cosmos, blue) the textbook readouts; and the closing capstone GG-A11-Origin (gold) the quarks-to-consciousness identity cascade. Two framed branches carry the phenomenology — the **Predictions Branch** (cyan), whose four papers hand their results along the row by the internal arrows (GG-P2-N3's $N{=}3$ into GG-P3-Electron, and that paper's ground-attractor method, GGA, on into GG-P6-Proton (Arisaka, 2026, P6-Proton)), and the **Cosmology & Astrophysics Branch** (teal). The spine is flanked by two columns of two, which do opposite jobs. On the left the **Foundations Branch** (violet) feeds it: GG-P4-Real, fed by GG-A2-GG6's Theorem 0 arena, manufactures the i that GG-A8-QM consumes, and GG-P5-CPT is the GG-A9-SM–GG-A10-Cosmos interface, fed by the ledger phases and feeding the η_B chain across the width of the map, with the shared-i edge descending solid between them. On the right the **White-Paper Branch** (brown) reads it: GG-W1-Symmetry (Arisaka, 2026, W1-Symmetry) walks the century from Noether to Hopkins–Singer, and GG-W2-Holonomy (Arisaka, 2026, W2-Holonomy) dissects the structure that arc arrives at — protection by topology rather than by symmetry — across sixteen open problems. The white papers derive nothing new, carrying every result at its parent's declared grade; the single dashed feed drawn into each names its heaviest source, not its only one. *Reading conventions:* solid arrows mark axiom-forced descent and carry an edge label naming the driving axiom or postulate where most illustrative; dashed arrows mark the deployment of derived results, or historical links inherited without being constitutionally forced.

roadmap because they deliver the values of the π-locked inventory in cosmic history. Part IV runs the launch and the terminus of the one open-system descent — inflation as the opening of the cosmic GRIP descent, with the derived plateau class, the computed tilt $n_s = 0.9634$, and the tensor target $r = 0.0038$, and dark energy as the ground-state terminus of the same descent, $\Omega_\Lambda = 1 - 1/\pi$. Part V populates the matter side of the ledger: Kaxion cold dark matter from the compact-angle sector with the π-locked relic abundance $\Omega_m = 1/\pi$, baryogenesis through the inherited neutrino-and-CP ledger, and Big Bang nucleosynthesis with the lithium-7 audit.

Three Parts then confront the derived universe with the sky. Part VI takes up the growth of structure and the present-day tensions — the CMB baseline and the scalar/tensor audit, the Hubble constant as one derived background with the distance ladder as the predicted outlier, the S_8/σ_8 matter-clustering question, and early-galaxy formation in the JWST era. Part VII extends the same boundary law to black holes and primordial relics — the information paradox under OGGEP/GDOS, primordial black holes as a feature-dependent supplementary channel, and the relict primordial gravitational-wave background. Part VIII gathers the whole history into the evolutionary chronology of the universe: the dictionary translating between t, μ, r, a, and z, and the master chronology table from the GUT boundary to the asymptotic de Sitter future.

Part IX is the paper's discipline: the cosmological codebook, in which every observable of the π-locked inventory is graded **Derived**, **Derived (conditional)**, or **Asserted**, each conditional entry tied to a named lemma. Part X, the third highlighted Part, is the topological completion — *why the universe has the inventory it does*: the closed-packet derivation of $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, and $\Omega_b = 1/(2\pi^2)$, and the universal-cosmic-identity reading that explains, layer by layer, why the cosmological coincidence is not a coincidence.

The remaining Parts weigh and close the case. Part XI is the assembly — the direct derivation of ΛCDM by GG-Theory, gathered in one place: the zeroth and the six silent axioms of the concordance model replaced by theorems and graded conditionals, the export chain walked in eight stages from the bulk action to the expansion history, and the assembled model displayed with the register of every entry and no fitted continuous parameter. Part XII is the Summary — the cross-paper coherence map, the comprehensive first-order codebook, and the second-order residual audit. Part XIII, the Discussions, states the significance, the comparisons with other cosmological programs, the predictions for falsification, and the open questions stated as open. Part XIV, the Conclusions, returns at the end of the climb to the summit view: the composition of the universe as geometry, not coincidence. The back matter is the workshop. Part XV collects the Technical Appendices — a graduate-level course in cosmology read through GG-Theory; Part XVI answers the Common Objections a working cosmologist, an inflation theorist, or a philosopher of cosmology would press hardest; and Part XVII closes with the annotated References, each carrying its Relevance note. Read in order, the seventeen Parts compose a single self-contained argument; read by need, each Part names its imports and can stand a visit on its own.

One thread runs through all seventeen problems, and it is worth stating in plain words before the formal work begins. Every one of them is, at bottom, the same question asked of a different number: *why this value, and not another?* A science that can only measure must treat that question as meaningless; a science that can derive must treat it as the whole point. The wager of this paper

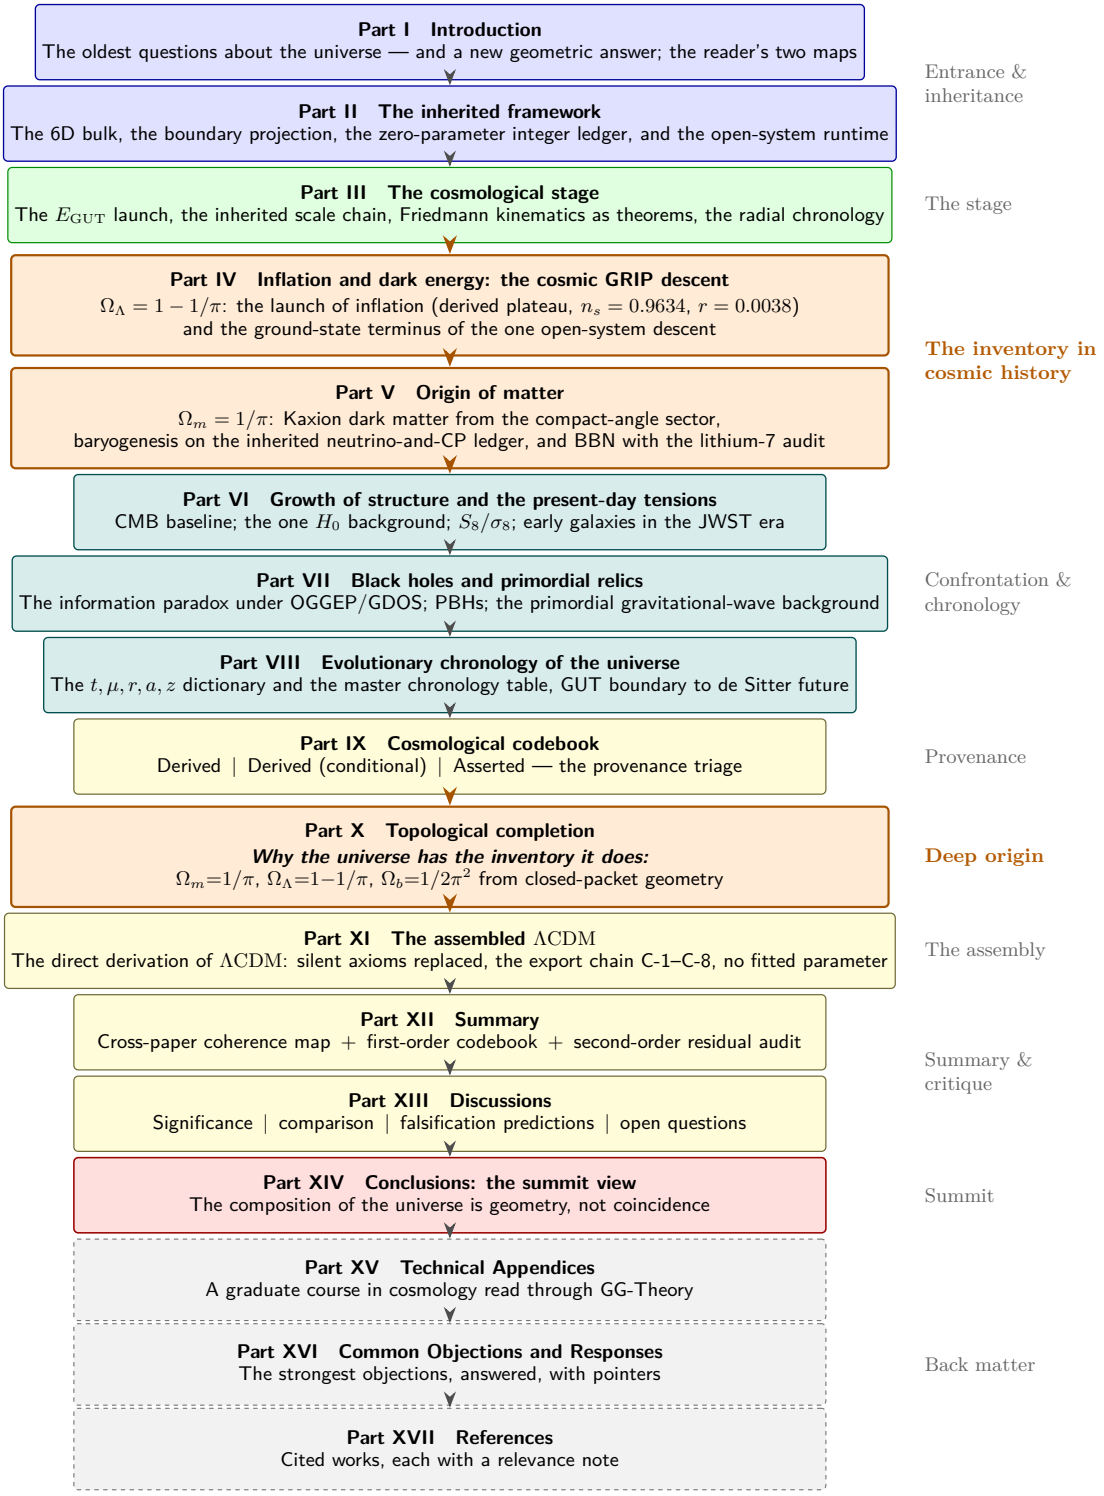

Figure 12: Roadmap of the seventeen Parts of GG-A10-Cosmos. The Parts run top to bottom; the sidebar names the ten phases. Parts IV, V, and X are highlighted because together they carry the π-locked cosmic inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$: derived in cosmic history in Parts IV–V, given its geometric origin in Part X, and assembled into the concordance model in Part XI. The full walk is Sec. 9.2.

is that cosmology is ready to move from the first kind of science to the second — and that the price of the move, paid in full in Part IX, is saying exactly which answers are earned and which are still owed.

This Part has done its work: it has established the concordance baseline, posed the open questions, and placed the paper within its twenty-two-member program. The body answers those questions, problem by problem, on a single inherited ledger; and the Conclusions return, at the end of the climb, to the view from the summit — the universe as geometry, not coincidence.

10 Part I in plain English: what the standard model of the universe does and does not explain

The reason this Part spends so long on what ΛCDM gets right is that the argument of the paper depends on it. A framework that offered a new account of the cosmic inventory while quietly needing the standard model to be wrong somewhere would be easy to dismiss and would deserve to be. The claim here is the harder one: that the concordance model is an excellent description, that its six fitted parameters are measured to a precision that is one of the real achievements of the century, and that the question it does not answer is why those six numbers have the values they do.

That question is not a complaint about precision. It is a question about provenance. A parameter that is measured is a parameter that could have come out otherwise, and the standard model has no machinery that would have objected. What follows proposes that the inventory is not a set of measurements at all but a set of open-system observables — quantities forced by the geometry that produces the boundary on which they are measured — and that if this is right, the correct description of a cosmic number is a closed form rather than a fitted value.

A reader should hold two things in mind through the rest of the paper. The first is that a closed form which agrees with a measurement to a few percent is a claim about structure, not a claim about precision, and it must never be sold as the second. The second is that this Part has stated the problem in the strongest available form: seventeen open problems, each of them a real one, and each of them a place the framework can be wrong. The Parts that follow are answerable to that list.

Part II
The inherited framework: GG-Theory in one Part

What the cosmology inherits — the six-dimensional bulk, the boundary projection, the zero-parameter integer ledger, and the open-system runtime

---

This Part answers a single question: *what is GG-Theory, and what does a cosmology built on it inherit?* Part I established the ΛCDM baseline, posed the open questions of cosmology, and mapped the twenty-two-paper program this cosmology belongs to; before we can claim geometric answers to those questions, the reader deserves a clear and unhurried picture of the machine the answers come out of. Nothing here is new physics. Everything in this Part was proved upstream, in the companion papers of the cohort, and is imported into GG-A10-Cosmos wholesale. The job of this Part is to make that inheritance *intuitive* — to hand a reader who has seen only Part I the four ideas that the rest of the paper leans on, in the order in which they build on one another. The actual cosmological history — the beginning at the GUT edge, the Friedmann kinematics, the expansion and cooling of the universe — is set up in the next Part, once the framework is in hand.

11 One geometry, no dials: the six-dimensional bulk

At the center of the whole program is a single geometric object — the constitution of GG-A1-Constitution (Arisaka, 2026, A1-Constitution) fixes what may count as an admissible claim, and GG-A2-GG6 (Arisaka, 2026, A2-GG6) derives the geometry itself — and it is worth meeting it in plain language before meeting it in equations. GG-Theory posits that reality is a *six-dimensional bulk*, a tapering "throat" of the form $\mathcal{M}_4 \times I_r \times S^1_\theta$. Three pieces make it up. The familiar four dimensions $\mathcal{M}_4$ are the spacetime we live in. A single radial direction I_r runs along the throat from a wide ultraviolet mouth down to a narrow infrared end; this is the energy axis, and it is best read as the renormalization-group scale, $r = \ln(E_{\mathrm{GUT}}/\mu)$ — moving along r is nothing other than running a coupling from high energy to low. And a small compact angle S^1_θ, the phase fiber, carries the locked gauge structure: its winding is where the dark sector will eventually live.

The word "warped" deserves a warning, because it is borrowed from a different picture. Here it does *not* mean a Randall–Sundrum extra spatial dimension (Randall and Sundrum, 1999) with a curved metric in the usual sense. It means RG running along r: the warp factor is the field-theory / holographic-RG reading of scale dependence, exactly as the rest of the cohort reads it. With that one substitution, the geometry stops being exotic and becomes familiar — "deeper into the throat" simply means "lower in energy."

The single most important intuition to carry forward is this: *our four-dimensional world is the boundary of this bulk.* We do not live inside the throat; we live on its edge, and what we call physics

is the shadow the bulk casts there. How that shadow is cast — the boundary projection itself — is the subject of the next section, where Figure 13 draws the throat, its radial scale axis, its compact fiber, and the four-dimensional boundary slab on which ordinary cosmology lives.

The geometry, made explicit. Behind the picture is a definite, parameter-free metric. In the natural units of the program the bulk line element is

$$ds_6^2 \;=\; e^{-2r}\,\eta_{\mu\nu}\,dx^\mu dx^\nu \;+\; dr^2 \;+\; d\theta^2, \qquad r \in [0, K], \quad \theta \in [0, 2\pi),$$

with $\eta_{\mu\nu}$ the Minkowski metric of the four-dimensional boundary (Arisaka, 2026, A2-GG6). The warp profile $A(r) = r$ is not chosen but forced, by demanding constant bulk curvature with a flat boundary ($A''(r) = 0$); there is no curvature scale, no compactification radius, and no brane tension to tune. The radial coordinate is literally the renormalization-group logarithm, $r = \ln(E_{\mathrm{GUT}}/\mu)$, running from the ultraviolet edge at $r = 0$ ($E_{\mathrm{GUT}} \simeq 2.3 \times 10^{16}\,\mathrm{GeV}$) down to the integer endpoint $r_{\mathrm{IR}} = K = 35$; the Planck–electroweak hierarchy is then the topological number e^{-K}, not a fine-tuned ratio.

Six dimensions are four plus one plus one. The count is forced role by role (Arisaka, 2026, A2-GG6): four Lorentzian directions $\mathcal{M}_4$ carry registered observations and the only length dimension; one non-compact *scale* direction I_r is the geometrized RG flow (a dimensionless energy logarithm, not space); and one compact *phase* circle S^1_θ carries integer-quantized Wilson-line holonomy — the dark sector lives in its winding. The three bulk directions are exactly what is needed to absorb the three dimensionful constants c, G_N, and $\hbar$; the boundary supplies the remaining two, k_B and the energy–information quantum $\mathsf{Q_B}$. This "three-plus-two" architecture is the spine of the whole ledger.

One geometry carries all six forces. On the bulk the gauge and gravitational structure is a single group,

$$G_{\mathrm{Bulk}} \;=\; \mathrm{Spin}(1,5) \;\times\; SU(5) \;\times\; U(1)_B \;\times\; U(1)_{Qi},$$

with the Lorentz/spin factor $\mathrm{Spin}(1,5)$ (the gravity sector) on equal footing with the internal factors (Arisaka, 2026, A2-GG6). Here $SU(5)$ is the minimal simple unification containing the Standard Model; $U(1)_B$ is gauged baryon number, which turns proton stability into a theorem; and $U(1)_{Qi}$ is the informational gauge force that the Qi-4 pillars later use. Projected onto the four-dimensional boundary this becomes the six forces an observer actually measures — strong, weak, electromagnetic, baryonic, informational, and gravitational — with three abelian invariances surviving after electroweak breaking. Crucially for cosmology, the same compact circle S^1_θ that quantizes these holonomies is where the Kaxion dark matter and the π-locked inventory are read off: the cosmos and the particle ledger are two projections of this one object, one story apart.

12 The boundary projection: how four-dimensional physics arises

If the bulk is the whole of reality, why do we only ever measure four-dimensional physics? The answer is the *boundary law*, and it is the deepest single result the cosmology inherits. An observer behind a cosmological horizon can never access the closed bulk; what a telescope registers is the

projection of the bulk onto the four-dimensional edge the observer inhabits. GG-A5-GDOS (Arisaka, 2026, A5-GDOS) raises this to a principle — the open-system geometric–gauge equivalence principle (OGGEP), the third axiom of the constitution — which requires every admissible theory to supply an explicit dictionary

$$\mathcal{D}_{\mathsf{OGGEP}} : \ (\mathcal{H}_{\mathrm{Bulk}}, \rho_{\mathrm{Bulk}}, U_{\mathrm{Bulk}}) \longmapsto (\mathcal{H}_{\mathrm{Bdy}}, \rho_{\mathrm{Bdy}}, \Phi_{t,s}, \{\mathcal{I}_\alpha\}),$$

carrying the inaccessible bulk data to a boundary state, an admissible (completely positive, trace-preserving) evolution channel $\Phi_{t,s}$, and the instruments that read it. The decisive recognition is that this bulk-to-boundary projection *is observation itself*: measurement is not a postulate appended to the theory but what the projection does.

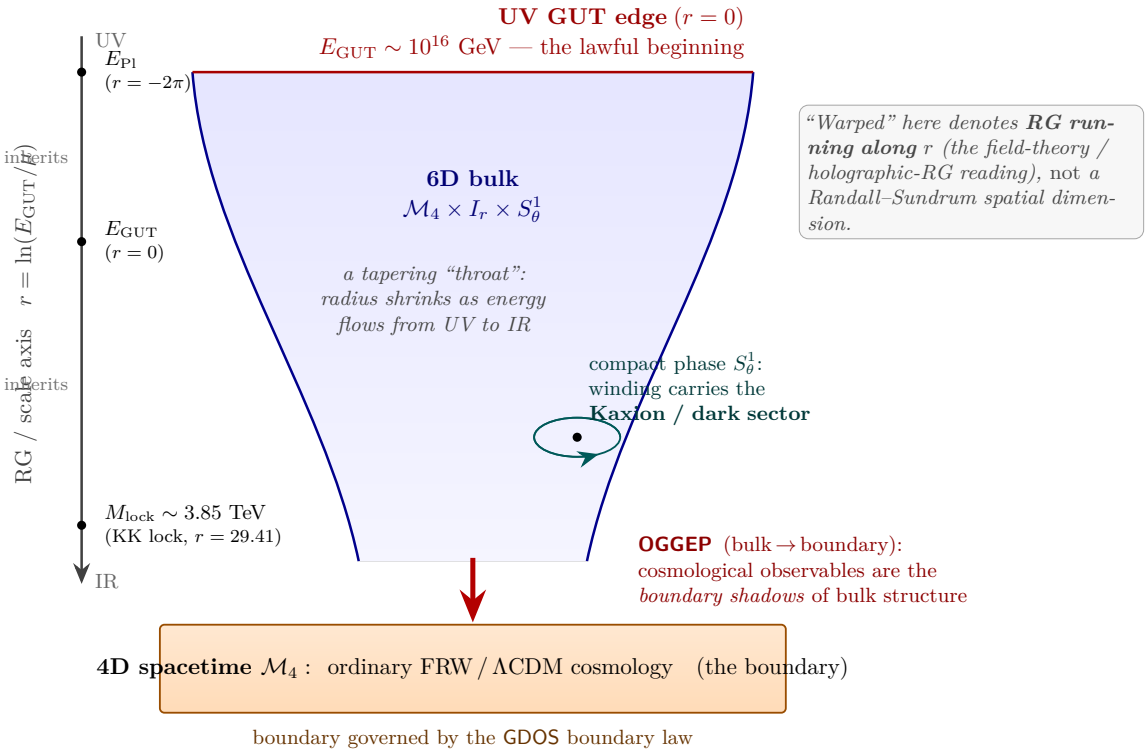

Figure 13: Schematic of the six-dimensional bulk $\mathcal{M}_4 \times I_r \times S^1_\theta$ and its OGGEP projection to the four-dimensional cosmological boundary. The radial axis $r = \ln(E_{\mathrm{GUT}}/\mu)$ is the renormalization-group / scale axis: the ultraviolet GUT edge sits at the top ($r = 0$, $E_{\mathrm{GUT}} \sim 10^{16}$ GeV, the lawful beginning) and low-energy physics lies toward the boundary, with the inherited scale chain $E_{\mathrm{Pl}} \to E_{\mathrm{GUT}} \to M_{\mathrm{lock}}$ ($M_{\mathrm{lock}} \sim 3.85$ TeV, the Kaluza–Klein lock scale) marked along it. The compact angle S^1_θ is the phase fiber whose winding carries the Kaxion / dark sector. The bold arrow is the OGGEP bulk-to-boundary map: cosmological observables on $\mathcal{M}_4$ (ordinary FRW/ΛCDM) are boundary shadows of bulk structure, with the boundary governed by the GDOS boundary law. "Warped" denotes RG running along r (the field-theory / holographic-RG reading), *not* a Randall–Sundrum spatial dimension.

GG-A5-GDOS then makes the projection dynamical and unique. Its boundary law, the geometric

dynamics of open systems (GDOS), is the Lindblad/GKLS generator

$$\mathcal{L}_{\mathsf{GDOS}}(\rho) \;=\; -\,i\,[H_{\mathrm{eff}},\rho] \;+\; \sum_{\alpha}\gamma_{\alpha}\Big(L_{\alpha}\,\rho\,L_{\alpha}^{\dagger} - \tfrac{1}{2}\{L_{\alpha}^{\dagger}L_{\alpha},\,\rho\}\Big), \qquad \gamma_{\alpha}\geq 0,$$

derived — not postulated — from the exact bulk dynamics by a Born–Markov reduction (Arisaka, 2026, A5-GDOS). The ordinary Schrödinger equation is recovered as its closed-system limit; openness is primary and closure the special case. In this language an *observable* is the value an open system reads off at a stationary seam, an attractor of the GDOS flow — and from this one law GG-A5-GDOS recovers four-dimensional quantum mechanics, general relativity, and quantum field theory as boundary theorems rather than separate assumptions.

The warm consequence, the one to hold onto, is that *ordinary FRW cosmology is recovered, not discarded.* Because four-dimensional general relativity is already a theorem of GG-A5-GDOS, the Friedmann–Lemaître–Robertson–Walker universe lives inside the framework automatically, with the four-dimensional Newton constant emerging as a scale-overlap integral over the throat (the overall unit set by G_N, with the bulk mass the derived combination, as in GG-A2-GG6; the explicit reduction to the Friedmann equations is in Part III and Appendix E). We do not choose between GG-Theory and the cosmology textbooks: the textbook equations are exactly what the boundary law produces, and what the framework adds is a fixed, parameter-free ledger of the source terms that feed them.

A boundary law fixes *what* the four-dimensional theory is; a cosmology also needs to know *how* its lawful possibilities are explored and stabilized over cosmic time. That is the geometric runtime, the four invariance pillars Gi-4 of GG-A6-GRIP (Arisaka, 2026, A6-GRIP), organized as a single cascade

$$\mathsf{TGG} \implies \mathsf{CCI} \implies \mathsf{GFF} \implies \mathsf{GRIP}.$$

The total geometric–gauge object TGG packages the closed bulk ledger as one composite connection on the G_{Bulk} bundle; the covariant conservation identity CCI hands the boundary observer a local conservation law and, under closure, forces the integer ledger of the next section; the geometric free-energy functional GFF = GEP + CEP ranks the admissible boundary configurations by a single scalar; and GRIP — the lifecycle DGI → GSSB → GA, lawful exploration to symmetry-breaking to stabilized attractor — fixes the order in which a branch is explored, committed to, and frozen. The next two sections put this machinery to work: first the fixed ledger it runs on, then the way it selects the cosmic branches that actually happen.

12.1 Which particles? The observer-indexed vacuum, and horizon thermality as openness

One further consequence of the projection deserves a named place before the ledger, because cosmology uses it on every page that mentions particle production. In an expanding or horizon-bearing spacetime there is famously no preferred vacuum: inequivalent positive-frequency conventions define inequivalent particle notions, related by Bogoliubov mixing. In the present framework this ambiguity has a fixed address. On the canonical bulk geometry inherited from (Arisaka, 2026,

A2-GG6) the constitutional time is globally static, so the bulk-level positive-frequency structure is unique; what varies from observer to observer is the *dictionary*: which modes an instrument counts is fixed by the observer's own time and horizon structure, and the instruments are part of the OGGEP boundary package by definition. "How many particles" is therefore an observer-indexed boundary observable, not a bulk quantity — the same status the framework already assigns to boundary states themselves. The physical output of the ambiguity is produced by the GDOS projection itself: tracing the inaccessible partner of a horizon's squeezed pairing yields an exactly thermal state (at $T = a/2\pi$ for the Rindler pairing), so horizon thermality is constitutional openness doing its ordinary work rather than new physics; and cosmological particle creation is the same structure run dynamically — a Gaussian pumping channel occupying the same slot as the dissipative generators of the boundary law (Arisaka, 2026, A5-GDOS). The full statements are carried in the real-bulk treatment of (Arisaka, 2026, A8-QM), whose standalone, full-length development is the Foundations Branch member paper GG-P4-Real (Arisaka, 2026, P4-Real). Two boundary cases are closed there at the model level, as entries R14–R15 of the series: the rotating case, where the ergoregion defeats the co-rotating selection and the uniqueness anchor moves to the asymptotic frame, superradiance becoming a symplectic pumping channel in the same GDOS slot; and the collapse epoch, which factorizes into a universal thermal channel times a decaying transient (the full-length treatment in the black-hole setting is Arisaka, 2026, C3-BH). What remains conjectural is the exact curved-space existence theory — the field's own, imported where used. The present paper uses only the two facts just stated: the bulk structure is unique, and the thermal content of horizons is the partial trace.

13 Zero free parameters: BI-6 and the integer ledger

A framework can be elegant and still be useless if it leaves the numbers free. What makes GG-Theory predictive rather than merely descriptive is that it has *no continuous dials* to turn, and two inherited results enforce this.

The first is the anomaly-locked bulk. Global consistency of the six-force closure requires a modified six-dimensional Bianchi identity, BI-6, for the antisymmetric tensor $H_3 = dB_2$ (Arisaka, 2026, A3-BI6; Arisaka, 2026, A4-Euclid):

$$dH_3 \;=\; \operatorname{tr} F_{SU(5)}^2 \;-\; F_B^2 \;-\; \tfrac{2}{5}\, F_B \wedge F_{Qi} \;-\; \tfrac{2}{5}\, F_{Qi}^2 \;+\; \operatorname{tr} R^2,$$

whose coefficient vector $(a, b, c, d) = (1, -1, -\frac{2}{5}, -\frac{2}{5})$ is fixed uniquely by three independent conditions — local closure ($dX_4 = 0$), global integrality of $[X_4]$, and anomaly cancellation by inflow through the single channel $I_8 = \frac{1}{2} X_4 \wedge X_4$. There is no continuous deformation: the trace-ratio lock $|a/c| = |a/d| = 5/2$ is rigid. The ultraviolet is legal, and it is legal in exactly one way.

The second is the integer ledger. GG-A4-Euclid (Arisaka, 2026, A4-Euclid) proves that the same closure forces a unique tuple of whole numbers, the OGN-5 ledger

$$(M, N, L;\, J, K) \;=\; (1, 3, 5;\, 7, 35),$$

through a five-step Diophantine cascade in which each integer is fixed by the one before: $M = 1$

(a single Green–Schwarz tensor channel), $L = 5$ (the minimal simple unification $SU(5)$), $N = 3$ (the three colors, forced by the dual-lock closure $L(2N+1) = M^2 + N^2 + L^2$, which at $M = 1$, $L = 5$ reads $N^2 - 10N + 21 = 0$ and admits $N = 3$ alone, the root 7 exceeding the block it sits in — GG-A3-BI6 §*OGN-5 as a theorem derived from BI-6*, Step 3), $J = 2N + 1 = 7$, and the closure integer

$$K \;=\; LJ \;=\; M^2 + N^2 + L^2 \;=\; 35.$$

Two remarks on N. The BI-6 trace ratio $-(L-N)/L = -2/5$ is GG-A3-BI6's *independent check* of the result, not its route; and the generation count $N_g = 3$ is a separate quantity the ledger does not fix (GG-P2-N3). That two independent arithmetic expressions for K agree — the *dual lock* — is what closes the ledger; it is generically false, and true here only because both are forced by the same axioms. These integers are not fitted; they are the arithmetic of the geometry. Together with the structural constants π and e and the dimensionful conversion factors $\{c, G_N, \hbar; k_B, \mathsf{Q_B}\}$, they are the entire stock from which every dimensionless number in the cohort is built: the same ledger that fixes the twenty-eight Standard-Model parameters of GG-A9-SM (Arisaka, 2026, A9-SM) fixes the cosmic inventory of this paper, with no new freedom in between. The topological step that turns K, the integers, and π into the π-locked inventory ($\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/2\pi^2$) is given in full in Part X and Appendix L, and the discipline that admits a number only if it is a cost-minimum element over this ledger is the MSMF/OGN principle. Every cosmological number must come out of this fixed ledger, or it is not allowed in. Tellingly, the same $K = 35$ that sets the scale-throat depth anchors the Kaxion's cap-loading ladder $\Xi_{\mathrm{cap}} = K + \Sigma$ (an ansatz, Part V): the Planck–electroweak hierarchy and the dark-matter window are read from the same integer. Exactly what each completed preprint contributes to this inherited ledger is tabulated alongside the cohort map in §9.1 (Table 3).

14 The runtime: how cosmic branches are selected

The framework so far is static: a bulk, a boundary law, a fixed ledger. But a cosmology is a *history* — it has to choose, epoch by epoch, which branch of the lawful possibilities actually happens. That choosing is what GG-Theory calls the runtime, and it is the last piece the cosmology inherits. The runtime has two components, both proved in the pillar papers GG-A6-GRIP and GG-A7-Qi4: the GRIP cascade, which fixes the order in which lawful branches may be explored, committed to, and stabilized; and the GFF ledger, the scalar ranking functional $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ that says *which* branch is favored. Together they select the cosmic branches that matter — inflation's exit, recombination, the growth of structure, and the final attractor the universe relaxes into. The GFF ledger is, in form, a free-energy functional, standing in the information-theoretic lineage of Shannon (1948), Jaynes (1957), Kullback and Leibler (1951), Landauer (1961), Amari and Nagaoka (2000), and Friston (2010); ranking branches by such a functional is what turns an arbitrary cosmic narrative into a definite branch-selection problem.

The runtime has two regimes, and the distinction is the one that matters most for a cosmology (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A7-Qi4). *GRIP-In* is descent on a *fixed* landscape: the system rolls down the GFF valleys by gradient flow and tunnels between basins by Bogomolnyi-

saturated instanton steps of action πk, halting uniquely at the deepest reachable attractor. *GRIP-Out* is the complementary event that *reorganizes the landscape itself* — it opens new strata and hands the next descent its starting point. The slogan of GG-A6-GRIP is exact: GRIP-Out opens the landscape; GRIP-In finishes the journey. Cosmic history is precisely this alternation. The scale-setting events — the exit from the GUT edge, the electroweak lock, the launch and end of inflation — are GRIP-Out: they build the landscape. Everything that then settles onto a fixed landscape — the recombination seam, the relaxation to the dark-energy attractor, the frozen present-day constants — is GRIP-In. Where the cosmic landscape itself came from is the GRIP-Out question this paper answers; the particle paper GG-A9-SM answered the complementary GRIP-In question, reading the Standard-Model constants off the landscape those same events had already built.

Which branch, among the many lawful attractors of a stratum, is the one actually realized? The selector is the set of four quantum-information pillars Qi-4 of GG-A7-Qi4 (Arisaka, 2026, A7-Qi4). A configuration is the observed *ground geometric attractor* only if it passes all four admissibility filters at once: *Qi-QUA*, that it carry the locked four-current identity $J_E/\mathsf{Q}_\mathsf{B} = J_e = J_B = J_{Qi}$ (the Arisaka equation); *Qi-EP*, that it be a genuine extremum of the one projected variational principle, not merely a local minimum of a coarse-grained potential; *Qi-CON*, that its scale signatures be those of a fixed point, surviving transport along r; and *Qi-HAL*, that the runtime can actually halt there, with no admissible descent remaining. The observed value is the unique four-certificate survivor of its stratum — the same criterion that makes every electron, proton, and Kaxion the identical object everywhere.

This is not an abstraction without a worked precedent: it is exactly how GG-A9-SM (Arisaka, 2026, A9-SM) fixed the Standard Model. There each parameter is a coordinate of a geometric attractor reached by this cascade — the electroweak vacuum selected as the stationary point of a Wilson–Hosotani holonomy potential (the GSSB step), the constants read off only after the system has descended onto that branch — and the abstract "selection" is made concrete as a single MSMF/OGN cost-minimum over the integer ledger, the unique admissible element of lowest arithmetic cost. GG-A10-Cosmos runs the identical machinery one story up: the cosmic branches — inflation's exit, recombination, the growth of structure, and the final attractor the universe relaxes into — are the GRIP readouts of the same runtime, now on the cosmological boundary stratum.

The whole inheritance can be stated once, as a single descent ladder:

$$\mathsf{BI\text{-}6} \implies \mathsf{TGG} \implies \mathsf{CCI} \implies \mathsf{GFF} \implies \mathsf{GRIP}.$$

Read left to right: the anomaly-locked bulk fixes the legal ultraviolet stage; TGG packages the total closed cosmic ledger; CCI hands the four-dimensional observer a local boundary conservation law; GFF ranks the admissible branches of that boundary theory; and GRIP fixes the order in which they are explored, selected, and stabilized. Figure 14 traces this same descent for the cosmological observables, and the inheritance theorem below states it as a rule.

The same content, sharpened to the cohort's seven formal inputs and stated as a theorem about admissibility, reads as follows. It is the load-bearing rule of the whole paper: a cosmic source term is admissible only if it can be traced backward through the same constitutional chain already used to build the particle and open-system papers (Arisaka, 2026, A1-Constitution; Arisaka, 2026, A2-GG6;

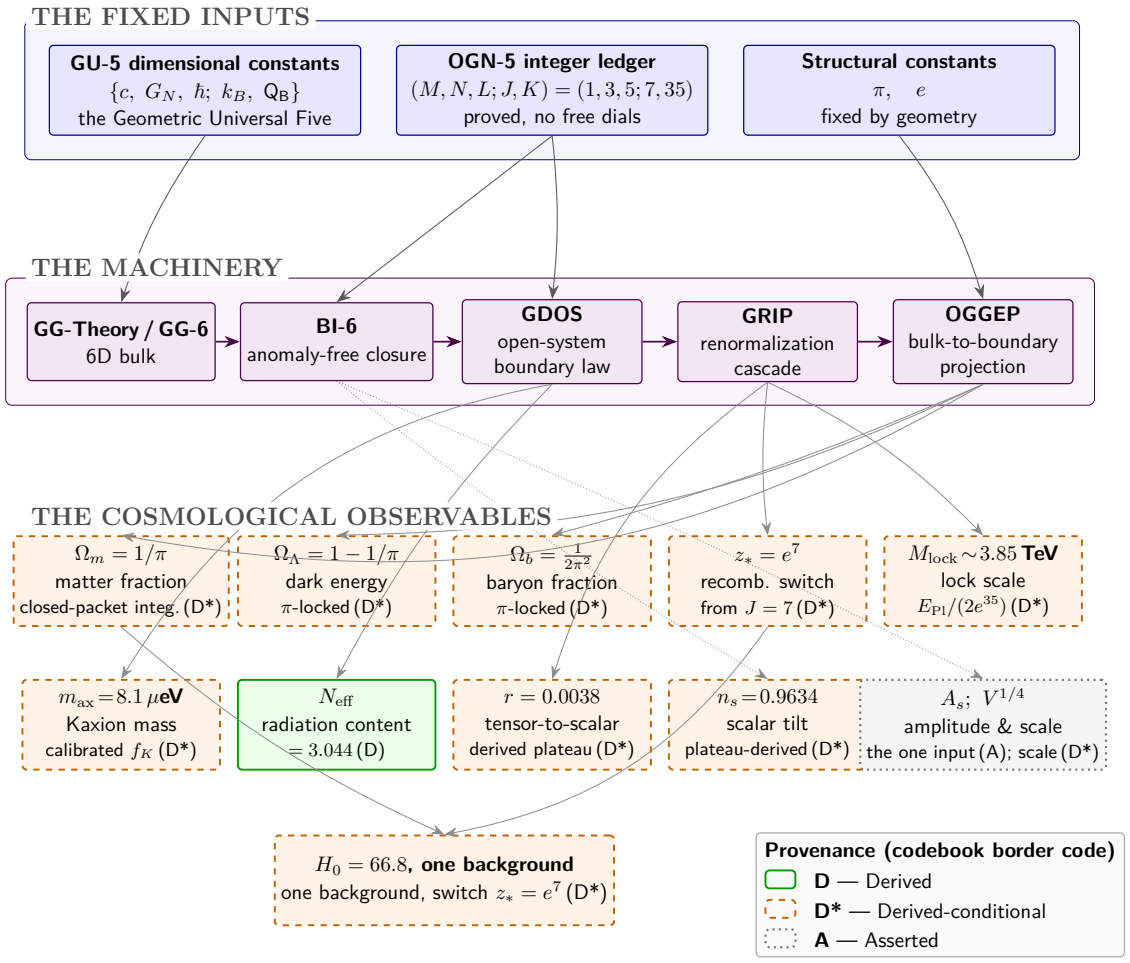

Figure 14: The zero-parameter derivation tree. Every cosmological observable (bottom tier) is traced back, through the GG-Theory machinery (middle tier), to a *fixed* input ledger (top tier) that contains *no continuous free parameters*: the GU-5 dimensional constants $\{c, G_N, \hbar; k_B, \mathsf{Q_B}\}$, the proved OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, and the structural constants π and e. Output nodes are border-coded by codebook provenance, and the colors are *not* cosmetic. *Solid green (D, Derived):* the radiation content $N_{\rm eff}$, a derived null. *Dashed amber (D*, Derived-conditional):* the π-locked fractions, the corridor-switch $z_* = e^7$, the lock scale $M_{\rm lock}$, the Kaxion mass, the one background rate H_0, and — computed on the derived plateau of Part IV — the tilt $n_s = 0.9634$, the tensor ratio $r = 0.0038$ and the inflationary scale $V^{1/4} = 8.1 \times 10^{15}$ GeV. *Dotted gray (A, Asserted):* the scalar amplitude A_s — the launch sector's one declared input, its warp candidate flagged. **Honesty caveat:** the dashed (D*) entries are *conditional on named inputs*—notably Conjecture VIII, whose closed-packet integration now closes *up to the declared epoch input* (the residual uniformity lemma exact to one part in 10^{15})—so the tree displays the derivation flow with that single calibration input, not an unconditional proof for the D* (or A) entries. Only the solid-green (D) nodes are claimed as fully derived.

Arisaka, 2026, A3-BI6; Arisaka, 2026, A5-GDOS; Arisaka, 2026, A6-GRIP; Arisaka, 2026, A7-Qi4; Arisaka, 2026, A9-SM).

Theorem 14.1 (Cosmological inheritance theorem)**.** *Every admissible cosmological source term in GG-A10-Cosmos must descend in the order*

$$\text{BI-6} \Longrightarrow \mathsf{TGG} \Longrightarrow \mathsf{CCI} \Longrightarrow \mathsf{GFF} \Longrightarrow \mathsf{GRIP} \Longrightarrow \textit{parameter block.}$$

No cosmic module is allowed to enter the Friedmann, perturbation, thermal, or information-processing sectors unless it survives each previous layer.

Proof. BI-6 fixes the legal ultraviolet stage and its coefficient structure. TGG packages the total closed cosmic ledger. CCI gives the local boundary conservation law seen by the four-dimensional observer. $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ then ranks the admissible branches of that boundary theory, and GRIP fixes the order in which those branches can be explored, selected, and stabilized. The remaining quantity is therefore not a primitive fluid or parameter, but a branch export. This is exactly the sense in which the present framework treats cosmological observables as theorem candidates rather than as a phenomenological menu. ■

Theorem 14.2 (Cosmological inheritance theorem (T-Cosmos-1))**.** *Every admissible cosmic sector in GG-A10-Cosmos must be expressible as a boundary theorem of the seven-paper chain in Table 3. In particular, no new free cosmological degree of freedom may be inserted unless it can be shown to descend from the GG-6/BI-6 bulk, survive the OGGEP/GDOS projection, enter the GFF ledger, and fit into an admissible GRIP factorization.*

Proof. The proof is a bookkeeping proof. GG-A1-Constitution fixes admissibility. GG-A2-GG6 and GG-A3-BI6 fix the legal six-dimensional stage and its anomaly structure. GG-A5-GDOS fixes what counts as an admissible boundary observable and an admissible reduced dynamics. GG-A6-GRIP fixes the runtime template $\mathsf{GRIP} = (\mathsf{DGI} \to \mathsf{GSSB} \to \mathsf{GA})$, which it states as universal and does not prove unique, and the scalar ranking functional $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$. GG-A7-Qi4 fixes the Qi-runtime, the selector architecture, and the hierarchy language by which internal reconstruction and higher-floor organization are read. GG-A9-SM fixes the particle ledger that every cosmological source term must now inherit. Any proposed cosmic module that cannot be placed into this chain is therefore not derived; it is merely appended. Such a module is inadmissible by construction. ■

With the descent ladder fixed, the remaining piece of runtime vocabulary is the *grammar* by which cosmic branches are actually selected — the part of the inheritance that decides how a given epoch transition is read, and where the cosmic history is allowed to come to rest.

14.1 Why GRIP and Wilson-loop GSSB now control the cosmology paper

GG-A10-Cosmos inherits not merely a list of sectors but a runtime grammar. The relevant grammar is GRIP, the Gauge-Runtime Invariance Pipeline,

$$\mathsf{GRIP} := (\mathsf{DGI} \to \mathsf{GSSB} \to \mathsf{GA}),$$

with the scalar ranking functional

$$F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}.$$

In the cosmology paper this matters because every important epoch is a branch event of one sort or another: inflationary launch and exit, relic selection, matter-radiation equality, recombination, halo collapse, black-hole Page turnover, and the late-time vacuum inventory all involve lawful exploration, branch commitment, and attractor stabilization. The universe is therefore not merely expanding; it is passing through a sequence of GRIP events.

One refinement of the terminus, established at runtime level in GG-A6-GRIP (Arisaka, 2026, A6-GRIP) and generalized across the layers of the program-wide cascade in the capstone GG-A11-Origin (Arisaka, 2026, A11-Origin), organizes the cosmological reading and is used throughout this paper. A GFF landscape generally admits many stable basins — many Geometric Attractors (GA's) — but on a stratum of the class GG-A6-GRIP calls admissible exactly one of them is the global minimum, the *Ground* Geometric Attractor,

$$\mathrm{GGA} := \arg\min F_{\mathsf{GFF}},$$

fixed by the landscape alone. (The four pillars of GG-A7-Qi4 characterize that object rather than select it; nothing in the cosmological reading below turns on them.) The distinction is exactly the one cosmology needs in order not to conflate two very different kinds of "attractor export." A *transient* cosmic history — a reheated hot patch, a frozen last-scattering record, a slice of the growth history — is a genuine GA: a stable record the projection can read, but not the deepest state of its sector, and not eternal. A *persistent relic* or final state — a dark-matter quantum that must survive to today, the asymptotic vacuum the universe relaxes into — is a GGA: the ground state of its stratum, with no deeper basin to decay to. The same theorem that makes every electron identical and the proton stable (GG-A9-SM) therefore underwrites, on the cosmological boundary, both the longevity of the dark matter (§35.1) and the finality of the late-time vacuum (Part IV): each is a GGA, not merely a GA.

A second piece of runtime vocabulary, defined at runtime level in GG-A6-GRIP (its GRIP-In/GRIP-Out regime definitions, together with the necessity of the second regime) and developed as the Lyapunov descent theorem T-GRIP-In-Lyapunov-Descent and the seventeen-event theory of GG-A11-Origin, is imported here because cosmology is where it pays the largest rent. The runtime operates in exactly two regimes, split by one question: *does the landscape change?* GRIP-In is the descent on a *fixed* landscape, halting uniquely at the stratum's GGA; GRIP-Out is the transient out-of-equilibrium event that *reorganizes the landscape itself* and supplies the initial datum of the next descent. Cosmic history is this alternation writ large, and the classification is sharp. The landscape-reorganizing transitions — the inflationary exit, reheating, the electroweak selection, QCD confinement, recombination — are GRIP-Out events: after each, the boundary carries strata and basins that did not exist before. The long epochs *between* transitions are GRIP-In: relaxation on the landscape the last transition built, with Hubble friction itself serving as the cosmological within-basin flow (the matter strata use the weak channel for the same job; GG-A9-SM). Crossovers such as matter–radiation equality reorganize nothing and are therefore *milestones of GRIP-In*, not

transitions. One word in that classification is fixed here, because the cohort's papers use it two ways. In GG-A6-GRIP's vocabulary a GSSB passage is a saddle crossing of GRIP-In on a landscape that does not move, and a GRIP-Out is a reorganization of the landscape itself — a bifurcation, not a stratum jump. Recombination is both in this paper's reading, and the two parts are named: the change of dictionary at $z_* = e^7$ is the GRIP-Out, and the passage of the plasma from the tightly coupled DGI phase to free streaming is the GSSB crossing on the landscape that switch makes. GG-A6-GRIP names the electroweak selection as its one GRIP-Out; GG-A11-Origin's event table lists cosmic cooling through the electroweak scale and through the Rydberg scale (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A11-Origin). The three lists are cross-referenced here and not reconciled, and one thing none of the three papers has shown is stated as such: that any of the transitions listed changes the discrete data of the stratum it acts on, which is what would make it a GRIP-Out in GG-A6-GRIP's strict sense rather than a GSSB passage on one landscape. And the chain halts exactly twice: the Kaxion relic and the de Sitter vacuum are the two points where cosmic GRIP-In terminates — the two GGA's of the preceding paragraph, now read dynamically. In one sentence: GRIP-Out opens the landscape; GRIP-In finishes the journey; and the universe we observe today is the journey, almost finished.

The electroweak paper already identified Wilson-line holonomy as the cleanest prototype of geometric spontaneous symmetry breaking (Aharonov and Bohm, 1959; Wilson, 1974; Hosotani, 1983; Hosotani, 2005; Arisaka, 2026, A6-GRIP; Arisaka, 2026, A9-SM). The cosmology paper now generalizes the lesson: the physically relevant branch variable is not necessarily a local scalar field chosen ad hoc, but can be a gauge-invariant loop or holonomy variable inherited from the compact GG-6 geometry. That observation narrows the early-universe story sharply. It means that phase selection in cosmology should be read, whenever possible, as Wilson-loop or holonomy selection on the GFF landscape, not as arbitrary scalar-field patchwork.

The cosmic specialization of this grammar — that every cosmological phase change is a lawful GRIP event, ranked by F_{GFF}, with persistent relics and final states singled out as Ground Geometric Attractors — is stated as the second cosmic theorem.

Theorem 14.3 (Cosmic GRIP theorem (T-Cosmos-2))**.** *Every cosmological phase change discussed in GG-A10-Cosmos must be readable as a GRIP event on the boundary state manifold, ranked by F_{GFF}, with GSSB realized by loss of stability of a previously admissible symmetric branch and with the selected post-transition history stabilized as a Geometric Attractor (GA) — and, in the case of a persistent relic or a final state, as the stratum's* Ground *Geometric Attractor (*GGA*; §14.1). In regime language: every such phase change is a GRIP-Out event — a transient reorganization of the boundary landscape — and every inter-transition epoch is GRIP-In relaxation on the landscape so created; the chain halts where GRIP-In terminates, at the two* GGA *endpoints.*

Proof. This is immediate once the seven-paper chain is accepted. GG-A5-GDOS provides the lawful boundary state manifold and its reduced dynamics. GG-A6-GRIP proves that any admissible runtime factorizes as GRIP and that branch ranking is governed by F_{GFF}. GG-A7-Qi4 extends this runtime with Qi-selector structure but does not alter the constitutional factorization. Therefore every cosmic epoch transition can be analyzed by the same three-stage logic: exploration of admissible channels, loss of stability of the symmetric branch at a threshold, and stabilization of the realized history

as an attractor — a GA for a transient epoch record, the stratum's GGA for a relic or final state that must persist (GG-A11-Origin (Arisaka, 2026, A11-Origin)). The regime clause follows from T-GRIP-Out-Necessity (GG-A11-Origin): gradient flow on a fixed landscape cannot cross a stratum boundary, so each genuine transition must be a landscape reorganization (GRIP-Out), and everything between transitions, being gradient flow plus its instanton stitching on a fixed landscape, is GRIP-In by definition (the GRIP-In definition of GG-A6-GRIP). ■

15 What this buys cosmology

Pulling the five ideas together gives the discipline that governs the rest of the paper. The bulk supplies one geometry; the boundary law makes ordinary cosmology a projection of it; the integer ledger leaves no continuous dials; and the runtime selects which lawful history happens. The single operative consequence is a rule with no exceptions: *every cosmological source term must descend the inheritance chain* — bulk, projection, ledger, runtime — and no new dial may be introduced along the way. A quantity that cannot be placed on this chain is not derived; it is merely appended, and is inadmissible by construction.

Table 4 makes the discipline concrete, mapping each open cosmological problem of Part I to the inherited components most relevant to it. It is a planning map, not a victory lap: it shows where each answer must come from, not that each answer is already complete.

Table 4: Working map from each of the seventeen open problems (Table 1) to the most relevant inherited components from the upstream papers (GG-A1-Constitution–GG-A6-GRIP, with the particle-sector ledger of GG-A9-SM where noted). This is a planning table for GG-A10-Cosmos rather than a claim that every problem is already solved.

No.	Problem	Most relevant inherited inputs
I. The two flagship tensions		
1	Hubble tension	Recombination history, sound horizon, early ledger, the one-background rate and the ladder's predicted systematic
2	S_8/σ_8 growth tension	Primordial spectrum, dark matter, neutrinos, growth history
II. The dark Universe and its fate		
3	The cosmological-constant problem	GFF = GEP + CEP; boundary GR; residual-ledger interpretation
4	Constant or evolving dark energy?	Residual vacuum export of the completed transition; late-time source evolution
5	Cosmic coincidence / "why now"	Full background closure; the π-locked shares; the epoch calibration
6	The nature of dark matter	Compact-angle ontology; Kaxion window; extra Abelian structure
III. The beginning and the primordial signals		
7	Origin / beginning	Constitutional distinction between model and observation; GG-6 bulk; $E_{\rm GUT}$ boundary scale; OGGEP
8	The cause of inflation	Scale chain; throat geometry; BI-6 normalization; GFF vacuum-selection logic
9	Primordial gravitational waves	Inflationary tensor sector, throat dynamics, CMB B-modes, relic spectrum
10	Baryogenesis	Derived neutrino masses, Majorana channel, leptonic CP block, exact baryonic protection (GG-A9-SM)
11	BBN / lithium-7	Derived baryonic ledger, expansion history, dark-sector allowances, $N_{\rm eff}$ consistency
IV. Structure, cosmic dawn, and the JWST surprises		
12	Early galaxies / JWST	Growth history, dark-matter interpretation, early star-formation environment, primordial spectrum
13	Little red dots / overmassive early black holes	Horizon ontology (OGGEP port); collapse thresholds; PBH windows; early accretion environment
14	Cosmic dawn / reionization / 21 cm	Baryonic ledger; growth history; early-structure timing from the inherited spectrum
15	Black-hole information	OGGEP, GDOS open-system semantics, bulk-to-boundary projection, GFF entropy ledger
V. Cracks in the foundations		
16	Isotropy / dipole / large-angle anomalies	Boundary-projection statistics (OGGEP/GDOS); large-scale sampling of the primordial spectrum
17	Neutrino masses and $N_{\rm eff}$	Derived neutrino block of GG-A9-SM; $N_{\rm eff}$ counting; expansion history

That last point is the one honest rule the paper holds itself to throughout. Inheritance constrains, but it does not by itself prove; so every cosmological entry in what follows carries an explicit provenance tag — *Derived* where the chain closes completely, *Derived-conditional* where it closes

only modulo a named lemma or conjecture, and *Asserted* where the framework points the way but the derivation is not yet in hand. The colors of Figure 14 are exactly these tags, and they are not cosmetic: they are the difference between a result and an intention.

There is a discipline hidden in this inheritance, and it deserves a plain statement. Because every tool used here was built — and paid for — in the particle-physics companion, the cosmology that follows is never free to invent what it needs: whatever explains the sky must already have explained the electron. Philosophically, this begins to dissolve the old division between laws and initial conditions — on a fixed geometry read through an open-system projection, what look like arbitrary starting values are the lawful shadows of structure. If that sounds too strong, the reader holds the antidote: the codebook of Part IX grades every claim this framework will make.

The framework is now in hand. The bulk, its boundary projection, the zero-parameter ledger, and the open-system runtime are everything the cosmology inherits and everything it is allowed to use. The cosmological history this framework supports — the lawful beginning at the GUT edge, the Friedmann kinematics on the boundary, and the radial ascent and cooling of the universe — is set up in the next Part.

16 Part II in plain English: what this paper borrows, and at what strength

Almost nothing in this Part is native to GG-A10-Cosmos. The six-dimensional geometry, the boundary projection, the anomaly-freedom accounting, and the runtime that selects among branches are all established in the upstream papers of the cohort, and they are set out here at exactly the strength their owning papers give them.

That last clause is the one to hold onto. An import re-narrated in a new paper's own words has not been verified; it has been copied, and copying is where the strength of a claim quietly drifts upwards. The rule followed here is that a result is quoted at its parent's grade, and where the parent hedges this paper hedges in the same place and with the same words. Several of the entries below are conditional on named conjectures, and the conditionality travels with them.

What the inheritance buys cosmology is a single thing, and it is worth saying plainly rather than leaving it to be inferred. It buys the absence of dials. If the geometry is fixed and admits no free continuous parameters, then a cosmic number is either forced by it or the framework has nothing to say about that number. There is no third option in which a parameter is adjusted to fit. That is the source of every prediction in this paper and it is also the source of every way the paper can fail: a framework with nothing to tune has nothing to retreat to.

Part III
The cosmological stage

The bulk as the arena of cosmic history — the inherited scale chain, the lawful beginning at the GUT edge, the Friedmann kinematics from the boundary law, and the expansion, cooling, and horizons that set the scene for inflation

17 The thesis: ΛCDM is a theorem of GG-Theory

The companion paper GG-A9-SM (Arisaka, 2026, A9-SM) closed a question that had stood open since the Standard Model was written down: it derived the model's twenty-eight free parameters — the couplings, the masses, the mixing angles — as outputs of one six-dimensional geometry, with no continuous input left to tune. This Part makes the cosmological claim that stands one story above it, and it makes it without hedging:

ΛCDM is a theorem of GG-Theory. The Friedmann–Lemaître–Robertson–Walker equations, and the form and source of every one of the six numbers the concordance model otherwise reads off the sky, are forced by the same constitution, the same bulk, and the same anomaly-locked ledger that fix the particle sector. The universe is not described by a model with six dials; it is the boundary readout of a geometry that has none.

The claim is established by a single descent — the same shape GG-A9-SM ran on the particle stratum, now run on the cosmological boundary stratum (Figure 15):

$$\underbrace{\text{GG-Constitution}}_{\text{GG-A1-Constitution}} \Rightarrow \underbrace{\text{GG-6} + \text{BI-6}}_{\text{GG-A2-GG6, GG-A3-BI6}} \Rightarrow \underbrace{\text{zero free parameters}}_{\text{GG-A4-Euclid}}$$

$$\Rightarrow \underbrace{\text{general relativity}}_{\text{GG-A5-GDOS}} \Rightarrow \underbrace{\text{the FRW equations}}_{\text{GG-A5-GDOS+GG-A6-GRIP} \to \text{here}} \Rightarrow \underbrace{\Lambda\text{CDM}}_{\text{here; Parts IV–X}} .$$

The first four arrows are theorems already proved in the companion papers, named where they are used; the last two are the work of this paper. The sections that follow climb the ladder rung by rung and end by stating the cosmological theorem and auditing it parameter by parameter.

A word, first, on what "theorem" is permitted to mean, because the word does real work here and must not be inflated. It means exactly what it means in GG-A9-SM: the framework *forces the form and the source* of every input, and a published provenance triage records, parameter by parameter, how completely each has been closed — *Derived* (D) where the chain closes end to end, *Derived-conditional* (D*) where it closes modulo a named lemma or conjecture (two for D4 and A15), and *Asserted* (A) where the framework fixes the structural slot but the closed form is not yet in hand. Of the six ΛCDM parameters, three are D* and three are A at this writing (§23). The boldness of the claim is therefore structural — ΛCDM is *forced, not fitted* — and its honesty is

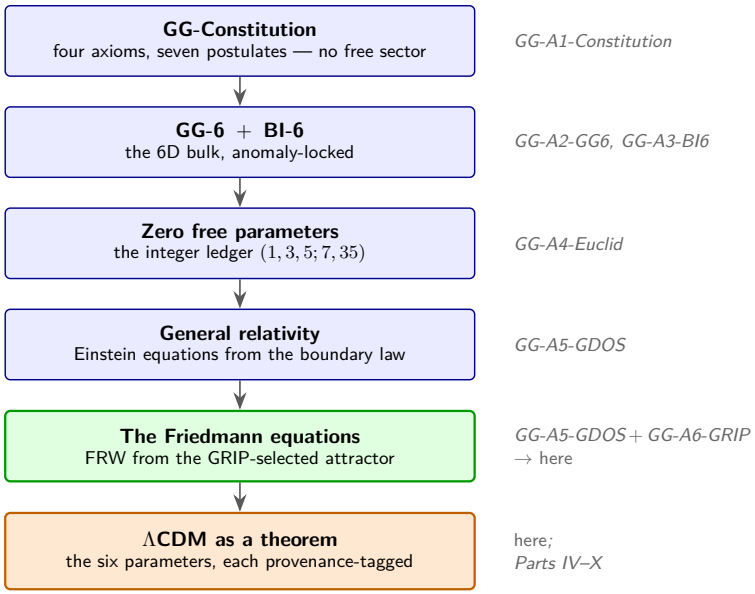

Figure 15: The derivation ladder. The cosmological theorem descends through six rungs, each a result proved in the companion paper named at right. The first four are inherited; the Friedmann line (green) and the ΛCDM capstone (gold) are this paper's own theorems, built on the inherited rungs and the GRIP runtime. This is the cosmological analog of the GRIP spine that GG-A9-SM ran on the particle stratum to derive the Standard Model.

granular: not one parameter is quietly upgraded, and the reader is handed the audit rather than asked to trust the headline. That is the standard GG-A9-SM met for the Standard Model, and the standard this Part holds itself to for the cosmos.

18 Why the universe began at the GUT edge, not the Planck scale

Ask where the universe began and the reflexive answer is the Planck scale, $E_{\rm Pl} \simeq 1.2 \times 10^{19}\,\text{GeV}$, where a hypothetical quantum theory of gravity is imagined to take over from a singular classical past. GG-Theory returns a different answer, and it is worth stating as bluntly as it deserves: *the observable universe begins at the grand-unification edge, $E_{\rm GUT} \simeq 2.3 \times 10^{16}\,\text{GeV}$, and the Planck scale is not a beginning at all.*

The reason is structural, not a matter of taste about quantum gravity. In this framework four-dimensional spacetime is not fundamental: it is the boundary projection of the six-dimensional bulk $\mathcal{M}_4 \times I_r \times S^1_\theta$, and the radial coordinate $r = \ln(E_{\rm GUT}/\mu)$ measures depth into that bulk as renormalization-group scale (Arisaka, 2026, A2-GG6; Arisaka, 2026, A5-GDOS). The ultraviolet edge of the geometry sits at $r = 0$, that is $\mu = E_{\rm GUT}$; the manifold occupies $r \in [0, K]$ and there is no part of it *above* the edge. Since the metric, the causal structure, and the very notions of "space" and "time" are features of the boundary projection that the edge bounds, they are simply *undefined* for $\mu > E_{\rm GUT}$. One cannot ask what the universe was doing at the Planck scale for the same reason one cannot ask what lies north of the North Pole: the coordinate that would label the question has

run out.

The Planck scale, in this reading, is not the floor of spacetime but a *derived* reading off the boundary. The four-dimensional Newton constant emerges as a scale-overlap integral over the throat, $M_{\rm Pl}^2 = \pi M_6^4(1 - e^{-2K})$ in the natural units of the geometry, the factor π coming from the full turn of the compact angle S^1_θ (Arisaka, 2026, A5-GDOS); the geometry even fixes the two scales and their ratio, $E_{\rm GUT} = E_{\rm Pl}\, e^{-2\pi}$, one monodromy of that same angle. But the Planck *energy* $E_{\rm Pl}$ is larger than $E_{\rm GUT}$ and therefore sits *above* the edge, in the region the manifold does not contain; and the Planck *mass* $M_{\rm Pl}$ is a derived coupling strength, not a location in time. Neither is a first moment. The one surface that is a first moment is the edge itself.

This dissolves the Big Bang singularity without denying any equation. Run the classical radiation-dominated solution $a(t) \propto t^{1/2}$ backward and the curvature invariants and the density formally diverge as $t \to 0^+$ (Kolb & Turner, 1990; Mukhanov, 2005); that is a true statement *inside* four-dimensional general relativity. But the four-dimensional Einstein equations are themselves a boundary theorem (§20), derived to hold at and below the ultraviolet surface, not above it; extrapolating them past $E_{\rm GUT}$ to reach $a = 0$ is using a result outside the domain on which it was proved. The framework therefore separates three things the textbook conflates:

$$\underbrace{a \to 0,\ \rho \to \infty}_{\text{formal backward limit}} \quad \neq \quad \underbrace{\mu = E_{\rm GUT},\ r = 0}_{\text{ultraviolet edge of the ledger}} \quad \longrightarrow \quad \underbrace{\Sigma^{\rm Bdy}_{\rm GUT}}_{\text{first admissible initial surface}} .$$

The singular limit is a formal consequence of equations pushed out of bounds; the GUT edge is where the ledger actually starts; and the first physically admissible initial-value surface, $\Sigma^{\rm Bdy}_{\rm GUT}$, carries *finite* inherited data — a finite scale factor, expansion rate, density, and curvature set by the geometry rather than diverging. Figure 16 contrasts the two pictures.

This is a statement about the domain of a description, not a metaphysics of creation. GG-Theory does not assert that nothing existed "before" $E_{\rm GUT}$; it observes that "before," "existed," and "time" are boundary notions the question presupposes and the geometry does not supply above its edge. The honest content is sharp and limited: the first surface on which a cosmology can be written down is $\Sigma^{\rm Bdy}_{\rm GUT}$, and everything this Part derives is anchored there.

Before any single cosmological observable can be derived, the stage on which the whole history plays out has to be built. That stage is the six-dimensional throat and its four-dimensional boundary, read through the inherited constitutional apparatus of Part II. This Part assembles it in four moves and then hands the assembled scene to the inflationary dynamics of Part IV. First we lay down the ruler: the scale chain $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock}$ that the particle papers already fixed, and what each rung of that ladder is responsible for. Second we ask where the story is actually allowed to start, and answer that it starts not at a four-dimensional Planck singularity but at the finite GUT edge of the throat, carrying finite initial data. Third we recover the Friedmann kinematics from the boundary Einstein equation, keeping the textbook equations untouched and constraining only the source ledger that feeds them. Fourth we run that ledger forward through expansion, cooling, and the horizon structure, and read the resulting arrow of time in the open-system language of GDOS/GFF. None of the equations in this Part are new; what is new is the disciplined inheritance chain that decides which terms are allowed to occupy them.

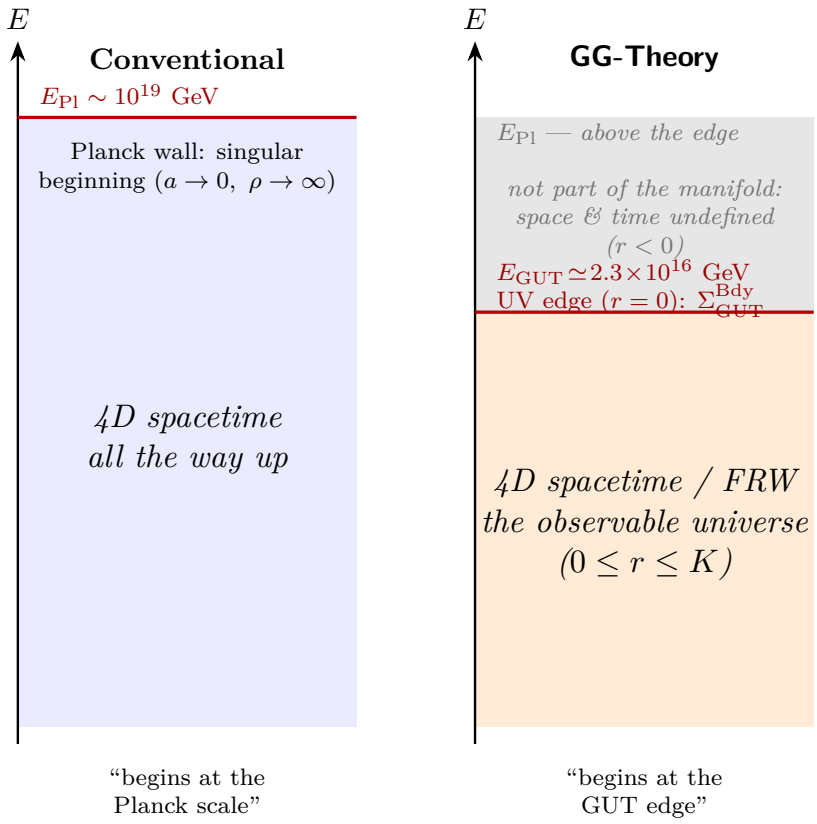

Figure 16: Where the universe begins. *Left:* the conventional picture — spacetime extends up to a Planck wall, with a singular beginning. *Right:* in GG-Theory four-dimensional spacetime is the boundary projection of the bulk and exists only at and below the ultraviolet edge $E_{\rm GUT}$ ($r = 0$); above the edge ($r < 0$, where the Planck *energy* lies) the manifold does not extend and the notions of space and time are undefined. The first admissible initial surface is $\Sigma^{\rm Bdy}_{\rm GUT}$, carrying finite inherited data; the Planck *mass* is a derived coupling, not a first moment.

19 The scale chain: $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock}$

Cosmology and particle physics are the same story told at two ends of one ruler, and in GG-Theory that ruler is literally the radial depth of the throat. The particle paper already engraved its tick marks. Three energy scales, fixed there by a short arithmetic ledger, descend unchanged into the cosmological sequel, and almost everything in this Part is a consequence of how far apart they sit. Before doing any dynamics it is worth setting the ruler down and reading what each mark controls — because a great deal of confusion in cosmology comes from asking a single "fundamental energy" to do several logically different jobs at once. Here the division of labor is clean.

19.1 The inherited scales

The particle paper fixed the three scales that matter most for the cosmological sequel:

$$E_{\rm Pl} = \sqrt{\frac{\hbar c^5}{G_N}}, \qquad E_{\rm GUT} = E_{\rm Pl} e^{-2\pi}, \qquad M_{\rm lock} = \frac{E_{\rm Pl}}{2e^{35}}. \tag{1}$$

Numerically, using the same values already adopted in GG-A9-SM, one has

$$E_{\rm Pl} \simeq 1.22 \times 10^{19}\ {\rm GeV}, \qquad E_{\rm GUT} \simeq 2.28 \times 10^{16}\ {\rm GeV}, \qquad M_{\rm lock} \simeq 3.85 \times 10^{3}\ {\rm GeV}. \tag{2}$$

The associated radial coordinate is

$$r = \ln\!\left(\frac{E_{\rm GUT}}{\mu}\right), \tag{3}$$

so that the full depth from the GUT edge to the lock scale is

$$\Delta r_{\rm GUT\to lock} = \ln\!\left(\frac{E_{\rm GUT}}{M_{\rm lock}}\right) = 35 - 2\pi + \ln 2 \simeq 29.41. \tag{4}$$

Equation (4) is one of the most important numbers in the paper. It says that there is a very large logarithmic hierarchy between the operational ultraviolet edge and the finite infrared lock scale even though both are fixed by the same short arithmetic ledger. The cosmological content of this Part is largely the story of what happens when a boundary universe inherits that large but finite depth.

19.2 What each inherited scale actually controls

Because three very different scales appear already at the beginning of the story, it is useful to state explicitly what each one controls. A great deal of confusion in cosmology comes from trying to make one scale do too many logically different jobs. In the present framework the division of labor is much cleaner. $E_{\rm Pl}$ is the dimensional anchor set by Newton's constant. $E_{\rm GUT}$ is the first ultraviolet scale at which the boundary description is turned on. $M_{\rm lock}$ is the finite infrared seam where the throat exports the low-energy hierarchy and which, crossed on the radiation descent, births the Kaxion and seeds the terminus; the inflationary plateau sits one e-fold below $E_{\rm GUT}$, not here (Part IV). The meV vacuum scale is then not a fourth independent input but a seesaw descendant of $M_{\rm lock}$.

Table 5: The four scales that organize the cosmological history of GG-A10-Cosmos. Only the first three are primitive inherited inputs; the fourth is a derived dark-energy scale candidate. The columns give each scale, its expression, its numerical size, and the role it plays in this Part. The fourth scale is the one to watch: it is assembled from the other three rather than carried in with them, so a reader tracking what the framework assumes should stop at the third row.

Scale	**Expression**	**Numerical size**	**Primary role in this Part**
Planck anchor	$E_{\rm Pl} = \sqrt{\hbar c^5/G_N}$	$1.22 \times 10^{19}\,\text{GeV}$	Converts the gravitational coupling into an energy anchor; not the operative boundary beginning of GG-6 cosmology.
Operational GUT edge	$E_{\rm GUT} = E_{\rm Pl} e^{-2\pi}$	$2.28 \times 10^{16}\,\text{GeV}$	First ultraviolet boundary at which the cosmological description is initialized.
Lock scale	$M_{\rm lock} = E_{\rm Pl}/(2e^{35})$	$3.85 \times 10^{3}\,\text{GeV}$	Hierarchy seam; the prompt second crossing that births the Kaxion and seeds the terminus (not the plateau height, which is fixed one e-fold below $E_{\rm GUT}$, Part IV).
Derived vacuum scale	$m_\Lambda^{\rm geom} = M_{\rm lock}^2/M_{\rm Pl,red}$	$\sim 6\,\text{meV}$	Natural infrared vacuum scale candidate for the late-time dark-energy residue.

The practical lesson of Table 5 is that the cosmology paper should not try to compress all questions into a single "fundamental energy." Beginning, inflation, reheating, and late acceleration are different dynamical questions. They inherit different combinations of the same short ledger.

Two clarifications are essential.

First, $E_{\rm GUT}$ and $M_{\rm lock}$ are not rival claims about the same physical instant. $E_{\rm GUT}$ is the ultraviolet edge at which the boundary description turns on. $M_{\rm lock}$ is the infrared scale at which the throat selects the low-energy seam and where the particle-physics and dark-sector consequences of the locked geometry are exported. The theory therefore contains a GUT-beginning statement and a GUT-adjacent inflation statement, with the lock as the later crossing — not the TeV-inflation statement an earlier version carried.

Second, the radial coordinate r is not identical to cosmological proper time t. The earlier papers interpret r as a geometric renormalization coordinate. Cosmological time is the proper time of the four-dimensional boundary metric. The beginning of the cosmological story is therefore not "a point at $r = 0$" but the first boundary initial-data surface inherited from the ultraviolet end of the throat.

20 The inherited foundation: the first three rungs

We now climb the ladder of Figure 15. The first three rungs are inherited from the companion papers and are recalled here only briefly — they are the arena, not the new physics. The fourth and

fifth rungs — general relativity and the Friedmann equations — are the heart of this paper, and the two sections after this one derive them in full, at the level of a graduate course, with the physical meaning of each step made explicit. The point of the brief recap is to fix in view the parameter-free base on which those derivations stand.

20.1 Rung 1 — the constitution fixes what may count as a cosmology (GG-A1-Constitution)

Before any geometry, a constitution. GG-A1-Constitution (Arisaka, 2026, A1-Constitution) lays down four axioms — local causality, the geometric–gauge equivalence principle, the open-system projection principle OGGEP, and the maximum-symmetry/minimum-freedom selector MSMF — and seven postulates that together define what counts as an admissible physical claim. For cosmology the decisive clause is MSMF: a theory may not carry a free sector it has not earned. This is what forbids the standard repair — bolting on a quintessence field, a tuned inflaton potential, or a dark-sector coupling to fit the data after the fact. The constitution does not yet say what the universe is; it says what kind of answer is allowed to count, and it rules out the parameter freedom ΛCDM otherwise leans on.

20.2 Rung 2 — the bulk and its anomaly lock (GG-A2-GG6, GG-A3-BI6)

The constitution compiles into a unique arena. GG-A2-GG6 (Arisaka, 2026, A2-GG6) derives the six-dimensional bulk with the parameter-free metric $ds_6^2 = e^{-2r}\eta_{\mu\nu}dx^\mu dx^\nu + dr^2 + d\theta^2$ and the unified group $\mathrm{Spin}(1,5) \times SU(5) \times U(1)_B \times U(1)_{Qi}$; GG-A3-BI6 (Arisaka, 2026, A3-BI6) proves that this six-force closure is anomaly-free through the modified Bianchi identity BI-6. The cosmological consequence is that the *stage* on which cosmic history plays out is fixed and rigid: one bulk, one boundary, one scale axis, no freedom to deform any of them. The arena of Part III is not chosen; it is the very object GG-A9-SM used for the particle ledger, read now at the largest scale.

20.3 Rung 3 — zero free parameters (GG-A4-Euclid)

The anomaly lock has an arithmetic shadow. GG-A4-Euclid (Arisaka, 2026, A4-Euclid) proves that BI-6 forces the integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, closed by the dual lock $K = LJ = M^2 + N^2 + L^2 = 35$; together with π, e, and the conversion factors $\{c, G_N, \hbar; k_B, \mathsf{Q_B}\}$ these whole numbers are the entire stock from which every dimensionless cosmological quantity is built. This is the rung that makes the eventual claim sharp: the source terms that feed the Friedmann equations — the density fractions, the baryon-to-photon ratio, the recombination redshift — are not numbers to be measured and inserted but closed-form expressions over a fixed ledger. A cosmology built on this rung cannot be retuned; a single wrong number falsifies it.

20.4 Rung 3, restated for cosmology

That third rung is what makes the eventual claim sharp: the source terms that feed the Friedmann equations — the density fractions, the baryon-to-photon ratio, the recombination redshift — are not numbers to be measured and inserted but closed-form expressions over a fixed ledger. With the

arena (Rungs 1–2) and the ledger (Rung 3) in hand, we turn to the two derivations that are the substance of this paper: gravity, and then the cosmology it carries.

21 General relativity as a theorem of GG-6/GDOS

In every other framework general relativity is a *starting point*: one writes the Einstein–Hilbert action by hand and proceeds (Misner *et al.*, 1973; Wald, 1984; Weinberg, 1972). In GG-Theory it is an *output*. The four-dimensional Einstein equations, the Newton constant that sets the strength of gravity, and the scale of the cosmological constant all descend from the six-dimensional bulk by a single controlled reduction (Arisaka, 2026, A5-GDOS). This is the load-bearing rung of the paper, so we carry the derivation out in full, at the level of a graduate course, recalling each standard object as we use it.

21.1 The starting point: curvature in six dimensions

The dynamics of any metric is governed by the Einstein–Hilbert action — the integral of the Ricci scalar R, the single number that measures how curved spacetime is, weighted by the invariant volume. In the bulk this reads

$$S_{\mathrm{EH}}^{(6)} = \frac{M_6^4}{2} \int d^6x \, \sqrt{-g_6}\, R_6,$$

with M_6 the (derived) six-dimensional gravitational scale, g_6 the bulk metric determinant, and R_6 its Ricci scalar. No other gravitational term may be added: the constitution admits only the anomaly-free closure of GG-A3-BI6, so the bulk is rigid. *Physical picture:* curvature costs action, and gravity is whatever dynamics minimizes that cost; all of four-dimensional gravity is folded into this one expression, and the task is to read it out on the boundary by integrating away the two extra directions.

21.2 Integrating out the throat: the scale-profile reduction

The bulk metric of GG-A2-GG6 is the warped product

$$ds_6^2 = e^{-2r}\, g_{\mu\nu}(x)\, dx^\mu dx^\nu + dr^2 + d\theta^2, \qquad r \in [0, K],\ \theta \in [0, 2\pi),$$

in which $g_{\mu\nu}(x)$ is the metric a boundary observer measures and the warp e^{-2r} encodes renormalization-group depth. Inserting this ansatz into $S_{\mathrm{EH}}^{(6)}$ and integrating over r and θ needs two standard facts about warped products. First, the volume element factorizes,

$$\sqrt{-g_6} = e^{-4r} \sqrt{-g_4},$$

since the warp multiplies each of the four boundary directions by e^{-2r}. Second, the curvature decomposes as

$$R_6 = e^{2r} R_4 - 20,$$

where R_4 is the curvature of $g_{\mu\nu}$, the prefactor e^{2r} is the rescaling of an inverse-length-squared quantity by the warp, and the constant -20 is the (negative) curvature the throat carries on its own. Assembling the integrand,

$$\sqrt{-g_6}\,R_6 = \sqrt{-g_4}\,(e^{-2r}R_4 - 20\,e^{-4r}),$$

the two extra integrals are elementary,

$$\int_0^K e^{-2r}\,dr = \tfrac{1}{2}(1-e^{-2K}), \qquad \int_0^K e^{-4r}\,dr = \tfrac{1}{4}(1-e^{-4K}), \qquad \int_0^{2\pi} d\theta = 2\pi,$$

and carrying them out collapses the six-dimensional action to a four-dimensional one,

$$S_{\rm EH}^{(6)} \;\longrightarrow\; \int d^4x\,\sqrt{-g_4}\,\Big[\tfrac{\pi M_6^4}{2}(1-e^{-2K})R_4 \;-\; 5\pi M_6^4(1-e^{-4K})\Big].$$

21.3 Reading off the four-dimensional theory

This is exactly the four-dimensional Einstein–Hilbert action with a cosmological constant,

$$S_{\rm EH}^{(4)} = \frac{M_{\rm Pl}^2}{2}\int d^4x\,\sqrt{-g_4}\,(R_4 - 2\Lambda_4),$$

once the two coefficients are read off:

$$\boxed{M_{\rm Pl}^2 = \pi M_6^4(1-e^{-2K})}, \qquad \Lambda_4^{\rm naive} = \frac{5(1-e^{-4K})}{1-e^{-2K}} \xrightarrow{K=35} 5 \;\;(\text{in } M_{\rm Pl}\text{ units}).$$

The Planck mass is not a postulate but the *overlap integral of the throat*, the factor π being the full turn of the compact angle S^1_θ. *Physical picture:* gravity is weak — $M_{\rm Pl}$ is enormous — precisely because curvature is diluted over the extra directions; the hierarchy between gravity and the other forces is geometric, not tuned. Varying $S_{\rm EH}^{(4)}$ together with the boundary matter action gives the four-dimensional Einstein equation; in its full inherited form it carries two small corrections from the heavy modes that were integrated out (Arisaka, 2026, A5-GDOS),

$$M_{\rm Pl}^2\,G_{\mu\nu} + \Lambda_4\,g_{\mu\nu} = \langle T_{\mu\nu}\rangle_{\rm Bdy} + \Delta T_{\mu\nu}^{\rm GS} + \Delta T_{\mu\nu}^{\rm KK},$$

with $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R_4\,g_{\mu\nu}$ the Einstein tensor, $\langle T_{\mu\nu}\rangle_{\rm Bdy}$ the projected boundary stress-energy, and the last two terms the Green–Schwarz and Kaluza–Klein inflow. At cosmological energies these are suppressed by $(E/M_{\rm lock})^2 \sim 10^{-15}$ and drop, leaving the textbook Einstein equation exactly,

$$\boxed{G_{\mu\nu} + \Lambda_4\,g_{\mu\nu} = 8\pi G_N\,T_{\mu\nu}}, \qquad G_N = \frac{1}{8\pi M_{\rm Pl}^2}.$$

21.4 Conservation, the Newton constant, and the vacuum scale

Three features of this result carry directly into cosmology. *(i) Conservation is a theorem.* The local law $\nabla_\mu T^{\mu\nu} = 0$ is not imposed by hand; it is the Ward identity of bulk diffeomorphism

invariance projected to the boundary, so the conservation of energy and momentum that the Friedmann continuity equation expresses is itself derived. *(ii) The Newton constant is fixed.* $G_N = 1/(8\pi M_{\rm Pl}^2)$ is set by the throat geometry, not chosen. *(iii) The vacuum scale arrives with a number.* The naive value $\Lambda_4^{\rm naive} \approx 5$ in Planck units is reduced by a boundary-localized seesaw to

$$m_\Lambda^{\rm geom} = \frac{M_{\rm lock}^2}{M_{\rm Pl,red}} \approx 6 \text{ meV},$$

which by itself collapses the notorious 10^{120} vacuum-energy discrepancy (Carroll, 2001; Weinberg, 1989) to a residual factor of a few in scale — and, once the ledger coefficient $\zeta_\Lambda = \pi/LK$ is selected, to the percent level (Eq. (112)). *Physical picture:* the cosmological constant is small not by cancellation to sixty decimal places but because a large scale ($M_{\rm Pl}$) divides a modest one ($M_{\rm lock}^2$), a seesaw of exactly the kind that makes neutrinos light. The order-unity coefficient and the sign of Λ are *not* yet closed in GG-A5-GDOS; pinning them is the work of Part IV, and we flag that debt rather than hide it. With that one acknowledged gap, general relativity — the Einstein equation, G_N, and the scale of Λ — is a theorem of GG-6/GDOS, derived rather than assumed.

22 The Friedmann equations as a theorem

The Einstein equation just derived holds for any spacetime. A cosmology is the special solution in which the universe looks the same at every place and in every direction — the cosmological principle (Peebles, 1993; Weinberg, 2008). In the textbook this principle is an *assumption*, justified after the fact by the smoothness of the microwave sky. In GG-Theory it is the geometric *attractor* of the runtime: the GRIP cascade DGI → GSSB → GA of GG-A6-GRIP (Arisaka, 2026, A6-GRIP), run on the cosmological boundary stratum, flows to the maximally symmetric configuration, so homogeneity and isotropy are *selected*, not posited. (That this attractor is *unique* is a D* result, resting on the attractor uniqueness GG-A6-GRIP states and the inflationary smoothing of Part IV; we state it, not bury it.) Granting the attractor, the derivation of the Friedmann equations is the standard graduate computation, which we now carry out because it is the second load-bearing rung.

22.1 The maximally symmetric metric

Homogeneity and isotropy fix the metric up to one function of time and a discrete curvature sign (Friedmann, 1922; Lemaître, 1927); the standard graduate derivations are collected in the textbooks of Dodelson (2003); Kolb & Turner (1990); Mukhanov (2005); Weinberg (2008). In comoving coordinates,

$$ds^2 = -dt^2 + a^2(t)\Big[\frac{dr^2}{1-kr^2} + r^2\, d\Omega_2^2\Big], \tag{5}$$

with $a(t)$ the scale factor — the single dynamical degree of freedom, the "size" of space — and $k \in \{-1, 0, +1\}$ the sign of the spatial curvature. *Physical picture:* $a(t)$ stretches every comoving separation in lockstep; nothing in the metric singles out a place or a direction, which is the cosmological principle made geometric.

22.2 The curvature of an expanding universe

The Einstein tensor of this metric is a textbook computation (Misner *et al.*, 1973; Wald, 1984); we record the load-bearing steps. With an overdot for d/dt, the non-vanishing Christoffel symbols are

$$\Gamma^0_{ij} = a\dot{a}\,\gamma_{ij}, \qquad \Gamma^i_{0j} = \frac{\dot{a}}{a}\,\delta^i_j, \qquad \Gamma^i_{jk} = \widehat{\Gamma}^i_{jk}(\gamma),$$

where γ_{ij} is the comoving spatial metric and $\widehat{\Gamma}$ its own connection. From these the Ricci tensor is

$$R_{00} = -3\frac{\ddot{a}}{a}, \qquad R_{ij} = \Big(\frac{\ddot{a}}{a} + 2\frac{\dot{a}^2}{a^2} + 2\frac{k}{a^2}\Big)g_{ij},$$

and contracting gives the Ricci scalar

$$R_4 = 6\Big(\frac{\ddot{a}}{a} + \frac{\dot{a}^2}{a^2} + \frac{k}{a^2}\Big).$$

Assembling $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R_4\,g_{\mu\nu}$, the independent components are

$$G_{00} = 3\Big(\frac{\dot{a}^2}{a^2} + \frac{k}{a^2}\Big), \qquad G_{ij} = -\Big(2\frac{\ddot{a}}{a} + \frac{\dot{a}^2}{a^2} + \frac{k}{a^2}\Big)g_{ij}.$$

Physical picture: G_{00} collects the expansion rate and the spatial curvature — it is the geometry's measure of how much energy is present; G_{ij} adds the acceleration $\ddot{a}$, the geometry's response to pressure.

22.3 The source: a perfect fluid from the inherited ledger

The same symmetry forces the boundary stress-energy to be that of a perfect fluid,

$$T^{\mu}{}_{\nu} = \mathrm{diag}(-\rho,\, p,\, p,\, p),$$

with ρ the energy density and p the pressure: isotropy forbids any momentum flux or shear. Here is the single place the framework departs from the textbook — not in the equation but in what may fill it. ρ and p are not free fluids to be fitted but the coarse-grained boundary projection of the bulk source ledger,

$$\rho = \rho_r + \rho_b + \rho_{\mathrm{DM}} + \rho_\Lambda + \cdots, \tag{6}$$

each sector's identity, equation of state $w_i = p_i/\rho_i$, and normalization fixed by the integer ledger of the inherited foundation rather than chosen.

22.4 The Friedmann equations

Inserting the Einstein tensor and the perfect-fluid source into $G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G_N T_{\mu\nu}$, the time–time (00) component yields the first Friedmann equation,

$$\boxed{H^2 \equiv \Big(\frac{\dot{a}}{a}\Big)^2 = \frac{8\pi G_N}{3}\,\rho - \frac{k}{a^2} + \frac{\Lambda}{3}} \tag{7}$$

and combining it with the trace of the space–space (ij) component gives the second, the acceleration equation,

$$\boxed{\frac{\ddot{a}}{a} = -\frac{4\pi G_N}{3}(\rho + 3p) + \frac{\Lambda}{3}} \tag{8}$$

The conservation law $\nabla_\mu T^{\mu 0} = 0$ — already a theorem from the previous section — gives the continuity equation,

$$\dot{\rho} + 3H(\rho + p) = 0, \tag{9}$$

which is not independent of the other two but states how each sector dilutes. *Physical picture:* the first Friedmann equation is a first integral of the motion, "energy density sets the expansion rate"; the second says pressure gravitates, so a sector with $\rho + 3p > 0$ decelerates the universe and a vacuum with $p = -\rho$ accelerates it; the continuity equation is simply dilution, $\rho_i \propto a^{-3(1+w_i)}$ (Dodelson, 2003; Kolb & Turner, 1990).

22.5 What is a theorem, and what is inherited

The three equations are identical in form to the textbook Friedmann system, and that identity is the point, not a limitation: GG-Theory modifies neither gravity nor kinematics. What it fixes is everything on the right-hand side. G_N is the throat overlap integral of the previous section; Λ carries the seesaw scale; and the contents of ρ — which sectors exist, with which equations of state, in which proportions — are read from the fixed ledger, not measured and inserted. The Friedmann equations are therefore a theorem of GG-Theory: the boundary Einstein equation, restricted to the runtime-selected symmetric attractor, sourced by a ledger with no free dials. The full line-by-line specialization — every Christoffel symbol and curvature component — is collected in Appendix E; what matters here is the chain, bulk action → Einstein equation → Friedmann system, each link forced. The concordance model is what this system becomes when the ledger is read out, which is the subject of the next section.

23 ΛCDM as a theorem: the six parameters

The two derivations just completed — general relativity, then the Friedmann equations — deliver the cosmological stage with a fixed source ledger; the concordance model is what that stage becomes when the ledger is read out. ΛCDM is specified by six numbers — the physical baryon and cold-dark-matter densities ω_b and ω_c, the acoustic angular scale θ_*, the reionization optical depth τ, and the amplitude A_s and tilt n_s of the primordial spectrum — from which the entire expansion history and perturbation pattern follow (Dodelson, 2003; Planck Collaboration, 2018). The claim of this Part is that all six are fixed *in form* by the geometry, so that the concordance model is a theorem of GG-Theory rather than a six-parameter fit. Table 6 states the claim parameter by parameter, with the provenance tag that keeps it honest.

Table 6: ΛCDM as a theorem: the six concordance parameters and the principal derived outputs, each with its GG-Theory closed form (or the structural slot it occupies), its provenance tag (D Derived, D* Derived-conditional, A Asserted), and where it is derived. The full twenty-one-entry audit is the codebook of Part IX.

ΛCDM parameter	GG-Theory form / origin	Prov.	Derived in
The six concordance parameters			
$\omega_b = \Omega_b h^2$	$\Omega_b = 1/2\pi^2$ (closed-packet baryon share)	D*	Part X
$\omega_c = \Omega_c h^2$	$\Omega_{\rm dm} = 1/\pi - 1/2\pi^2$	D*	Part X
θ_* (acoustic scale)	r_s/D_M on the inventory at the carried rate (the switch $z_* = e^7$ annotated, not scored)	D*	Part VI
τ (reionization depth)	reionization-seam slot fixed; closed form owed	A	Part VI
A_s (scalar amplitude)	inflationary two-point slot fixed; closed form owed	A	Part IV
n_s (scalar tilt)	$n_s = 0.9634$ exact on the plateau ($1 - 2/N^* = 0.9620$ at leading order; Lemma IV, revised)	D*	Part IV
Principal derived outputs			
Ω_m	$1/\pi \approx 0.318$	D*	Part X
Ω_Λ	$1 - 1/\pi \approx 0.682$	D*	Part X
H_0	one background 66.8 on every ladder-free probe; the ladder's 73.5 the predicted outlier	D*	Part VI
age t_0	FRW integral over the π-locked inventory	D*	Part VIII

The audit is deliberately unflattering where it must be. Four of the six — the two matter densities, the acoustic scale, and the scalar tilt — are *Derived-conditional*: their closed forms ($\Omega_b = 1/2\pi^2$, $\Omega_{\rm dm} = 1/\pi - 1/2\pi^2$, the acoustic scale's r_s/D_M on that inventory, and $n_s = 0.9634$ exact, $1 - 2/N^* = 0.9620$ at leading order) follow from the integer ledger, the π-locked inventory of Part X, and the derived plateau of Part IV, each modulo the named conjecture that Part states explicitly (two for the dark-matter density). One — the primordial amplitude A_s — is the launch sector's one *Asserted* declared input, its warp candidate (75%) flagged and its closing computation (Lemma V′) named. The sixth, the reionization optical depth τ, is astrophysics — a late-time reionization quantity outside the fundamental claim, carried at its measured value. This is the precise sense in which ΛCDM is a theorem: its *form* is forced — six closed-form slots over one fixed ledger, none free to retune — while the completeness of each form is audited openly, four derived-conditional, one declared input, one astrophysical. The bold claim and the honest ledger are the same statement read at two resolutions; the chapters that follow set about closing the gap.

24 Expansion, cooling, and horizons

Given the equations and the ledger, the rest is kinematics — but kinematics that does an enormous amount of work. Run a homogeneous universe with positive energy density forward and three things happen automatically: it expands, it cools, and it continually reshapes which regions can talk to which. This section is the thermal-history backbone of the paper, the standard one (Dodelson, 2003; Kolb & Turner, 1990; Mukhanov, 2005), read through the GG-6/GDOS source ledger. We derive the scaling laws for each component, the cooling law from comoving entropy, the temperature–time dictionary that every later milestone rests on, the horizon structure that motivates inflation and dark energy, and the universal Γ/H competition that decides every freeze-out. The equations are textbook — Friedmann's dynamical solutions (Friedmann, 1922) and the Hubble law (Hubble, 1929) read against standard rulers such as the baryon acoustic oscillation scale (Cole *et al.*, 2005; Eisenstein *et al.*, 2005; Peebles & Yu, 1970) — and the discipline is in the ledger that fills them. The full from-scratch derivation — comoving-entropy conservation, the relativistic degree-of-freedom count $g_*(T)$, and the freeze-out condition $\Gamma \sim H$ — is given in Appendix F; here we collect the results on which the later milestones rest.

24.1 Exact scaling law for a constant equation-of-state component

Equation (9) is enough to derive the standard scaling law. For a component with $p = w\rho$,

$$\dot{\rho} + 3H\rho(1+w) = 0. \tag{10}$$

Using $H = \dot{a}/a$, this becomes

$$\frac{\dot{\rho}}{\rho} = -3(1+w)\frac{\dot{a}}{a}. \tag{11}$$

Integrating both sides yields

$$\rho(a) = \rho_0\, a^{-3(1+w)}. \tag{12}$$

This compact formula already encodes the entire coarse-grained thermal history:

$$\text{radiation } (w = 1/3): \qquad \rho_r \propto a^{-4}, \tag{13}$$

$$\text{pressureless matter } (w = 0): \qquad \rho_m \propto a^{-3}, \tag{14}$$

$$\text{curvature:} \qquad \rho_k \propto a^{-2}, \tag{15}$$

$$\text{vacuum } (w = -1): \qquad \rho_\Lambda = \text{const.} \tag{16}$$

The extra factor of a^{-1} for radiation compared with matter is the cooling term. Matter number density dilutes like volume, a^{-3}. Radiation number density also dilutes like a^{-3}, but each photon's energy redshifts by another factor of a^{-1}. That is why the universe cools.

If spatial curvature is negligible, Eq. (7) together with Eq. (12) can be integrated exactly for any constant $w \neq -1$:

$$a(t) \propto t^{\frac{2}{3(1+w)}}. \tag{17}$$

Thus

$$\text{radiation era:} \qquad a(t) \propto t^{1/2}, \qquad H = \frac{1}{2t}, \tag{18}$$

$$\text{matter era:} \qquad a(t) \propto t^{2/3}, \qquad H = \frac{2}{3t}. \tag{19}$$

For a true vacuum term with $w = -1$, Eq. (7) gives $H = \text{const}$, so

$$a(t) \propto e^{H_\Lambda t}. \tag{20}$$

In one line: radiation gives rapid decelerating expansion, matter gives slower decelerating expansion, and vacuum energy gives accelerating exponential expansion.

24.2 Cooling law from comoving entropy conservation

The statement that the universe cools is usually abbreviated to $T \propto a^{-1}$. It is worth deriving that relation explicitly because later sections on freeze-out, reheating, BBN, and recombination all depend on it.

For a relativistic plasma in approximate equilibrium,

$$\rho_r(T) = \frac{\pi^2}{30} g_*(T) T^4, \qquad p_r(T) = \frac{1}{3}\rho_r(T), \tag{21}$$

and the entropy density is

$$s(T) = \frac{\rho_r + p_r}{T} = \frac{2\pi^2}{45} g_{*s}(T) T^3. \tag{22}$$

If entropy is conserved in a comoving volume, then

$$\frac{\mathrm{d}}{\mathrm{d}t}\left[s(T) a^3\right] = 0. \tag{23}$$

When g_{*s} is approximately constant, Eq. (23) reduces to

$$T \propto a^{-1}. \tag{24}$$

More generally,

$$T(a) \propto g_{*s}^{-1/3}(T)\, a^{-1}. \tag{25}$$

This small correction matters whenever species become nonrelativistic and dump entropy into the still-coupled plasma. Conceptually, however, the law remains simple: absent significant entropy production, expansion means cooling.

Combining Eqs. (7) and (21) in the radiation era gives the standard temperature–Hubble relation,

$$H(T) = \sqrt{\frac{8\pi G_N}{3}\frac{\pi^2}{30} g_* T^4} \simeq 1.66\, g_*^{1/2} \frac{T^2}{M_{\mathrm{Pl}}}. \tag{26}$$

Using $H \simeq 1/(2t)$ during radiation domination then gives

$$t \simeq 0.301\, g_*^{-1/2} \frac{M_{\rm Pl}}{T^2}. \tag{27}$$

Equation (27) is one of the most useful equations in practical cosmology. It lets us translate temperatures into approximate times throughout the early universe. BBN, recombination, neutrino decoupling, electroweak freeze-out, and many later audit steps all rest on this relation.

24.3 Why the expansion cools every sector differently

Although $T \propto a^{-1}$ is the right law for a coupled relativistic plasma, not every component cools in the same way. It is important not to conflate the photon temperature, the matter kinetic temperature, and the dilution of vacuum energy.

Radiation. For photons and any tightly coupled relativistic bath, the typical particle energy redshifts as $E \propto a^{-1}$, and the temperature follows Eq. (24).

Nonrelativistic matter. For a decoupled nonrelativistic gas, the number density still scales as a^{-3}, but the kinetic temperature scales approximately as

$$T_{\rm nr} \propto a^{-2} \tag{28}$$

after decoupling because momenta redshift as $p \propto a^{-1}$ while kinetic energy is $p^2/2m$.

Vacuum energy. A true cosmological constant does not cool at all:

$$\rho_\Lambda = \text{const}, \qquad p_\Lambda = -\rho_\Lambda. \tag{29}$$

This is why vacuum energy eventually wins against matter and radiation. It does not dilute while they do.

Open-system caveat. Within GDOS one should remember that "temperature" is sometimes an effective boundary variable rather than a complete microscopic description. Nevertheless, for the eras relevant to ordinary radiation, recombination, and large-scale expansion, the textbook thermodynamic language remains the correct leading coarse-grained one.

24.4 Why the universe expands at all

One sometimes speaks as if "expansion" itself required a special new force. Equations (7) and (8) show that this is not the right way to think. Once a homogeneous metric and nonzero background energy density are present, a static universe is highly nongeneric. In fact, for $k = 0$ and $\rho > 0$, the first Friedmann equation implies $H^2 > 0$, so the natural solutions are expanding or contracting. The historical universe is observed to be on the expanding branch.

The deeper question is therefore not "why does a homogeneous universe expand at all?" but rather:

(a) why did the boundary universe start on the expanding rather than the contracting branch;

(b) why did it later experience accelerated phases;

(c) and how are the source terms behind those phases constrained by the inherited ledger?

GG-A10-Cosmos addresses those sharper questions rather than treating expansion itself as a mystery separate from general relativity.

24.5 The comoving Hubble radius and the meaning of horizon growth

A great deal of cosmological intuition is encoded in the comoving Hubble radius

$$(aH)^{-1}. \tag{30}$$

During ordinary radiation and matter domination,

$$\text{radiation:} \qquad (aH)^{-1} \propto a, \tag{31}$$

$$\text{matter:} \qquad (aH)^{-1} \propto a^{1/2}. \tag{32}$$

Thus the comoving Hubble radius grows with time in decelerating eras. By contrast, during quasi-de Sitter inflation where H is nearly constant and $a \propto e^{Ht}$,

$$(aH)^{-1} \propto e^{-Ht}, \tag{33}$$

so the comoving Hubble radius shrinks.

This distinction is the kinematic core of the horizon problem and its solution. In a purely decelerating universe, regions now far apart were never in mutual causal contact. An accelerated phase can reverse that logic by pushing previously microscopic, causally connected regions to cosmological sizes before standard radiation and matter evolution resume.

In the language of this Part, this kinematic fact will later become the main observational reason for wanting an inflationary stage at all. Whether that stage is driven by a free inflaton potential or by a throat modulus is secondary to the kinematic need. The need is fixed by Eq. (33).

24.6 Particle horizon, event horizon, and causal size

The word "horizon" is often used loosely in cosmology, but several distinct integrals matter. The first is the particle horizon,

$$d_p(t) = a(t) \int_{t_{\rm ini}}^{t} \frac{\mathrm{d}t'}{a(t')}, \tag{34}$$

which measures the maximum physical distance from which a light signal could have reached the observer by time t. The second is the event horizon,

$$d_e(t) = a(t) \int_t^\infty \frac{\mathrm{d}t'}{a(t')}, \tag{35}$$

which measures the maximum distance from which future signals can ever reach the observer. The two are not the same, and their distinction is crucial when one talks about inflation, dark energy, and ultimate fate.

For a power-law background $a \propto t^n$ with $0 < n < 1$, Eq. (34) gives

$$d_p(t) = \frac{t}{1-n} \times c, \tag{36}$$

while the event horizon diverges, so no future event horizon exists. In particular,

$$\text{radiation } (n = 1/2): \qquad d_p = 2ct, \tag{37}$$

$$\text{matter } (n = 2/3): \qquad d_p = 3ct. \tag{38}$$

By contrast, in de Sitter space with $a \propto e^{Ht}$, Eq. (35) converges and gives

$$d_e = \frac{c}{H}. \tag{39}$$

That single formula foreshadows the two opposite causal properties of the universe emphasized elsewhere in the paper: early inflation shrinks the comoving Hubble radius and solves the horizon problem, while late dark-energy domination creates a genuine future event horizon.

The reason to include these derivations here is not pedagogical excess. They keep the causal story of the paper precise. "Expansion and cooling" is not merely the dilution of a thermal bath. It is also the continual reshaping of which regions can causally communicate with which other regions, and the reshaping of that communication pattern is exactly what makes both inflation and late-time acceleration cosmologically consequential.

24.7 Freeze-out, decoupling, and the universal competition Γ/H

Once Eqs. (26) and (27) are in hand, nearly every event in the thermal history reduces to one basic competition:

$$\frac{\Gamma(T)}{H(T)}. \tag{40}$$

If $\Gamma/H \gg 1$, microscopic interactions are fast compared with the expansion and the process tracks equilibrium efficiently. If $\Gamma/H \lesssim 1$, the process freezes out, decouples, or falls out of equilibrium. This universal criterion will reappear throughout the remainder of the paper:

- it controls when reheating completes after inflation;
- it controls whether heavy neutrino decays can generate baryon asymmetry;

- it controls neutrino decoupling and any extra-radiation effects;
- it controls the nuclear reaction network of BBN; and
- it controls recombination and photon last scattering.

At a conceptual level, the entire cosmic thermal history is just the competition between internal relaxation rates and the external dilution rate set by the expansion.

24.8 A compact table of expansion and cooling laws

Table 7: Useful background scaling laws for Part II. The equations are standard; what changes later is the interpretation of the source sectors. The columns give the component, its equation-of-state parameter, the density scaling, the expansion law, and a cooling note. Nothing in the table is new; it is set down so that the later reinterpretation can be checked against the standard result line by line rather than asserted against it.

Component	w	$\rho(a)$	$a(t)$	**Cooling / comment**
Radiation	1/3	a^{-4}	$t^{1/2}$	$T \propto a^{-1}$, comoving entropy approximately conserved.
Pressureless matter	0	a^{-3}	$t^{2/3}$	Number density dilutes as volume; decoupled kinetic temperature scales as a^{-2}.
Curvature	—	a^{-2}	—	Effective bookkeeping term, not a fluid.
Vacuum / cosmological constant	-1	constant	e^{Ht}	No cooling; late-time accelerator.

25 The radial chronology and the arrow of time

The kinematics tell us how the stage moves; they do not yet tell us why it moves in one direction only. The spine that orients the whole history is the radial coordinate $r = \ln(E_{\mathrm{GUT}}/\mu)$ of Eq. (3): cosmic history is an ascent out of the ultraviolet edge of the throat, and the open-system descent of GDOS/GFF gives that ascent a preferred direction. Read this way, the source sectors are not an arbitrary menu but inherited terms, the thermodynamic arrow is the macroscopic face of one monotone ledger, and the successive eras — radiation, matter, vacuum — are attractor phases of a single descent. With that orientation fixed, the stage is fully built, and we can hand it to Part IV, where the inflationary modulus and the late F_{GFF} vacuum inventory take over the action.

25.1 Reinterpreting the source sectors in GG-6/GDOS

The standard equations explain expansion and cooling once the components are given. GG-6/GDOS enters when we ask what those components are allowed to be. At the level of this Part the

leading reinterpretation is:

$$\rho_r \leftrightarrow \text{relativistic boundary degrees of freedom descending from the zero-mode and light descendant spectrum,} \tag{41}$$

$$\rho_m \leftrightarrow \text{baryons plus the compact-angle dark sector,} \tag{42}$$

$$\rho_\Lambda \leftrightarrow \text{residual GFF/CEP vacuum functional after GSSB and lock selection.} \tag{43}$$

This still looks schematic, but it is already a strong restriction. It means that the cosmological source ledger is not an arbitrary menu:

(i) matter is not free to be any imaginable hidden field if the compact-angle ontology already supplies a constrained candidate;

(ii) baryogenesis is not free to live in any arbitrary CP-violating sector if the neutrino block is already fixed;

(iii) and the accelerating term is not free to be a wholly inserted scalar potential if the GFF ledger already provides a natural residual vacuum interpretation.

Once again, the equations of expansion and cooling stay standard. What narrows is the meaning of ρ_i.

25.2 Expansion, cooling, and the arrow of time in the GDOS language

The standard FRW equations explain the kinematics of expansion, but GG-A10-Cosmos also needs to explain why the expanding thermal history has a preferred direction. This is where the open-system language of GDOS and the GFF ledger become more than philosophical packaging.

For a homogeneous comoving volume $V_c = a^3$, the first law of thermodynamics may be written as

$$\mathrm{d}(\rho a^3) + p\,\mathrm{d}(a^3) = T\,\mathrm{d}(s a^3), \tag{44}$$

where s is the entropy density of the coarse-grained plasma. If the expansion is adiabatic at the level of the tracked component, the right-hand side vanishes and one recovers the continuity equation. If there is entropy production through irreversible processes, decays, or reheating, then $\mathrm{d}(sa^3) > 0$.

Equation (44) is the thermodynamic translation of the GDOS statement that the boundary universe is an open system descending on an GFF ledger. In that language,

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathsf{GFF} \leq 0 \qquad \text{during a dissipative relaxation stage,} \tag{45}$$

while the entropic-curvature contribution drives the coarse-grained arrow of time. Inflation, reheating, matter domination, recombination, and late vacuum domination are therefore not only different source-dominance regimes. They are also different attractor phases of one open-system descent.

This observation helps organize the whole cosmological story. Radiation domination is the phase in which entropy density is high and the bath is tightly coupled. Freeze-out events are the moments when one subset of degrees of freedom decouples from the common GFF descent because

its microscopic rate falls below the global dilution rate. Reheating is the moment when a formerly vacuum-like contribution is converted into a thermalized radiation bath, thereby increasing sa^3. Late dark-energy domination is the approach to a different kind of attractor, one in which ordinary matter becomes negligible and the residual vacuum functional dominates the boundary ledger.

The underlying message is simple but important: the arrow of cosmological time in GG-A10-Cosmos is not added by hand on top of the Friedmann equations. It is the macroscopic face of the same open-system descent that the earlier GDOS paper already elevated to theorem level for general subsystems. That arrow is consumed here, not created: per the constitutional custody chain, AX-3 owns the arrow constitutionally — the one-way reading of the dictionary is a clause of the constitution (the AX-3 synthesis of Arisaka, 2026, A1-Constitution) — with the canonical law-level derivation in Part IX of the CPT member paper (Arisaka, 2026, P5-CPT), carried out over a bulk that is itself exactly T-symmetric, and the dynamical implementation in the GDOS companion (Arisaka, 2026, A5-GDOS); the arrow of cosmic time read off the radial descent is the cosmological face of that one constitutional arrow, a consumer of AX-3's non-invertibility, not a private one.

25.3 The fiber without holonomy: why the chronology cannot run backward

The preferred direction that this section has attributed to the GDOS/GFF descent has a root one level deeper, and it is worth one subsection because it is an exact topological statement rather than a thermodynamic plausibility: *holonomy requires closed loops, and an interval has none.* The bulk's two non-spatial directions are opposite objects in precisely this respect. The compact circle S^1_θ is nothing but closed loops: everything it carries is holonomy — Wilson phases, integer windings, the quantized charges of the gauge sector — gauge-invariant and topologically protected. The radial interval, $r = \ln(E_{\rm GUT}/\mu)$, can hold no holonomy at all: no loop on an interval ever closes, so the fiber has no gauge-invariant memory, no quantized charge, no protected content. Transport along it cannot be a group action — there is nothing invertible to return to — and can only compose forward: a semigroup, which is the GDOS generator's origin and the reason the radial chronology of this Part is one-way at the level of law rather than of initial conditions. The universe does not fail to run backward because entropy happens to have been low; it fails to run backward because the direction it moves along is a fiber on which "backward" stores nothing that could be returned to.

The pairing of the two fibers is then the anatomy of the whole history this paper narrates. The circle *remembers*: it holds the letters and the charges — the flavor lattice whose chirality makes CP violation unavoidable, the Majorana parities, the center winding that is baryon number — so the history has definite content to deposit. The interval *forgets*: it gates, dissipates, and never returns, so each deposit is final — the prompt lock crossing, the zero-dilution clause, the crossover that welds the sphaleron door are all this one statement wearing three epochs' clothes. Read together: *cosmic history is the one-way reading of holonomy data* — the descent along the fiber without holonomy, depositing on the boundary the letters stored on the fiber with it. Cosmic expansion, cooling, and redshift are the boundary shadow of that motion, not its cause; the causal direction is the constitutional one, bulk to boundary. And Sakharov's three conditions, which this paper's baryogenesis appendix assembles into η_B, are in this reading not three arranged ingredients but the anatomy of the two fibers themselves: nonequilibrium is the gated interval, CP violation is the

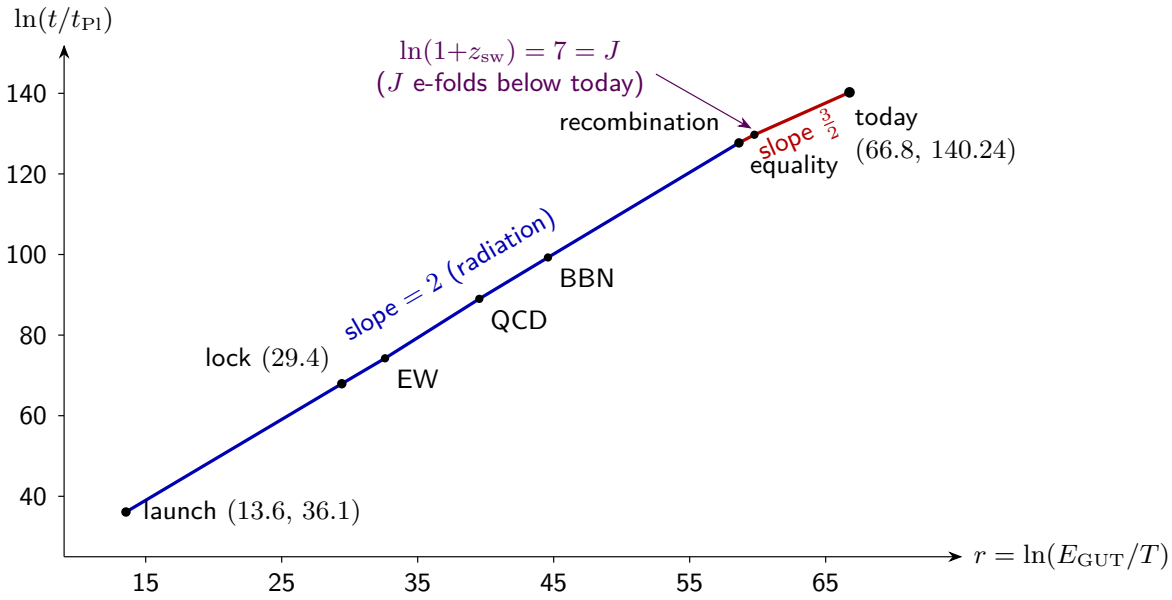

Figure 17: The whole history of time, on one line: boundary time (Planck units, logarithmic) against this paper's radial coordinate ($E_{\rm GUT} = 2.3 \times 10^{16}$ GeV), for eight epochs from the launch to today. The radiation era is a straight line of slope exactly 2; the slope kinks to $\frac{3}{2}$ at equality; the lock crossing sits at $r = 29.4$ (equal to $K = 35$ in the half-Planck convention of the proton and electron companions); the corridor switch sits $J = 7$ e-folds of boundary scale below today; the endpoint's audited decomposition is given in the text. Boundary time is the exponential shadow of the linear descent. (after Arisaka, 2026, P5-CPT)

chiral circle, and the lawful ΔB channel is the circle's topology meeting the weak sectors. The full deployment of this observation — symmetry breaking as symmetry reading, with the determinism graded and its falsifiers dated — is carried by the cohort's symmetry review, GG-W1-Symmetry; this Part records the clause that is cosmology's own: the chronology cannot run backward because its direction is the fiber that cannot remember.

The one-way reading has now been drawn, in the CPT companion's expanded Part IX (Arisaka, 2026, P5-CPT), as a single quantitative figure, and this Part adopts it (Figure 17) because it is the radial chronology of this section made visible. In the radiation era the Friedmann clock and the radial coordinate obey $d\ln t/dr = 2$ exactly — boundary time is the exponential shadow of the linear descent, multiplying by $e^2 \simeq 7.4$ per unit of r — and the whole history from the launch to today is one straight line with a single kink to slope $\frac{3}{2}$ at matter–radiation equality. The endpoint is an audit that closes: $\ln(t_0/t_{\rm Pl}) = 140.24 = 36.1$ (launch) + 91.6 (radiation) + 12.5 (matter+Λ) — the age of the universe as an e-fold ledger on the descent, from standard kinematics and this paper's derived $T_{\rm rh}$, nothing else. (The companion also files, at flagged grade and no higher, the compression $140.24 = 4K \times 1.0017$; its would-be derivation is this paper's e-fold budget program.)

And the three dark episodes of this paper sit on the same anatomy at its largest scale (§31.2): the interval's coordinate opened the chronology, the circle's coordinate still rings within it, and the terminus is the chronology at rest.

25.4 Bottom line: the stage is set

The cosmological stage is now complete. The universe expands because a homogeneous spacetime with nonzero background energy density generically has $H^2 > 0$. It cools because comoving entropy is approximately conserved while the scale factor grows, so the energy per relativistic quantum redshifts as a^{-1}. It runs in one direction because the open-system GFF ledger descends monotonically along the radial ascent $r = \ln(E_{\rm GUT}/\mu)$ out of the GUT edge. The equations are classical and standard; what GG-A10-Cosmos adds is not a replacement of those equations but a constrained inheritance chain for the components that occupy them, beginning at the finite GUT-edge surface and orientated by the open-system descent.

The significance and novelty of Part III. It is worth pausing, before the cosmic history begins, to say plainly what this Part has accomplished and why it is new, because the result is easy to understate. The conventional division of labor in cosmology is sharp: general relativity and the Friedmann equations are *given* — handed down from 1915 and 1922 — and the science consists of measuring the six numbers that fill them. Part III has dissolved that division. It has shown that the gravitational law itself, the cosmological equations it produces, and the form of the numbers that feed them are not three separate gifts but consequences of a single six-dimensional geometry, and it has shown this by derivation rather than assertion.

The novelty is concrete, and worth itemizing. *First*, general relativity is an *output*: the four-dimensional Einstein equation, the Newton constant $G_N = 1/8\pi M_{\rm Pl}^2$, and the scale of the cosmological constant all fall out of the Kaluza–Klein reduction of one six-dimensional Einstein–Hilbert action, with the Planck mass $M_{\rm Pl}^2 = \pi M_6^4(1 - e^{-2K})$ revealed as a geometric overlap integral rather than a fundamental input — a derivation of gravity in a sense no other framework attempts, since in every textbook the Einstein equation is where one *starts*. *Second*, the Friedmann equations are themselves derived, not posited: the cosmological principle that singles them out is read as the geometric attractor of the runtime, so the smoothness of the universe is a selected outcome rather than an initial assumption. *Third*, the beginning is relocated — the observable universe starts at the grand-unification edge $E_{\rm GUT}$, not the Planck scale, because four-dimensional spacetime is the boundary projection of the bulk and is simply undefined above its ultraviolet edge; the Big Bang singularity is not denied as a formal backward limit but denied the standing of a physical initial surface, a clean departure from a century of received intuition. *Fourth*, the source ledger carries no free dials: the contents of ρ and their proportions are closed-form expressions over the fixed integer ledger of GG-A4-Euclid, so the cosmology cannot be retuned to fit. *Fifth*, the cosmological-constant problem is not patched but closed to the percent level — the lock-scale seesaw $m_\Lambda = M_{\rm lock}^2/M_{\rm Pl,red} \approx 6\,\text{meV}$ removes the canonical 10^{120} vacuum-energy discrepancy down to a residual factor of a few, and the ledger coefficient $\zeta_\Lambda = \pi/LK$ takes that residue to $\rho_\Lambda^{1/4} = 2.23\,\text{meV}$ against a measured 2.24 (Derived cond.).

Before the actors enter, notice what the stage itself has already settled. A beginning without a singularity means the oldest question — *what happened at time zero?* — was never quite the right question: on the boundary there is no time zero, only a finite edge where the geometry's description lawfully begins. And an arrow of time that points down an open-system descent means the future is not a mystery appended to physics but a direction built into it. These are philosophical dividends

of a technical construction, and they cost nothing extra.

Taken together, these results promote ΛCDM from a six-parameter fit to a *theorem* of GG-Theory, in exactly the structural sense in which GG-A9-SM (Arisaka, 2026, A9-SM) promoted the Standard Model: the form of every input is forced by the geometry, and a per-parameter provenance triage records, in the open, how completely each form has so far been closed. That is the significance — and it is also where the honesty of the claim matters most. Part III does not assert that every cosmological number is computed today. The uniqueness of the homogeneous attractor, the order-unity coefficient and the sign of Λ, and the primordial-spectrum parameters remain, respectively, conditional or open — flagged, not buried, and handed forward to Parts IV through X, which take up the unfinished closures one at a time. What Part III has settled is the harder and more fundamental point: that there *is* a single derivation to complete, that gravity and the Friedmann stage are theorems of a parameter-free geometry, and that the cosmos and the electron descend from one ledger one story apart. With the stage so understood — the arena fixed, the beginning lawful and finite, the kinematics derived, and the arrow of time orientated by the open-system descent — Part IV brings on the two leading actors: the inflationary modulus that drives the early quasi-de Sitter stage, and the residual F_{GFF} vacuum that fixes the π-locked dark-energy fraction of the present epoch.

26 Part III in plain English: why the universe starts at the GUT edge and not at the Planck scale

The claim of this Part is easy to state and easy to underestimate. General relativity and the Friedmann equations are not assumed here; they are derived as the four-dimensional shadow of the bulk dynamics, and the derivation is set out in full because a reader is entitled to check it rather than take it.

The part of the Part that is genuinely new is the starting point. The textbook history begins at the Planck scale because that is where the known physics runs out, and the interval between the Planck scale and the GUT scale is carried as an unexplained early era. Here the beginning sits at the GUT edge, and it sits there for a structural reason rather than a practical one: the scale chain that fixes the lock scale also fixes where the chain can start. That is a real consequence of the framework and it is exposed — if the universe demonstrably has physics between the two scales that this reading forbids, the reading fails.

The radial chronology introduced at the end of the Part is a bookkeeping device and should be read as one. Writing cosmic history against the radial coordinate rather than against time does not add physics; it makes visible that the early universe is enormous in radial extent and brief in years, and the late universe the reverse. That the arrow of time and the direction of the radial descent point the same way is the substantive claim in the section, and it is argued rather than assumed.

Part IV
Inflation and dark energy: the cosmic GRIP descent

From the lawful beginning of Part III, the universe runs as one open-system descent — the launch of inflation and the terminus of dark energy — governed by the same GRIP geometry that fixes the electron

27 One descent, from the electron to the cosmos

Part III handed us a lawful beginning. The operative origin of the universe is not the formal $a \to 0$ singularity but the finite ultraviolet boundary surface at E_{GUT}; from that edge, ordinary boundary general relativity, sourced by the inherited ledger, sets the universe expanding and cooling. We do not re-derive that beginning here. We take it as given and watch what happens next.

What happens next is a single story. Every system in nature is open, and an open system is read by OGGEP, evolves by GDOS, and runs as the GRIP cascade: it explores its admissible channels (DGI), ranks them by the GFF ledger, breaks the symmetric branch by GSSB, and settles into a geometric attractor. This is not a metaphor borrowed for cosmology. It is the same cascade, the same functional, the same attractor language that the companion electron paper (GG-P3-Electron) uses to fix the electron's charge and mass as a *local* ground-state attractor. The thesis of this Part is that the universe's whole history is that one descent, run at cosmic scale: a launch and a terminus. Inflation is the opening stretch of the descent — the explore phase, the DGI that throws the boundary onto a flat plateau. Dark energy is its terminus — the global, late-time attractor that the boundary relaxes into, the Ground Geometric Attractor, the GGA. The same geometric principle that fixes one electron fixes the vacuum of the cosmos.

Honesty has to be tiered, and we name the tiers once, here, so the reader can carry them through the whole Part. The electron ledger of GG-P3-Electron is *derived*; the Kaxion relic of GG-C1-Kaxion is carried on one *calibrated* scale, declared as such, and is *derived-conditional* after it. The inflationary tilt n_s is now *derived-conditional* on the plateau class, and the amplitude A_s is the launch sector's one *declared input* (its warp candidate flagged in the codebook). The π-locked late-time inventory and the dark-energy fraction are *derived-conditional*: they close up to one named topological step, Conjecture VIII (closed-packet integration), whose final cancellation has since been traced to the declared epoch input — the residual uniformity lemma exact to one part in 10^{15} — and we will say so plainly at the point where it bites. Everything robust in this Part survives even if that one step fails; everything that depends on it is fenced. With the tiers named, we let the descent run.

28 The launch: inflation as the opening of the descent

28.1 Why the universe needs a launch

The descent that began at the $E_{\rm GUT}$ edge cannot reach the universe we observe without an early stretch of accelerated expansion. This is not a stylistic preference; it is forced by four stubborn facts that any successful cosmology must explain:

1. the near flatness of the present universe;
2. the large-scale homogeneity of the CMB across regions that are naively out of contact in a purely decelerating history;
3. the existence of primordial perturbations with an almost scale-invariant scalar spectrum;
4. and the absence, so far, of evidence for a large primordial tensor signal.

The first two are the classic flatness and horizon problems. The third is the origin-of-structure problem. The fourth is already a discriminator among inflationary mechanisms. The historical resolution is well known: Guth in 1981 (Guth, 1981) proposed an "old inflation" driven by a metastable false vacuum, and Linde in 1982 (Linde, 1982) together with Albrecht and Steinhardt in 1982 (Albrecht and Steinhardt, 1982) replaced the false-vacuum picture with "new inflation" driven by a slow-rolling scalar field, repairing the graceful-exit problem of the original proposal.

GG-A10-Cosmos therefore does not try to evade inflation by rhetorical means. It accepts the need for an accelerated early phase, but it denies that the phase must be implemented by introducing a wholly free scalar sector with a potential chosen from a wide menu. Under MSMF that route is too unconstrained. The correct question is not whether inflation happens, but whether the same geometry already needed for the particle and anomaly papers can supply an inflationary stage with much less freedom. In the language of the descent: inflation is the DGI opening move — the boundary exploring the flat throat branch — and the job here is to read that opening off the inherited geometry rather than to bolt a new field onto it.

28.2 Why a free inflaton is forbidden by the inherited program

The earlier preprints already fixed the working rule: do not add a new continuous sector unless the inherited ledger truly cannot do the job. A free inflaton potential of the form

$$V_{\rm inflaton}(\phi) = \text{"choose your favorite polynomial, plateau, hilltop, or hybrid form"} \tag{46}$$

is therefore constitutionally disfavored. It is not mathematically illegal, but it violates the spirit of MSMF unless one can show that no inherited GG-6 modulus can generate the required accelerated phase.

The inherited program already contains two natural vacuum-selection variables:

(i) the throat/radial modulus — the Σ_r scale coordinate — controlling the effective lock depth;[1] and

(ii) the Green–Schwarz normalization modulus Ω, which the anomaly paper says is driven to $\Omega = 1$ during the descent (in its original TeV-attractor language; in the two-basin reading the drive completes by the lock crossing).

For the purpose of cosmology these need not be treated as independent scalar fields. They are different descriptions of the same coarse-grained vacuum-selection direction of the throat. In geometric language one may speak of the Σ_r scale modulus (the coordinate the wider literature would call the radion — see the footnote above for why GG-6 reserves that word). In anomaly language one may track the effective Green–Schwarz normalization Ω. The physical content is the same: there is an inherited modulus whose motion changes the effective four-dimensional vacuum budget and which is forced to settle at a distinguished post-inflationary minimum.

This is what it means for inflation to be lawful rather than inserted. The throat does not need to borrow an inflaton; it already carries the one direction along which the boundary can drift slowly. The launch is built from a part the universe already had.

The campaign's two-fiber reading (§28.15) will sharpen this prohibition into a census: the bulk has exactly one open non-spatial fiber, and hence exactly one candidate inflaton — the problem with the inflaton menu was never taste; it was counting.

28.3 The minimal effective action for the inflationary modulus

At the level of the four-dimensional boundary, the appropriate coarse-grained effective action is

$$S_\varphi = \int \mathrm{d}^4x\sqrt{-g}\left[\frac{M_{\mathrm{Pl,red}}^2}{2}R - \frac{1}{2}(\partial\varphi)^2 - V_{\mathrm{eff}}(\varphi)\right], \tag{47}$$

where φ is the canonically normalized modulus representing the relevant throat/Green–Schwarz vacuum direction. This is not the introduction of a new primitive field. It is the low-energy description of a modulus already present in the inherited geometry.

The detailed microscopic derivation of V_{eff} is deferred to a dedicated follow-up analysis, but the qualitative origin is fixed by the earlier papers and by the geometric hints already carried through the project:

$$V_{\mathrm{eff}} = V_{\mathrm{curv}} + V_{\mathrm{Casimir}} + V_{\mathrm{IR\,boundary}} + V_\Omega. \tag{48}$$

Each term has a clear meaning.

- V_{curv} is the vacuum energy associated with the six-dimensional curvature.
- V_{Casimir} is the finite-size contribution from the compact throat geometry.

[1] A terminological reconciliation, once and for all. The particle papers state, correctly, that GG-6 has *no radion*: the interval length K is a theorem-fixed integer, and no light radius-stabilization scalar exists in today's spectrum (Arisaka, 2026, A4-Euclid; Arisaka, 2026, P1-KK). The cosmological inflaton is not that (non-existent) particle: it is the Σ_r *configuration coordinate* of the radial-scale stratum (Arisaka, 2026, A6-GRIP) — dynamical during the descent, frozen at the ground attractor today, its present-day excitations being the lock cluster, not a light scalar. Where the wider literature says "radion," this paper says the Σ_r scale modulus.

- $V_{\text{IR boundary}}$ encodes the infrared boundary conditions and their later shift under GSSB.
- V_Ω encodes the tendency of the Green–Schwarz normalization to relax toward its physical post-inflationary value $\Omega = 1$.

The crucial point — and this is where the treatment earns the word *derived* — is that the GDOS boundary law fixes the *class* of this potential with no freedom. The genuine conformal anomaly of a boundary law supplies only the Weyl-squared and Euler densities; the scalaron/plateau sector can therefore arise only from the non-conformal curvature coupling of a scalar, and the scale modulus φ — the one-dimensional radial stratum Σ_r — is the only such scalar the inherited geometry possesses. The result is the Starobinsky/ξ class (Mukhanov, 2005; Starobinsky, 1980), with shape and slope fixed by four-dimensional gravity:

$$V_{\text{eff}}(\varphi) = \frac{3}{4}\, M^2 M_{\text{Pl,red}}^2 \left(1 - e^{-\sqrt{2/3}\,\varphi/M_{\text{Pl,red}}}\right)^2 , \tag{49}$$

where the single scale M — the scalaron mass, the curvature of the plateau's edge — is the launch sector's *one declared input*: fixing it by the observed scalar amplitude gives $M = 3.10 \times 10^{13}\,\text{GeV}$, i.e. a plateau height $V_0^{1/4} \simeq 8.1 \times 10^{15}\,\text{GeV}$. (A warp candidate for this one number exists and is recorded in the codebook as flagged, not claimed: $M = M_5\, e^{-K/3}$ reproduces the required value to 87%, equivalently $A_s = 21.5\, e^{-2K/3}$ reaches 75% of the observed amplitude. The tower route of §28.6 scales differently, $M \propto e^{-K/2}$; the two coincide only near $K = 35$, and where the warp form is a flagged guess the tower form rests on a theorem — the per-level coefficient will decide between them.) Everything else about (49) — the exponential flattening, the $\sqrt{2/3}$, the graceful steepening at the edge — carries no dial. This plateau is the surface the descent explores: broad and gently sloped, the geometric arena on which DGI plays out.

28.4 The two basins: a GUT-adjacent plateau, and a lock that is crossed rather than camped

The natural first reading of the descent — and the program's working hypothesis before the present derivation — was that the vacuum sector driving inflation was controlled not by E_{GUT}^4 but by a lock-scale plateau $V_0 \sim M_{\text{lock}}^4$ — *TeV-scale slow-smooth inflation.* That assertion is now retired, and the reason is recorded here at full strength because it is a sign-level result, not a preference. A plateau of height M_{lock}^4 cannot source the observed primordial spectrum on two independent counts. First, the amplitude: $A_s = V/(24\pi^2 \epsilon_V M_{\text{Pl,red}}^4)$ with $V = M_{\text{lock}}^4$ demands $\epsilon_V \sim 10^{-53}$ — a pure dial with no geometric origin. Second, and fatally, the tilt: the approach to a plateau of the form $V_0(1 + e^{-\beta\varphi})$ has *positive* curvature, so the spectrum comes out *blue*, $n_s - 1 \sim +10^{-26}$, against the measured red tilt $n_s - 1 = -0.0351 \pm 0.0042$ — wrong in sign, and by twenty-four orders of magnitude in size. No parameter choice repairs a sign. The corresponding codebook conjecture (the former Lemma V) is retired in Part IX with this refutation attached.

The derived potential (49) instead places the launch where the descent begins. The inflationary

Hubble scale of the primary plateau is

$$H_1^2 \simeq \frac{V_0}{3M_{\rm Pl,red}^2} = \frac{M^2}{4}\left(1 - e^{-\sqrt{2/3}\,\varphi_*/M_{\rm Pl,red}}\right)^2, \qquad H_1 \simeq 1.5 \times 10^{13}\ \text{GeV}, \tag{50}$$

and the plateau sits at $r_{\rm inf} = \ln(E_{\rm GUT}/V_0^{1/4}) = 1.04$ — *one e-fold below the GUT edge*. Inflation is literally the opening stretch of the boundary's descent, the explore phase that the GRIP-Out at $E_{\rm GUT}$ launches; the poetic sentence that opened this Part was, all along, the correct energy assignment.

The lock has lost none of its cosmological standing — it has gained precision. The lock-scale basin is real: the radial Morse function flattens toward its unique critical point, and the curvature scale of that middle basin is

$$H_2 = \frac{M_{\rm lock}^2}{\sqrt{3}\,M_{\rm Pl,red}} \simeq 3.5 \times 10^{-12}\ \text{GeV} \simeq 3.5\ \text{meV}, \tag{51}$$

which the reader should recognize at once: it is the *terminus* seesaw scale $M_{\rm lock}^2/M_{\rm Pl,red}$ of the dark-energy section below. The old TeV-inflation corridor was, unknowingly, a statement about the wrong end of the descent — the number belongs to dark energy's ancestry, not to the CMB's. And the middle basin is *crossed, not camped*: the crossing completes promptly ($N_2 = 0$) on three independent grounds — any GRIP-ledger nucleation action (π, 2π, 3π) lies fifty-six or more orders of magnitude in rate above the supercooling threshold $S_{\rm crit} = 4\ln(M_{\rm lock}/H_2) \simeq 138$; the lock potential as written is a barrier-free saddle; and the large-N enhancement that drives supercooling in warped-throat models is structurally absent here. No sustained secondary inflation occurs, no e-folds are stolen from the primary window, and no relic is diluted. The lock keeps its three jobs — the particle hierarchy, the birth of the Kaxion, and the seeding of the terminus — and gives up only the fourth it never could have performed.

Theorem 28.1 (Two-basin slow-smooth inflation theorem (T-Cosmos-4, revised))**.** *Inflation is the GFF-dominant opening stretch of the Σ_r descent: a plateau of the derived class (49), GUT-adjacent in height, whose graceful exit is the failure of slow roll at the plateau's edge. The lock-scale basin is a distinct, later GRIP-Out crossed promptly ($N_2 = 0$). The slow-roll observables are*

$$n_s = 1 - \frac{2}{N^*}, \qquad r = \frac{12}{N^{*2}},$$

with N^* fixed by reheating matching and undiluted by the crossing (leading order; the exact values on the plateau are in §28.10).

Proof. BI-6 forbids introducing an unrelated free inflaton sector once the throat and its holonomy structure already exist. GDOS fixes the potential class: the conformal anomaly supplies no scalaron, so the plateau must come from the one non-conformal scalar, Σ_r, yielding (49). The sign-level no-go above excludes the lock-scale height; the amplitude then fixes the height as the one declared input. GRIP interprets the entrance onto the plateau as branch exploration and the exit as slow-roll failure, with reheating as the first post-exit attractor export; the nucleation bound forces the later lock

crossing to be prompt, so N^* is the undiluted reheating-matched value. ■

28.5 The first e-fold: the cold start and the forced throw

The edge carries more energy than the plateau can hold, and following that surplus settles two questions at once: what the universe was doing in its first e-fold, and why inflation required no trigger.

The first admissible surface inherits the density of the edge itself, $\rho_{\rm edge} \simeq E_{\rm GUT}^4$ (§18). Measured against the plateau height this is

$$\frac{\rho_{\rm edge}}{V_0} = \left(\frac{E_{\rm GUT}}{V_0^{1/4}}\right)^4 = e^{4r_{\rm inf}} \simeq 63, \tag{52}$$

the fourth power of the same one-e-fold ratio that places the plateau below the edge. The surplus must be stored somewhere, and the derived class offers exactly three channels. The potential cannot take it: $V_{\rm eff}$ is bounded by V_0 and holds one part in sixty-three. A thermal bath cannot take it, twice over: nothing exists on $\Sigma_{\rm GUT}^{\rm Bdy}$ to *be* a bath — its data are geometric (§18) — and radiation is conformally invariant, sourcing no Ricci scalar, so a radiation-dominated edge would park the scalaron at the bottom of (49) and the launch would never occur at all. A hot start does not merely fail to help; in the derived class it forbids inflation permanently. What remains is the kinetic channel:

$$\tfrac{1}{2}\,\dot\varphi^2\big|_{\rm edge} \simeq \rho_{\rm edge} - V_0 \simeq 62\,V_0. \tag{53}$$

Two conclusions deserve to be said plainly. First, *the universe began cold*: no particles, no bath, no temperature — pure geometry, with all but a sixty-third of its energy in the motion of one modulus. The familiar phrase "hot Big Bang" names not the beginning but the reheating exit, some seventy e-folds later (fifty-nine of inflation and fourteen of modulus decay), when the modulus decays and the bath first appears. Second, the sentence that opened this Part — the DGI that *throws* the boundary onto a flat plateau — is not a metaphor. The throw is the forced kinetic component of (53): bookkeeping, not poetry.

The opening stretch of history is therefore a kination sprint. Kinetic energy redshifts as a^{-6}, so the sprint lasts $\Delta N = \ln(62)/6 \simeq 0.69$ e-folds before the plateau height dominates, and while it lasts the field coasts $\Delta\varphi = \sqrt{6}\,\Delta N \simeq 1.69\,M_{\rm Pl,red}$ against Hubble friction. At the end of the sprint V_0 is the only component that has not diluted, and accelerated expansion begins — which is the honest answer to the question of what triggered inflation: *nothing did*. There is no false vacuum, no tunneling event, no switch thrown. A bounded plateau in the presence of red-shifting kinetic energy is a trigger-free ignition; inflation starts because the plateau is the first thing in the descent that cannot dilute.

28.6 The entrance: what the geometry must still supply

Where on the plateau does the descent begin? For an inflaton introduced by hand this is the initial-conditions problem, answered by assumption. Here it is sharper, because the field is not an independent primitive: it is the canonical coordinate of the same radial stratum whose depth is r.

Wherever the scalaron is curvature-dominated, the Einstein-frame dictionary of the derived class, $\varphi = \sqrt{3/2}\, M_{\rm Pl,red} \ln(1 + R/3M^2)$, together with $R = 4\rho/M^2_{\rm Pl,red}$ and the descent parametrization $\rho(r) = E^4_{\rm GUT} e^{-4r}$, closes into a depth-to-field map:

$$\varphi(r) = \sqrt{\tfrac{3}{2}}\, M_{\rm Pl,red} \ln\!\big(1 + e^{4(r_{\rm inf} - r)}\big), \qquad N(r) = \tfrac{3}{4}\big(1 + e^{4(r_{\rm inf} - r)}\big). \tag{54}$$

This is a useful dictionary on the plateau, where the potential dominates and the identification of R with 4ρ is the correct one, and it is how the depth coordinate and the canonical field are to be compared once slow roll has begun. (Stated exactly: the scale the map descends in is $\mu^4 = M^2_{\rm Pl,red} R_J/4$, the curvature scale of the Jordan frame, which coincides with the total energy density only after the exit — the identification is a plateau statement, not a general one.)

It is worth being explicit that (54) may *not* be evaluated at the edge, because the edge is precisely where its premise fails. Section 28.5 showed that the inherited surplus is carried by the kinetic channel, and for a scalar the Einstein-frame trace gives $R = (4V - \dot\varphi^2)/M^2_{\rm Pl,red}$, which at $\frac{1}{2}\dot\varphi^2 \simeq 62\, V_0$ is *negative*. A kination epoch has the opposite curvature sign from a vacuum-dominated one; reading the entrance value off the vacuum relation would be using the dictionary exactly where the paper's own theorem says it does not apply. The entrance is therefore not fixed by the map.

Nor is it fixed by the energy budget, and the reason is structural. The plateau of (49) is asymptotically *flat*: $V(\varphi) \to V_0$ for any sufficiently large φ. The edge condition $\rho_{\rm edge} = 63\, V_0$ therefore fixes the *speed* $|\dot\varphi_0|$ and says nothing whatever about the *position* φ_0. The honest inventory is that the entrance requires two inherited data — the initial displacement and the sign of the initial velocity — of which the budget supplies one, and the geometry is so far asked for the other without answering.

What can be stated exactly is the requirement they must meet. The sprint of §28.5 moves the field by $|\Delta\varphi| = 1.69\, M_{\rm Pl,red}$, and $N^* = 52.6$ e-folds of slow roll require the field to be at $\varphi_* = 5.21\, M_{\rm Pl,red}$ (leading order; 5.30 exact) when the sprint ends. Hence

$$\varphi_0 \;\geq\; 3.5\, M_{\rm Pl,red} \quad \text{(throw up the plateau)}, \qquad \varphi_0 \;\geq\; 6.9\, M_{\rm Pl,red} \quad \text{(throw down it)}, \tag{55}$$

the second being the first plus twice the sprint, since a descending field must begin far enough up to survive the loss. Both are satisfiable: the plateau has no upper end, and no principle of the construction caps the initial displacement. The orientation of the throw is accordingly *not* the decisive question it might appear to be — it moves the bar in (55) by $2|\Delta\varphi|$ and does not by itself admit or exclude the launch. The cascade reading does carry an expectation, that the descent opens at the symmetric configuration and moves away from it, which places the start high and the motion downward; but the identification of the symmetric stage with a position on the radial coordinate is made in the family sector rather than this one, and it is recorded here as an expectation and not imported as a result.

The displacement is bounded from above as well, and the upper bound is constitutional rather than observational. On the plateau the quantum kick per e-fold, $H/2\pi$, overtakes the classical drift,

$\sqrt{3/2}\, M_{\rm Pl,red}/N$, at

$$N_{\rm et} \;=\; 4\pi\sqrt{\tfrac{3}{2}}\,\frac{M_{\rm Pl,red}}{M} \;\simeq\; 1.2\times 10^{6}, \qquad \varphi_{\rm et} \;\simeq\; 17.5\, M_{\rm Pl,red}. \tag{56}$$

Beyond that point the launch is in stochastic self-reproduction: the field's next move is drawn rather than driven, and the descent *samples*. Sampling is precisely what the constitution forbids — the cascade selects unique attractors and never draws from an ensemble, which is the clause the third horizon problem (§73) holds to observational account. A framework whose answer to that problem is determinism cannot place its own launch in the sampling regime. The requirement (55) therefore closes into a corridor:

$$3.5 \;\le\; \varphi_0 \;\le\; 15.8 \quad \text{(up the plateau)}, \qquad 6.9 \;\le\; \varphi_0 \;\le\; 19.2 \quad \text{(down it)}, \tag{57}$$

in units of $M_{\rm Pl,red}$, equivalently $N^* = 52.6 \lesssim N_{\rm tot} \lesssim 1.2\times 10^{6}$. The floor is observational, the ceiling is constitutional, and the entrance is bracketed from both sides.

What fixes the value inside the corridor is, on present evidence, the same physics that owes the amplitude — and both now have a named mechanism. The R^2 coefficient is the threshold sum of the lock-anchored tower of Arisaka (2026, P1-KK): the radial spectrum is a Bessel tower whose base is fixed by that paper's Theorem T-KK-1 at $M_{\rm lock} \simeq 3.85$ TeV, with asymptotic spacing $\pi \times 1.0$ TeV. Summing that tower with standard per-mode threshold coefficients brackets the required $c = M_{\rm Pl,red}^2/12M^2 = 4.6\times 10^{8}$ within an order of magnitude on either side — after the three standard routes miss it outright (minimal matter loops by nine orders, holographic counting at the theory's anchors by five, which is the closure already recorded above made quantitative, and the unwarped interval by eleven). What remains of the amplitude is a single signed number: the per-level heat-kernel coefficient of the tower's field content, required to land in $[1/772,\, 1/100]$ — a finite enumeration against standard tables, cross-checkable through the six-dimensional a_6 coefficient on the warped interval.

The bend no longer waits on that number. A mode contributes to the local R^2 coefficient only while $m^2 \gtrsim R$; as the descent climbs in curvature, $\sqrt{R} = 2\sqrt{N}\, M$ sweeps up the tower and depletes the coefficient that sustains the plateau,

$$\frac{c(N)}{c_0} \;=\; 1 \;-\; \frac{2\sqrt{N}\, M}{M_{\rm top}}, \tag{58}$$

with $M_{\rm top}$ the tower's ultraviolet end. At the pivot the depletion is at or below the two-percent level for either natural choice of $M_{\rm top}$, and the running it induces is at the 10^{-4} level: the observed plateau, and with it n_s and r, are untouched. The bend itself, $N_{\rm bend} = (x\, M_{\rm top}/2M)^2$ for depletion threshold x, lands *inside* the corridor of (57) across the entire reasonable band: for the wider ultraviolet end every threshold from 0.3% to 39% does so, and for the narrower one the tower is stripped entirely by $N \simeq 1.2\times 10^{5}$, likewise inside. The test staked above is therefore *passed*, and passed by shape alone — the tower's spacing, not its per-level coefficient. The plateau ends because rising curvature strips the tower that built it: the R^2 era dismantles its own scaffolding. The exact value of φ_0 inside the corridor still awaits the per-level number and the smoothed crossover, and

(57) remains its honest statement.

Two consequences survive the correction untouched, because neither depends on the value. Whatever φ_0 is, it is geometric data, and geometric data are patch-uniform: the descent enters the plateau at the same point everywhere, so the standard objection that plateau inflation requires improbable initial conditions dissolves structurally — there is no ensemble of initial conditions for the entrance to be improbable in. And the launch remains an event of the throat's first stretch, the same first e-fold in which the overlap integral accrues 86.5% of the Planck mass (Problem E.5.2): gravity and the launch are, in this precise sense, both features of the throat's rim.

28.7 Slow-roll equations and the meaning of "slow-smooth"

The standard slow-roll equations (Lyth & Riotto, 1999; Mukhanov, 2005) follow from Eq. (47):

$$3H\dot{\varphi} + V'_{\rm eff}(\varphi) \simeq 0, \tag{59}$$

$$H^2 \simeq \frac{V_{\rm eff}(\varphi)}{3M^2_{\rm Pl,red}}. \tag{60}$$

The standard potential slow-roll parameters are

$$\epsilon_V := \frac{M^2_{\rm Pl,red}}{2}\left(\frac{V'}{V}\right)^2, \qquad \eta_V := M^2_{\rm Pl,red}\frac{V''}{V}. \tag{61}$$

For the plateau form (49), and in the regime where $e^{-\beta\varphi/M_{\rm Pl,red}} \ll 1$, one obtains to leading order (with $\beta = \sqrt{2/3}$ and $c_1 = 2$ for the form (49), so that $\epsilon_V \simeq \frac{4}{3}e^{-2\beta\varphi/M_{\rm Pl,red}}$ and $N_e \simeq \frac{3}{4}e^{\beta\varphi_*/M_{\rm Pl,red}}$)

$$\epsilon_V \simeq \frac{1}{2}c_1^2\beta^2 e^{-2\beta\varphi/M_{\rm Pl,red}}, \tag{62}$$

$$\eta_V \simeq -c_1\beta^2 e^{-\beta\varphi/M_{\rm Pl,red}}. \tag{63}$$

The number of e-folds from φ_* to the end of slow roll is

$$N_e \simeq \frac{1}{M^2_{\rm Pl,red}}\int_{\varphi_{\rm end}}^{\varphi_*}\frac{V}{V'}{\rm d}\varphi. \tag{64}$$

For the leading plateau approximation this gives

$$N_e \simeq \frac{1}{c_1\beta^2}\exp(\beta\varphi_*/M_{\rm Pl,red})\,, \tag{65}$$

so that

$$\eta_V \simeq -\frac{1}{N_e}, \qquad \epsilon_V \simeq \frac{1}{2\beta^2 N_e^2} \quad \text{(up to order-unity model factors)}. \tag{66}$$

This is the mathematical meaning of *slow-smooth*. The slope and curvature are exponentially small on the plateau, so the modulus drifts rather than falls. The field does not need an exquisitely tuned polynomial inflection point. It only needs the broad exponential flattening naturally produced by the bulk geometry. The descent, on this stretch, is a long gentle glide, not a plunge.

28.8 Scalar tilt, amplitude, and the tensor target: two predictions and one input

The standard scalar spectral tilt (Mukhanov & Chibisov, 1981; Mukhanov *et al.*, 1992) is

$$n_s = 1 - 6\epsilon_V + 2\eta_V. \tag{67}$$

Using Eq. (66), one obtains the characteristic plateau result

$$n_s \simeq 1 - \frac{2}{N_e} + \mathcal{O}\left(\frac{1}{N_e^2}\right). \tag{68}$$

The e-fold count is no longer a corridor to be quoted but a number to be computed: reheating matching (next subsection) gives $N^* = 52.6$, undiluted because the lock crossing is prompt, and hence

$$n_s = 1 - \frac{2}{N^*} = 0.9620, \tag{69}$$

against the measured 0.9649 ± 0.0042 — a deviation of -0.30%, a pull of -0.7 at this leading order, -0.4 for the exact value of §28.10. The honest grade on this number deserves emphasis, because the grade itself is the point: *derived-conditional*, closed modulo the GDOS plateau-class derivation and the prompt-crossing lemma, where the program's working codebook could previously write only *asserted*. The tilt is no longer inherited from the data; it is risked against them.

The scalar amplitude is

$$A_s = \frac{V_*}{24\pi^2 M_{\rm Pl,red}^4 \epsilon_V(\varphi_*)}. \tag{70}$$

The tensor-to-scalar ratio is

$$r = 16\epsilon_V. \tag{71}$$

Eliminating ϵ_V between Eqs. (70) and (71) gives

$$r = \frac{2V_*}{3\pi^2 A_s M_{\rm Pl,red}^4}. \tag{72}$$

This relation is extremely useful because it shows directly that once the inflationary plateau height is fixed, the tensor sector is fixed. On the derived plateau the slow-roll identities collapse it to a pure function of the e-fold count,

$$r = 0.0038 \quad (12/N^{*2} = 0.0043 \text{ at leading order}), \tag{73}$$

a factor of nine below the current BICEP/Keck bound $r < 0.036$ (BICEP/Keck Collaboration, 2021) and several times *above* the projected $\sigma(r) \simeq 10^{-3}$ design floors of the Simons Observatory's enhanced surveys and LiteBIRD (Ade et al., 2019; LiteBIRD Collaboration, 2023). Where the retired TeV reading predicted $r \sim 10^{-52}$ — safety by invisibility — the derived launch stakes a number that this decade's instruments will reach.

The scalar amplitude (70) remains the sector's *one declared input*: it fixes the plateau height (through the scalaron mass M), exactly as the epoch fixes the π-inventory's snapshot. Its warp candidate — $A_s = 21.5\, e^{-2K/3}$, reaching 75% of the observed amplitude with no parameter — is

recorded in the codebook as flagged, not claimed. The falsification handle is now two-sided and sharper than before: a tensor measurement inconsistent with $r = 0.0038$ beyond the launch's spread 0.0037–0.0040, in either direction, kills the derived launch; a tilt measurement excluding 0.9634 at three standard deviations in either direction kills the count (a value below 0.955 would in addition demand a diluted window, $N_2 > 8$, and so kill the prompt-crossing lemma); and a loud stochastic gravitational-wave background attributable to a strongly supercooled TeV-scale transition would refute the same lemma from the opposite side. The plateau that fixes the numbers is the plateau that fires the switches — the virtue and the vulnerability remain the same statement.

28.9 How the descent does the cosmological work

With the launch at the top of the descent, the classical worries invert. The horizon and flatness problems are resolved by the duration of the accelerated stage and the shrinkage of the comoving Hubble radius; and the notorious initial-patch problem that afflicts genuinely low-scale inflation — the need for a pre-inflationary homogeneous region exponentially larger than the inflationary Hubble volume — does not arise here, because the primary plateau begins one e-fold below the GUT edge, where the required patch is microscopic.

The practical criterion is simply

$$\frac{\mathrm{d}}{\mathrm{d}t}(aH)^{-1} < 0, \tag{74}$$

maintained for a sufficient number of e-folds. In a plateau model the total number of e-folds is governed by the field excursion across the flat throat segment, not by a high vacuum scale. Therefore the model can still smooth curvature and establish causal contact across the observed universe provided the plateau is wide enough in the canonical field direction.

In the present program the width is controlled by the logarithmic throat depth and by the same vacuum-selection dynamics that later drives $\Omega \to 1$. That is a coherent geometric mechanism for obtaining an order-10^1 to order-10^2 accelerated interval without introducing a separate inflaton sector; what turns the interval into a number is the reheating history, evaluated below. (An earlier version of this paragraph stated, as "the correct statement at the current stage," that the module merely targets the broad $N_e \sim 50$–60 corridor with the plateau height fixed by the lock scale. Both halves of that sentence are retired: the height is fixed one e-fold below the GUT edge by the declared amplitude, not by the lock scale, and the count is computed, not targeted.)

The symbolic discussion of N_e becomes more informative once one writes the standard matching formula between horizon exit during inflation and the present Hubble scale. For a comoving mode k, one may write schematically

$$N_k \simeq 67 - \ln\!\left(\frac{k}{a_0 H_0}\right) + \frac{1}{4}\ln\!\left(\frac{V_k^2}{M_{\mathrm{Pl,red}}^4 \rho_{\mathrm{end}}}\right) + \frac{1-3w_{\mathrm{rh}}}{12(1+w_{\mathrm{rh}})}\ln\!\left(\frac{\rho_{\mathrm{rh}}}{\rho_{\mathrm{end}}}\right) - \frac{1}{12}\ln g_{\mathrm{th}}, \tag{75}$$

where ρ_{end} and ρ_{rh} denote the energy densities at the end of inflation and after reheating, w_{rh} is the effective reheating equation-of-state parameter, g_{th} counts the relativistic degrees of freedom at reheating, and the pivot is $k = 0.05\ \mathrm{Mpc}^{-1}$, the scale at which n_s, r and A_s are quoted (Planck Collaboration, 2018). (An earlier version printed this relation with the constant 61.6 and without

the pivot term; the two omissions, each about five e-folds, canceled, and the count below was right by compensation. The form above is the standard one.) Equation (75) is standard but important: the required e-fold number is not fixed by the inflationary vacuum scale alone. It also depends on how efficiently the universe reheats and how one matches the inflationary stage to the later hot Big Bang.

Here the matching is evaluated rather than parameterized. The modulus decays at Planck strength, $\Gamma = M^3/M_{\rm Pl,red}^2$ with M the scalaron mass of Eq. (49) (the only channel consistent with the warped structure), so the bath takes over at $T_{\rm rh} \simeq 6 \times 10^{10}\,{\rm GeV}$ ($H = \Gamma$) after roughly fourteen e-folds of matter-like oscillation; inserting $V_0^{1/4} = 8.1 \times 10^{15}\,{\rm GeV}$, $\rho_{\rm end} = V(\epsilon_V = 1)$, $w_{\rm rh} = 0$ over the fourteen e-folds and $g_{\rm th} = 106.75$ into Eq. (75) at the pivot gives $67 - 5.41 - 5.41 - 3.54 - 0.39 = 52.3$. The paper carries

$$N^* = 52.6 \pm 0.5,$$

and because the lock crossing is prompt ($N_2 = 0$), no later episode dilutes it.

The chain closed on itself, and what the half e-fold is made of. The loop can be closed rather than evaluated once: fix N^*, let the declared amplitude set the scalaron mass at the *exact* $\epsilon_V(y_*)$ of §28.10, let that mass set Γ and hence $T_{\rm rh}$, and require Eq. (75) to return the count it started from. That loop has one fixed point,

$$N^* = 52.26, \qquad M = 3.12 \times 10^{13}\,{\rm GeV}, \qquad V_0^{1/4} = 8.11 \times 10^{15}\,{\rm GeV}, \qquad T_{\rm rh} = 6.0 \times 10^{10}\,{\rm GeV},$$

giving $n_s = 0.9632$ and $r = 0.0039$ — a third of an e-fold below the carried value, 0.0002 in the tilt and one unit in the last digit of r. The paper keeps 52.6 because it is the value the codebook, the figures and the companions carry and because nothing observable turns on the difference; but the reader should know what the declared half e-fold is made of, and it is not geometry. It is the coefficient in the width. "Planck strength" means $\Gamma = M^3/M_{\rm Pl,red}^2$ with coefficient one, and this paper does not compute the coefficient: a perturbative gravitational width carries a loop factor of order 48π, which lowers $T_{\rm rh}$ to $5 \times 10^9\,{\rm GeV}$ and the count to 51.4, and the tilt to 0.9626. Two further things are worth saying. First, "the reheating temperature" is itself a convention: the instantaneous estimate $H = \Gamma$ gives $6 \times 10^{10}\,{\rm GeV}$ and is the matching formula's input, while the paper's own Boltzmann integration of the decaying modulus (the computation behind Figure 3) puts the completion of reheating, $\rho_r = 3\rho_\phi$, at $3 \times 10^{10}\,{\rm GeV}$ — the value earlier versions printed — with the bath peaking at $T_{\rm max} \simeq 7 \times 10^{12}\,{\rm GeV}$ on the way; the two are 0.23 e-folds apart in the count, and the completion value returns $N^* = 52.02$, just below the declared interval's floor of 52.1 — which is one reason the interval is quoted as a half e-fold and not less. Second, that interval is not symmetric about the physics: a factor of two in $T_{\rm rh}$ moves N^* by 0.23, so the declared $[52.1, 53.1]$ reaches only a factor 1.6 below the fixed point's reheating temperature and excludes the 48π alternative by 0.7 e-folds. Across the full range a coefficient could take, the count spans about one e-fold, the tilt moves by 0.0006 and the tensor ratio by 4%, neither of which this decade's instruments resolve. A reader who prefers the closed chain should read 52.26 with the three quantities above and change nothing else. The e-fold count that feeds n_s and r above is therefore not a phenomenological corridor but a closed piece of bookkeeping — which is precisely why the present paper keeps the inflation

section inside a full cosmological ledger instead of treating it as a stand-alone potential-fit exercise: the same reheating history that fixes N^* sets the bath on which the baryogenesis of Part V runs.

28.10 The tilt and the tensor ratio evaluated exactly on the plateau

The relations $1-2/N^*$ and $12/N^{*2}$ in Eqs. (69) and (73) are the leading terms of the large-N expansion, and the codebook once carried them as the prediction. On the plateau (49) the slow-roll parameters are closed forms in $y \equiv e^{-\sqrt{2/3}\,\varphi/M_{\rm Pl,red}}$: $\epsilon_V = \frac{4}{3}\,y^2/(1-y)^2$, $\eta_V = -\frac{4}{3}\,y(1-2y)/(1-y)^2$, the end of inflation is $y_{\rm end} = 1/(1+2/\sqrt{3}) = 0.464$, and the e-fold count from y_* to $y_{\rm end}$ is $N = \frac{3}{4}[1/y_* + \ln y_* - 1/y_{\rm end} - \ln y_{\rm end}]$. Solving for y_* at $N^* = 52.6$ and evaluating $n_s = 1-6\epsilon_V+2\eta_V$, $r = 16\epsilon_V$ without expansion gives

$$n_s = 0.9634, \qquad r = 0.0038 \qquad (N^* = 52.6;\ \text{leading order } 0.9620,\ 0.0043), \tag{76}$$

the next-order terms moving the tilt up by 0.0014 and the tensor ratio down by twelve percent. The exact pair is what the framework predicts and what the Dev tables score; the leading-order pair is retained where it is quoted as the Starobinsky relation. Against *Planck* 2018 (0.9649 ± 0.0042) the exact tilt sits at -0.4σ.

The 2025 tilt, and the tension it states. The comparison must also be made against the measurements that exist now, and the paper does not soften it; the dedicated microwave-background companion (Arisaka, 2026, C5-CMB) made it first. The joint analysis of SPT-3G, *Planck*, ACT and BICEP/Keck gives $n_s = 0.9682 \pm 0.0032$, and with DESI baryon-acoustic oscillations added, 0.9728 ± 0.0029 (arXiv:2512.10613, 2025); the derived 0.9634 sits 1.5σ and 3.2σ below them. There is no input in the chain that moves the tilt upward: the plateau class is forced, the amplitude is declared, and a lower reheating temperature moves N^* down and the tilt further from the data.

So the position must be stated without a rider, because the obvious rider does not hold. One might say the two 2025 combinations should be made to agree before either is used; they do agree — they differ by 0.0046 on a combined uncertainty of 0.0043, which is 1.1σ. The honest statement is therefore that **this framework's tilt criterion, three standard deviations in either direction, is met today by the DESI-combined value (3.2σ) and not by the CMB-only value (1.5σ)**: the tilt is the framework's one currently failing test, the failure is two-sided and cannot be relieved by any dial the paper holds, and that it will be settled — in one direction or the other — by the next round of ground-based CMB polarization and by whether the DESI baryon-acoustic preference for a higher tilt survives its own $w_0 w_a$ degeneracy. A framework that computes the tilt has to live with the measurement of it; that is what computing it means.

28.11 Graceful exit as a geometric spontaneous symmetry breaking

The phrase *graceful exit* means that inflation ends dynamically and continuously rather than by an abrupt collapse of the theoretical description. In the descent language it is the break: the GSSB that closes the explore phase. In GG-6 this exit is the failure of slow roll at the edge of the primary plateau; the GSSB grammar below describes it and, later, the seam crossing.

One sentence of bookkeeping before the mechanism, because the derivation separates two events

that the retired TeV reading fused. The *graceful exit of inflation* is the failure of slow roll at the edge of the primary plateau — $\epsilon_H = 1$, the modulus rolling to its minimum — and it happens at the top of the descent. The *lock* GSSB — the infrared seam transition at $M_{\rm lock}$ that births the Kaxion and seeds the terminus — is a distinct, later GRIP-Out, crossed promptly during the radiation descent. The symmetry-breaking machinery below describes the general grammar of both; where earlier text attributed the end of inflation to the infrared seam, the seam event now stands on its own as the second crossing.

The guiding idea is simple. During the inflationary plateau, the modulus φ drifts slowly, and the Green–Schwarz normalization sector and compact-angle holonomy have not yet reached their final post-inflationary branch. As the modulus reaches the plateau's edge, the effective mass matrix of the residual vacuum sector changes sign along one direction. The order parameter can be represented abstractly by a holonomy variable Θ or, equivalently, by the deviation of the physical normalization from its attractor value. The effective free-energy functional then takes the schematic form

$$\mathcal{F}_{\rm IR}(\Theta) = \mathcal{F}_0 + \frac{1}{2} m_\Theta^2(\varphi)\Theta^2 + \frac{\lambda_\Theta}{4}\Theta^4 + \cdots , \tag{77}$$

with

$$m_\Theta^2(\varphi) > 0 \quad \text{on the plateau}, \qquad m_\Theta^2(\varphi_c) = 0, \qquad m_\Theta^2(\varphi) < 0 \quad \text{after the critical point.} \tag{78}$$

At $\varphi = \varphi_c$, the symmetric branch loses stability and the system selects one of the broken branches. This is exactly the GDOS language of DGI plus GSSB plus GA translated into cosmology — the same three-move anatomy the electron paper uses, now playing out on the cosmic plateau.

Operationally, the consequences are threefold.

(i) The slow-roll conditions fail because the plateau potential steepens as the modulus descends toward its minimum ($\epsilon_H \to 1$).

(ii) The quasi-de Sitter phase ends at the edge of the primary plateau; the infrared seam transition at $M_{\rm lock}$ is the later, separate crossing named above, not the end of inflation.

(iii) The modulus oscillates about its minimum and decays at Planck strength, $\Gamma = M^3/M_{\rm Pl,red}^2$: the stored energy reaches ordinary relativistic degrees of freedom after roughly fourteen e-folds of matter-like oscillation, at $T_{\rm rh} \simeq 6 \times 10^{10}$ GeV ($H = \Gamma$) — the same fourteen-e-fold reheating history that fixes N^* above, not a dump at the end of inflation.

The end of inflation may be parameterized by the standard condition

$$\epsilon_H := -\frac{\dot{H}}{H^2} = 1 \tag{79}$$

or equivalently by the breakdown of the slow-roll inequalities. The physical interpretation, however, is richer than in ordinary scalar inflation: inflation ends because the modulus reaches the edge of the plateau that the throat's vacuum selection laid down.

28.12 Reheating and the return to the hot Big Bang

Once inflation ends, the universe must return to an ordinary thermal history. Let $\xi_{\rm rh}$ be the fraction of the vacuum energy at the end of inflation transferred efficiently into the relativistic plasma. The reheating temperature is then estimated by

$$T_{\rm rh} \simeq \left(\frac{30}{\pi^2 g_{*,\rm rh}} \, \xi_{\rm rh} V_{\rm end} \right)^{1/4} . \tag{80}$$

The decay is Planck-strength — $\Gamma = M^3/M_{\rm Pl,red}^2$, the modulus's only consistent channel — so $\xi_{\rm rh}$ is not order unity: the oscillating condensate leaks for about fourteen e-folds, the bath spiking to $T_{\rm max} \simeq 7 \times 10^{12}\,\text{GeV}$ and matched at

$$T_{\rm rh} \simeq 6 \times 10^{10}\ \text{GeV}, \tag{81}$$

seven orders above the TeV-scale heavy states the resonant mechanism uses (the $\gtrsim 10^9\,\text{GeV}$ window of hierarchical leptogenesis is not this paper's), comfortably below the GUT edge (monopoles stay diluted), and far above the lock (which the cooling bath then crosses at $t \simeq 1.6 \times 10^{-14}\,\text{s}$, T of order a few TeV, releasing the Kaxion's angle and producing the KK sector as ordinary cooling thresholds). This is one of the most attractive features of the scenario, and it is the first attractor export of the descent: the hot Big Bang is not discarded but recovered immediately after the break, with a reheating temperature derived from the modulus decay itself — the cooling bath then crossing, in due course, the same lock scale that governs the particle hierarchy. Later baryogenesis and dark-matter calculations can therefore begin from a thermodynamic stage whose characteristic scale was not chosen separately for cosmology.

28.13 A compact comparison and the bottom line of the launch

Table 8: Schematic comparison between ordinary high-scale inflation and the minimal GG-6 inflation module. Each row poses one question and answers it twice, in the left column for a typical high-scale inflaton model and in the right for the minimal implementation here. The comparison is schematic by design: it is a map of where the two accounts differ in kind, not a numerical competition between them, and no row should be read as a measurement.

Question	**Typical high-scale inflaton model**	**Minimal GG-6 implementation**
Primitive field content	New scalar sector usually inserted by hand	No new primitive inflaton; low-energy modulus of inherited throat/GS sector.
Potential freedom	Large menu of tunable potentials	Plateau generated by curvature + Casimir + IR-seam competition.
Characteristic vacuum scale	Often near 10^{16} GeV, chosen	$V^{1/4} = 8.1 \times 10^{15}$ GeV, fixed one e-fold below $E_{\rm GUT}$ by the derived class + the one input.
Tensor prediction	Can be observable depending on scale and slope	$r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order): a firm two-sided target, decisive for the Simons Observatory and LiteBIRD.
Secondary structure	Typically none, or ad hoc	The lock basin: real, crossed promptly ($N_2 = 0$); births the Kaxion; seeds the terminus.
Exit mechanism	Model-dependent waterfall or steepening	Slow-roll failure at the plateau's edge ($\epsilon_H = 1$); modulus decay and reheating; the IR-seam GSSB is the later crossing.
Relation to later dark sector	Often separate from dark matter and dark energy	Same event chain naturally linked to Kaxion and residual vacuum term.

The inflationary module of GG-A10-Cosmos is therefore built from one sharp proposition:

> The GDOS boundary law forces the plateau class; the amplitude declares its one input; everything else — the tilt, the tensor ratio, the e-fold count, the reheating temperature, the promptness of the lock crossing — is computed. Inflation is the opening stretch of the same descent whose second crossing stocks the dark sector and whose terminus is dark energy.

The retired proposition is kept on the record for honesty:

> A TeV-scale slow-smooth inflationary stage can be generated by the same six-dimensional throat that already underlies the particle hierarchy, with the relevant modulus represented either as the throat radion or as the Green–Schwarz normalization direction, and with

graceful exit occurring through a geometric spontaneous symmetry breaking at the infrared seam.

Its most important quantitative consequence is now Eq. (73): $r = 0.0038$, reachable this decade. The retired reading's untestable $r \sim 10^{-52}$ is replaced by a number that can lose.

28.14 The launch, graded

The derivation chain of the launch, with the grade of every link stated in the codebook's vocabulary, runs as follows (Figure 18 draws its two halves: the derived potential and the (n_s, r) predictions): the inflaton is the Σ_r scale modulus (*forced* — the only one-dimensional modulus the theory possesses); the potential class is Starobinsky/ξ (*derived* — the GDOS non-conformal route); $n_s = 0.9634$ and $r = 0.0038$ (*derived-conditional* — given the class and the prompt crossing); $N^* = 52.6$ and $T_{\rm rh} \simeq 6 \times 10^{10}\,\text{GeV}$ (*derived-conditional* — reheating matching at Planck-strength width, with the width's coefficient taken as one; §28.10); $N_2 = 0$ (*derived-conditional* — the ledger-action bound, with fifty-six orders of margin); the amplitude A_s (*the one declared input*, warp candidate at 75% flagged); and the entrance displacement φ_0 (*open* — bracketed on both sides by the corridor of §28.6, its selecting bend now mechanized by the mode depletion of (58) and landing inside the corridor by shape alone). Two routes toward paying the amplitude have been attempted and closed — warped decay (excluded by seventeen orders) and holographic central-charge counting (no large-N regime at the theory's anchors) — and one remains open and sharply posed: the signed per-level heat-kernel coefficient of the lock-anchored tower of Arisaka (2026, P1-KK), required in $[1/772,\ 1/100]$ — the amplitude's mechanism being the tower's threshold sum, with the bend test already passed by the depletion law of §28.6, and with the same number deciding between the flagged warp form $M \propto e^{-K/3}$ and the tower form $M \propto e^{-K/2}$. The falsifiers are staked above; the retired TeV reading and its codebook conjecture are recorded, with their refutation, in Part IX.

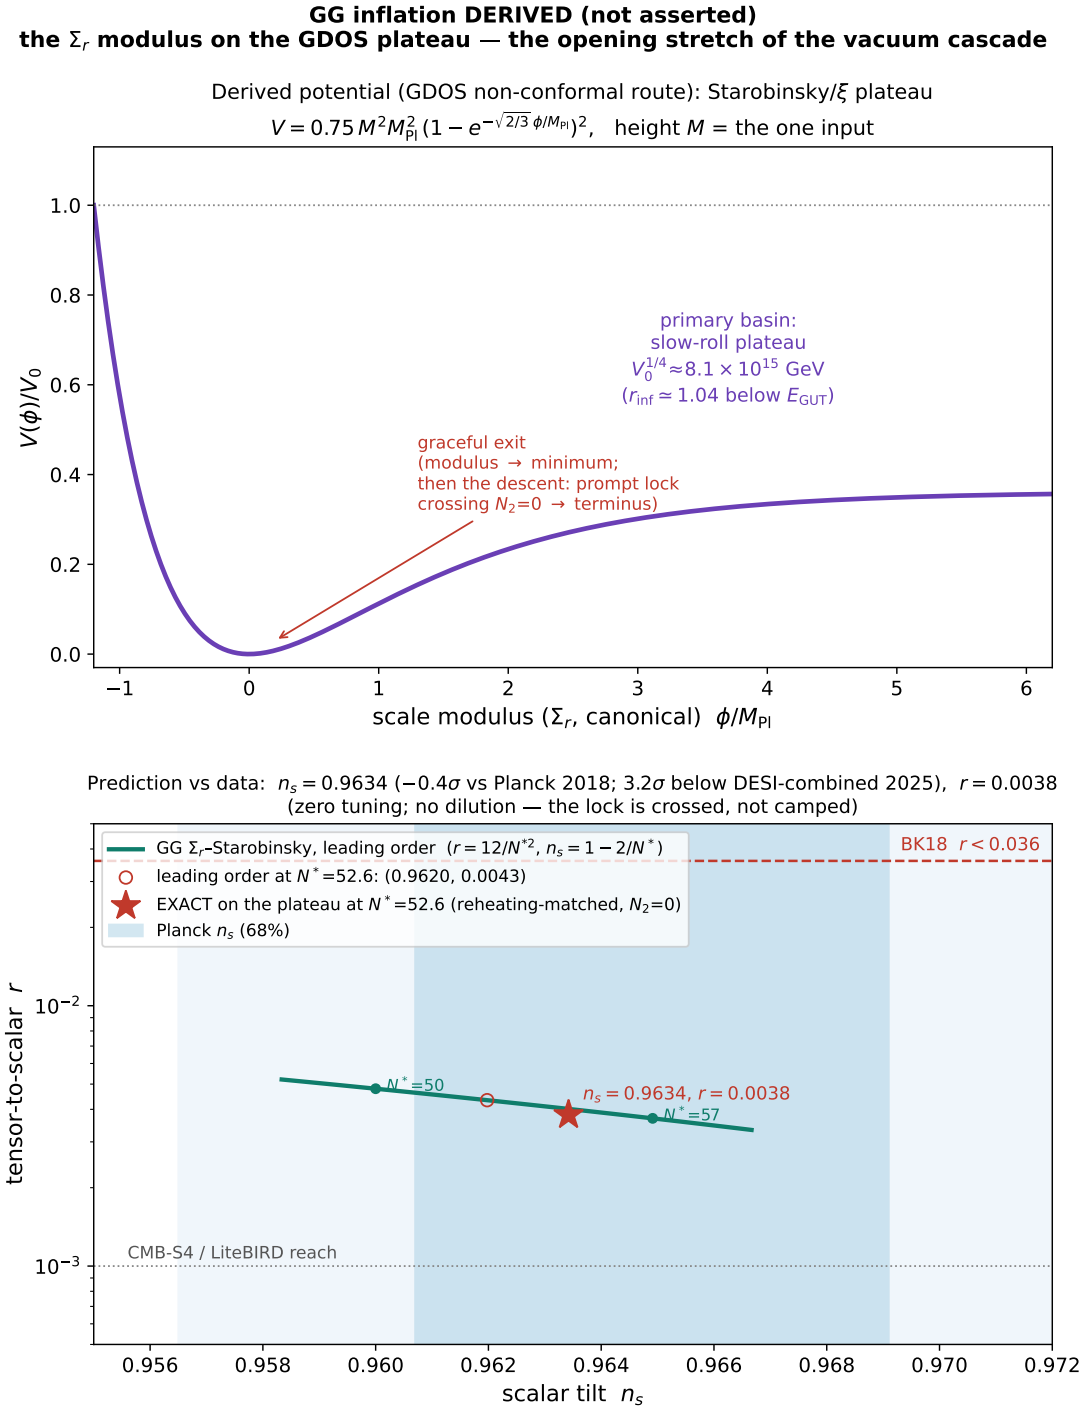

Figure 18: The derived launch. *Top:* the Σ_r plateau of the derived class (49) at $V_0^{1/4} \approx 8.1 \times 10^{15}$ GeV, one e-fold below the GUT edge — shape fixed by gravity, height the one input; graceful exit, prompt lock crossing, descent to the terminus. *Bottom:* $N^* = 52.6$ (star): $n_s = 0.9634$ (0.4σ, bands), $r = 0.0038$ — a factor of nine below BICEP/Keck (dashed), decisive this decade (Simons Observatory/LiteBIRD, dotted).

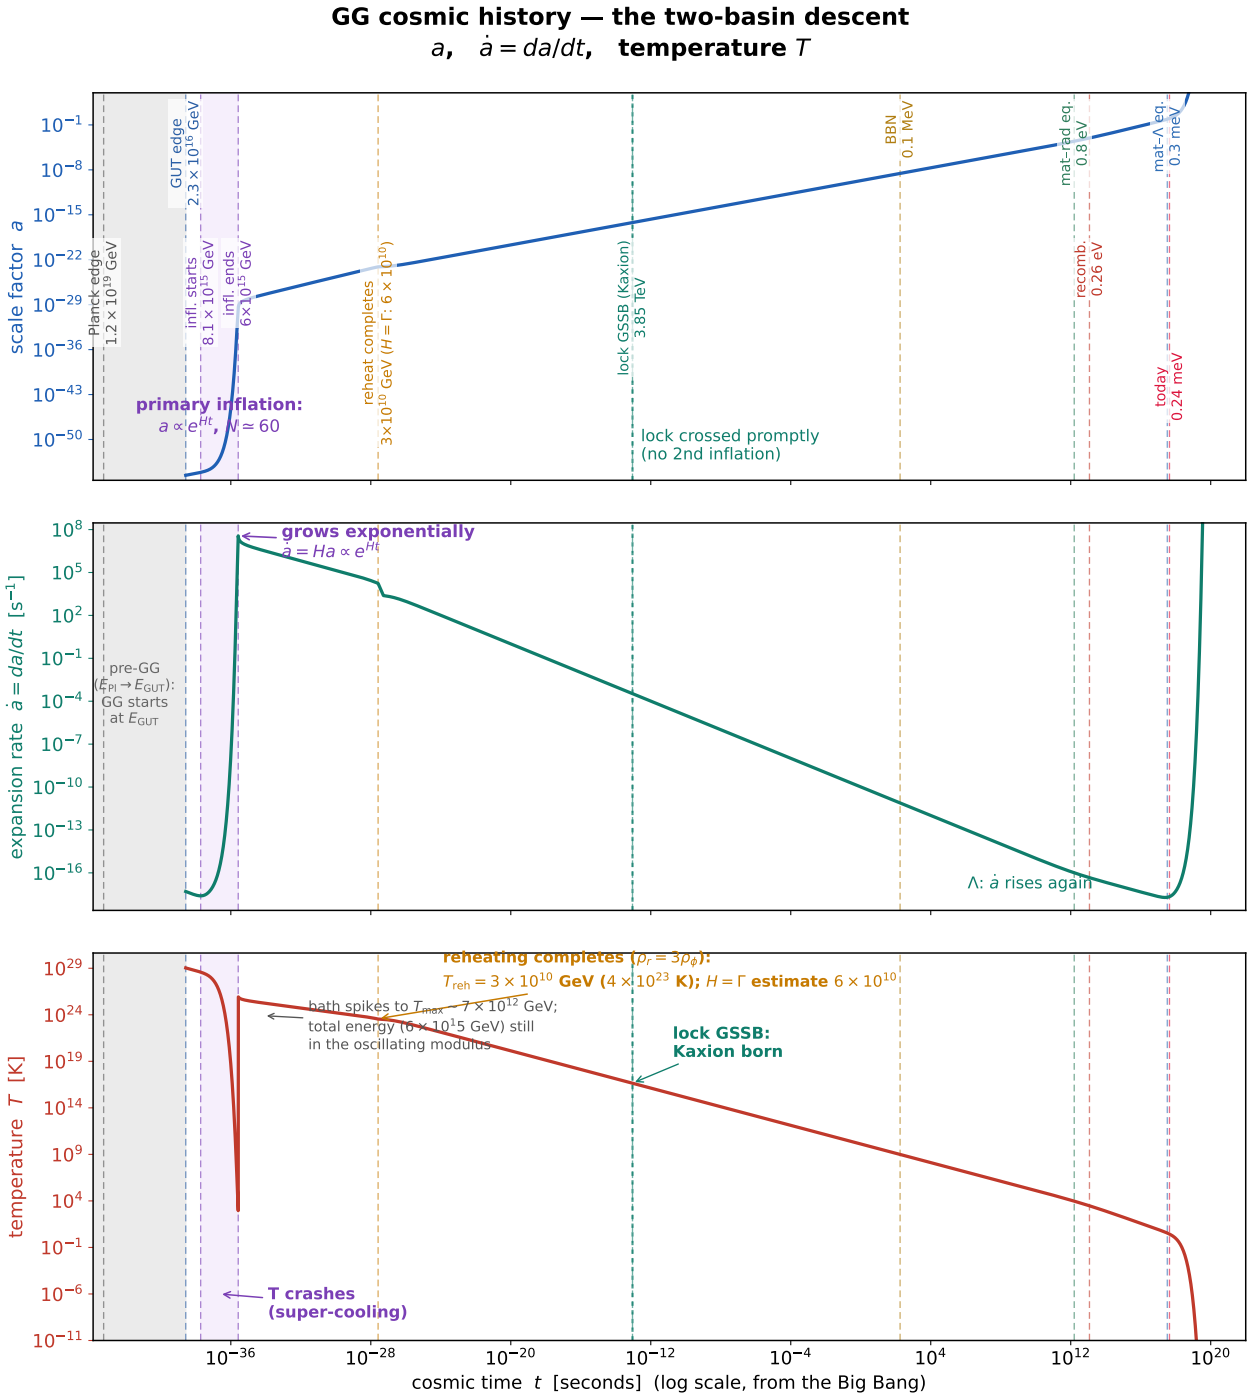

Figure 19: The two-basin descent, numerically integrated (Friedmann + inflaton + reheating fluid, smooth by construction). Scale factor a, expansion rate $\dot{a} = da/dt$, and temperature T from the Planck edge to 10^{22} s. The gray band is pre-GG ($E_{\rm Pl} \to E_{\rm GUT}$: the geometry begins at the GUT edge); the purple band is the primary inflation ($N \simeq 60$, $a \propto e^{Ht}$); the temperature crashes through super-cooling, the bath spikes to $T_{\rm max} \simeq 7 \times 10^{12}$ GeV while the modulus decays, and the bath takes over at $T_{\rm rh} \simeq 6 \times 10^{10}$ GeV ($H = \Gamma$; completion, $\rho_r = 3\rho_\phi$, at 3×10^{10}); the lock is crossed promptly at $t \simeq 1.6 \times 10^{-14}$ s (Kaxion angle released, no second inflation); and in the far future $\dot{a}$ rises exponentially again — the two de Sitter epochs bracketing the whole history. Every vertical epoch line carries its energy scale.

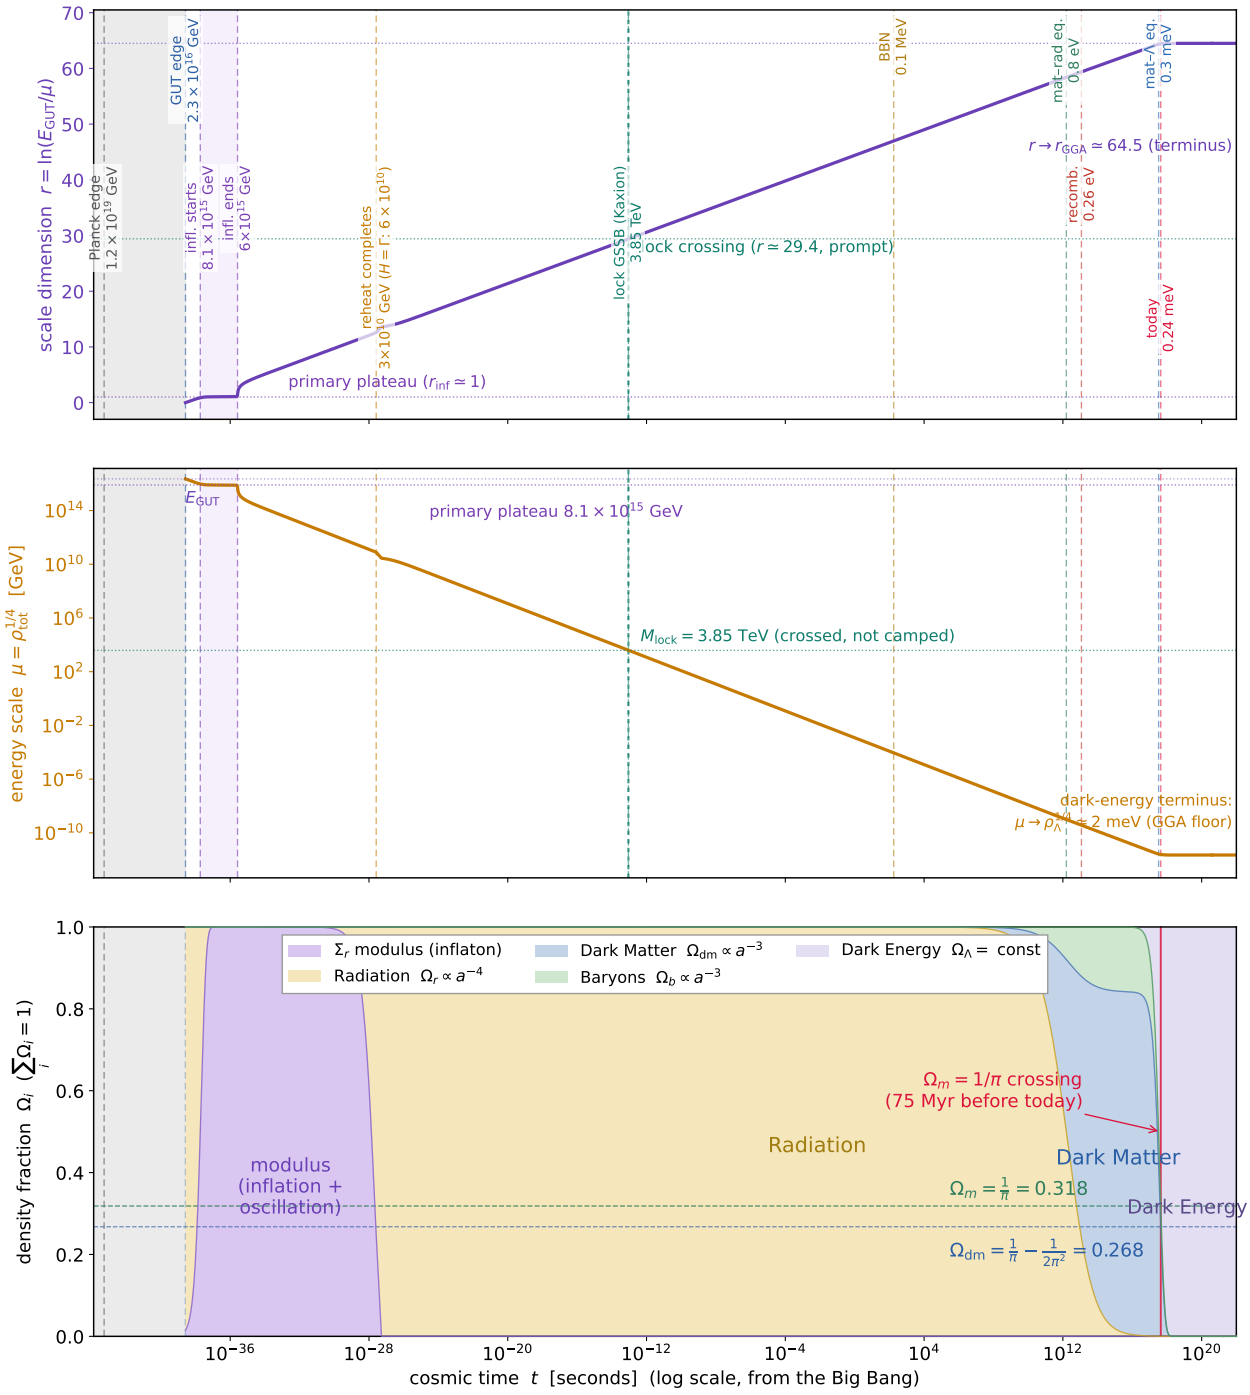

Figure 20: The descent's own variables. *Top:* the scale dimension $r = \ln(E_{\rm GUT}/\mu)$: flat on the primary plateau ($r_{\rm inf} \simeq 1$), descending monotonically through the prompt lock crossing ($r \simeq 29.4$, crossed, not camped), and flattening at the terminus $r_{\rm GGA} \simeq 64.5$. *Middle:* the energy scale $\mu = \rho_{\rm tot}^{1/4}$, from the primary plateau ($8.1 \times 10^{15}\,$GeV) through $M_{\rm lock}$ to the dark-energy floor ($\approx 2\,$meV). *Bottom:* the cosmic energy budget $\Omega_i(t)$ on the seconds axis — the Σ_r modulus band through inflation and reheating, the long radiation era, the matter interlude, and dark energy's take-over, with the $\Omega_m = 1/\pi$ crossing (75 Myr before the present) marked in red.

The relaxation relic: a signature at the largest scales (F25). One further prediction rides on this module, inherited from the quantum-mechanics companion (Arisaka, 2026, A8-QM) and recorded in the falsifier ledger as F25. In the cohort's reading the Born statistics of quantum outcomes are not postulated but produced: the boundary law's own dissipation relaxes the bulk configurational statistics to the Born measure, monotonically, at a per-exchange rate set by the port coupling. A mode that leaves the horizon during inflation decouples from the port stream and freezes whatever divergence from the Born measure it still carries, so the earliest- exiting — largest-scale — modes retain the most. The generic relic is a power law anchored at the largest observable scales, a primordial-statistics deficit $\xi(k) = 1 - A\,(k/k_0)^{-\gamma}$, with k_0 the first observable mode out and γ set by the collision rate and mixing weight on the plateau — and, because the initial configuration is constitutionally determined rather than free, the amplitude A is in principle *computable*, not a fit parameter. Detection of a large-scale statistical anomaly incompatible with the computed form — or a relic where the determined initial data forbid one — falsifies the relaxation account; exact Born statistics at all scales bounds the pre-inflationary relaxation time from below. The quantification of A and γ within this module's determined descent is left to a dedicated study.

28.15 Inflation on the two fibers: the interval's coordinate, rolling once

The launch is now graded; this closing subsection reads it on the anatomy the cohort's recent campaign has made exact (GG-W1-Symmetry; Arisaka, 2026, P5-CPT; Arisaka, 2026, C1-Kaxion): *holonomy requires closed loops, and an interval has none.* The compact circle S^1_θ is nothing but closed loops — everything on it is holonomy, quantized and protected, its transport a group: the fiber that remembers. The scale interval I_r can hold no holonomy — its transport composes only forward, a semigroup: the fiber that forgets. The dark-matter companion has already drawn one portrait on this anatomy: the Kaxion is *the anomaly ledger's own phase — the boundary dual of the tensor the ledger forces — promoted to a boundary field* (Arisaka, 2026, C1-Kaxion). The launch supplies the mirror, and it was written into this Part's own derivation before the campaign named it: the inflaton φ is the canonically normalized coordinate of the radial stratum Σ_r — *the interval's own coordinate, promoted to a boundary field.* The two great scalar episodes of cosmology are the bulk's two non-spatial coordinates, each taking the stage exactly once.

The census: one interval, one inflaton; one circle, one axion. The uniqueness step of this Part deserves restating in the campaign's terms, because it is the structural prize. The GDOS boundary law's genuine conformal anomaly supplies only the Weyl-squared and Euler densities; a plateau sector can therefore arise only from the non-conformal curvature coupling of a scalar; and the inherited geometry possesses exactly one such scalar — the radial modulus, because the bulk has exactly one open non-spatial fiber. Meanwhile the compact sector supplies the tensor's dual and the circle's lifted holonomies, exactly one light angle among them (Arisaka, 2026, C1-Kaxion), because the bulk has exactly one closed non-spatial fiber. Read together: *one interval ⇒ one inflaton; one circle ⇒ one axion.* The two great embarrassments of riches in the literature — the inflaton zoo ("choose your favorite potential") and the string axiverse (hundreds of light axions) — are, on this anatomy, the same mistake made twice: uncounted fibers. The dimensional census of the bulk *is*

the scalar field content of cosmology, and MSMF's prohibition of a free inflaton (§28) is thereby sharpened from a parsimony rule into a counting statement.

The episode's shape is the fiber's character. Why does inflation *end*, while the Kaxion *rings on*? Because each coordinate inherits its parent fiber's topology. The circle's coordinate is periodic: displaced, it oscillates — there is always somewhere to come back to; protected by periodicity, it cannot get heavy; its relaxation is a ringing that persists to this day. The interval's coordinate is non-periodic and unprotected: its potential is a plateau with an edge, it traverses once, slow roll fails at $\epsilon_H = 1$, and the episode is over — *inflation is transient because the interval is an interval.* Even the two potentials wear their fibers' faces: the Kaxion's is a cosine — topology-generated, periodic — while the inflaton's is the warp exponential $\left(1 - e^{-\sqrt{2/3}\,\varphi/M_{\mathrm{Pl,red}}}\right)^2$ — the e-folding of the descent itself, bent into a potential. And the once-only logic of the descent applies at the launch with full force: the roll happened once, in one direction, and cannot recur, because the fiber it ran along has no inverse transport with which to re-ascend the plateau — the same gated-interval logic by which the later lock crossing is prompt and unrepeated (Arisaka, 2026, P5-CPT). (Stated, like the rest of this subsection, as a reading of proved structures, not a new theorem.)

The interval's numbers. Sorted on the anatomy, the launch's observables are pure interval bookkeeping. The e-fold count $N^* = 52.6$ is a length of descent, undiluted because the crossing is prompt; the tilt $n_s = 1 - 2/N^* = 0.9620$ and the tensor ratio $r = 12/N^{*2} = 0.0043$ (at leading order; 0.9634 and 0.0038 exact) are functions of that length alone; and the one declared input — the plateau height, with its flagged $e^{-K/3}$ warp candidate — is an *address* on the fiber that forgets, exactly where a two-fiber theory expects its declared inputs to sit. Nothing quantized appears anywhere in the launch's output, and the anatomy says why: the interval has no holonomy to write with. What inflation makes — the stretched, frozen perturbation spectrum — carries the descent's numbers; what inflation cannot touch is the ledger: charges, windings, the integers of BI-6 ride through the launch unaltered, because the launch is the other fiber's episode.

The circle's two cameos. Inflation is the interval's solo, but the circle appears exactly twice, at the boundaries of the act, and both cameos are already this Part's. *The exit is the circle's first word*: the graceful exit is a GSSB whose order parameter is the holonomy variable Θ — as the modulus approaches the seam, the compact-angle mass matrix flips sign, and the circle's symmetry breaking ends the interval's episode. *The handoff seeds the next act*: the later lock GSSB — the descent's second saddle — leaves the compact angle displaced, and that displacement, relaxing, is the dark matter (§35.8). The two fiber-episodes are chained: the interval's exit hands the circle the displacement that becomes its ring. In one sentence: *inflation is the fiber that forgets, opening the story; dark matter is the fiber that remembers, carrying it; and the terminus (§29.10) is what remains when both have finished speaking.*

29 The terminus: dark energy as the Ground Geometric Attractor

29.1 What has to be explained at the end of the descent

The launch left a hot, expanding, nearly flat universe and a tiny residual vacuum term. We now run the descent forward to where it stops. The late-time acceleration problem contains two logically distinct questions.

1. Why does the present universe contain an approximately constant energy-density component whose pressure is sufficiently negative to accelerate the expansion?

2. Why is that component numerically comparable to the matter density precisely in the present epoch rather than only in a remote past or future era?

The first is the dark-energy question. The second is the cosmic coincidence question. A framework can answer one without answering the other. Part IV aims to at least formulate both within the same inherited ledger.

The historical arc is familiar. The cosmological-constant term Λ was introduced by Einstein in 1917 (Einstein, 1917) to allow static cosmological solutions and was famously withdrawn after Hubble's expansion measurement; Weinberg's 1989 review (Weinberg, 1989) surveyed the dramatic mismatch between the observed value of the vacuum energy and naturalness estimates from quantum field theory. The observational reintroduction of Λ came in 1998–1999 with the Type-Ia supernova results of Riess et al. (Riess et al., 1998) and Perlmutter et al. (Perlmutter et al., 1999), which established that the late-time expansion is accelerating. The present section retains the empirical Λ but reinterprets its scale through the residual GFF ledger of the inherited bulk. The cosmological-constant and dynamical dark-energy literature, against which the framework's residual-vacuum branch is positioned, and the de Sitter horizon thermodynamics relevant to the late vacuum-dominated phase, are reviewed in (Carroll, 2001; Copeland *et al.*, 2006; Gibbons & Hawking, 1977; Peebles & Ratra, 2003).

The background condition for acceleration follows immediately from Eq. (8): for a flat universe without an explicit curvature term,

$$\frac{\ddot{a}}{a} > 0 \qquad \Longleftrightarrow \qquad \rho + 3p < 0. \tag{82}$$

A component with equation-of-state parameter $w < -1/3$ can therefore drive acceleration. A strict cosmological constant has $w = -1$. The standard ΛCDM answer is simply to write

$$\rho_\Lambda = \text{const}, \qquad p_\Lambda = -\rho_\Lambda, \tag{83}$$

and fit the constant. GG-A10-Cosmos asks a narrower and more structural question: can the same term be re-read as a residual boundary export of the inherited GFF and BI-6 geometry rather than as an inserted constant with no pedigree? In the descent language, the question is whether Λ is the global attractor the boundary relaxes into — the GGA.

29.2 Residual vacuum energy as a GFF/CEP boundary functional

The earlier papers already suggested the key conceptual route. The GFF ledger is

$$\mathsf{GFF} = \mathsf{GEP} + \mathsf{CEP}, \tag{84}$$

where GEP is the geometric-energy block and CEP is the entropic-curvature block. In the particle paper this ledger organizes masses, couplings, and vacuum selection. In cosmology the natural next step is to interpret the late-time vacuum term as a small residual zero-mode export of the same ledger.

At the level of the four-dimensional effective action one writes the residual term as

$$S_{\Lambda,\mathrm{eff}} = -\int \mathrm{d}^4x\sqrt{-g_4}\,\rho_\Lambda^{\mathrm{eff}}. \tag{85}$$

The GG-6 interpretation is that

$$\rho_\Lambda^{\mathrm{eff}} = \rho_{\mathsf{GFF},\mathrm{res}} = \rho_{\mathsf{GEP},\mathrm{res}} + \rho_{\mathsf{CEP},\mathrm{res}}, \tag{86}$$

where "res" denotes the tiny leftover zero-mode contribution after the inflationary GSSB transition, the selection of the physical normalization, and the coarse-graining of heavy throat modes. In the minimal version of the program the dominant late-time term is expected to be the CEP residue, because the entropic-curvature block is exactly where the framework stores slow, coarse-grained, late-time free-energy information. But the decomposition in Eq. (86) is the honest statement: at the present stage one should not overclaim a complete microscopic split.

What matters first is that the late-time vacuum term is no longer conceptually independent. It is a remnant of the same throat and same vacuum-selection event that already appeared in the launch. The terminus is made of the descent's own leftovers — the same GFF functional, read at its floor.

29.3 A natural meV scale from the lock-scale seesaw

The most notorious feature of dark energy is its tiny scale (Carroll, 2001; Weinberg, 1989). Observationally one may summarize it by saying that

$$\rho_\Lambda^{1/4} \simeq 2.224\ \mathrm{meV}, \tag{87}$$

read off the derived background $\Omega_\Lambda = 1 - 1/\pi$ at $H_0^{(0)} = 66.6$, the two-road reading (Route I 66.8); 2.243 meV on the Planck 2018 inventory — the difference is Planck's 67.4 and fractions read against the derived ones, not a disagreement about the vacuum. Any claimed derivation should therefore supply a meV-type quantity without introducing a new arbitrary scale by hand.

One structural remark belongs here. The seesaw scale $m_\Lambda = M_{\mathrm{lock}}^2/M_{\mathrm{Pl,red}}$ is now *uniquely* the terminus's number: the retired TeV-inflation reading had, unknowingly, assigned the same combination to the launch (its $H_{\mathrm{inf}} = M_{\mathrm{lock}}^2/\sqrt{3}\,M_{\mathrm{Pl,red}}$ is this seesaw, up to $\sqrt{3}$), double-booking one scale at both ends of the descent. In the two-basin reading the bookkeeping is clean: the

combination $M_{\rm lock}^2/M_{\rm Pl,red}$ is the curvature scale of the middle basin and the density scale of the terminus — the lock episode briefly *visits* the terminus's vacuum energy before the crossing drains it, and dark energy is, in this reading, the undrained residue of that drainage (a flagged interpretation: it gives the open coefficient ζ_Λ a physical home — the drainage inefficiency — without yet computing it).

GG-6 provides an immediate dimensional candidate from the already-derived scale chain. The simplest combination of $M_{\rm lock}$ and the Planck scale with dimensions of energy is

$$m_\Lambda^{\rm geom} := \frac{M_{\rm lock}^2}{M_{\rm Pl,red}}. \tag{88}$$

The reduced mass, not the full one, is the right normalization here, and the choice is load-bearing rather than cosmetic: it moves the answer by $\sqrt{8\pi} = 5.01$. It is the reduced mass because this quantity is destined for the Friedmann equation and for $H_\Lambda = \sqrt{\rho_\Lambda/3M_{\rm Pl,red}^2}$, where $M_{\rm Pl,red}$ is what appears; equivalently it is the normalization in which the boundary Einstein–Hilbert action of §21 was read off. Using Eq. (2),

$$m_\Lambda^{\rm geom} \simeq 6 \times 10^{-3}\ {\rm eV} = 6\ {\rm meV}. \tag{89}$$

This is already in the correct observational corridor. The corresponding vacuum density scale is

$$\rho_\Lambda^{\rm geom} \sim \zeta_\Lambda \left(\frac{M_{\rm lock}^2}{M_{\rm Pl,red}}\right)^4, \tag{90}$$

with ζ_Λ the suppression factor carrying the exact zero-mode overlap and residual-ledger bookkeeping. Uncorrected, $(6.083\,{\rm meV})^4$ (6.09 meV at the rounded lock scale 3.85 TeV, the figure GG-A5-GDOS prints) overshoots the measured ρ_Λ by a factor of 56 — a factor of 2.73 in scale, set against the 10^{120} the seesaw has already removed. That residue is not left open: §29.8 selects $\zeta_\Lambda = \pi/LK = \pi/175 = 0.017952$ from the frozen ledger, with the mechanism stated there, and closes the scale at $\rho_\Lambda^{1/4} = 2.23$ meV against the 2.224 meV of Eq. (87) (Eq. (112)). What Eq. (90) contributes at this stage is the scale; the coefficient is the business of §29.8. No new scale beyond the inherited lock scale is required for either.

This "lock-scale seesaw" is one of the cleanest places where the cosmology sequel benefits from the particle derivation. Without the particle paper, the appearance of the meV scale would look like a numerological accident. With the particle paper, it is the first nontrivial infrared scale one can build from the already-fixed hierarchy. The same lock that held the launch at TeV now sets the floor of the terminus at meV — one geometry, both ends.

29.4 Why the equation of state is exactly $w = -1$: the transition is complete

If the late-time term is truly the residual zero-mode export of a selected vacuum functional, then it should be nearly constant on cosmological timescales. The stress tensor of a constant vacuum term is

$$T^\mu{}_\nu(\Lambda) = -\rho_\Lambda\, \delta^\mu{}_\nu, \tag{91}$$

which implies

$$p_\Lambda = -\rho_\Lambda, \qquad w_\Lambda = -1. \tag{92}$$

Therefore the most natural minimal prediction of the framework is not a rolling quintessence-like field with arbitrary $w(z)$ (Copeland *et al.*, 2006; Peebles & Ratra, 2003), but a vacuum-like term extremely close to -1. This is the signature of an attractor that has been *reached*: a rolling field is a descent still in motion, but $w = -1$ is the descent at rest. The vacuum-residual exactness says the transition is complete — the boundary has arrived at its terminus and stopped.

This is also where the anomaly logic of GG-6 becomes conceptually relevant. The earlier papers are built on a tightly constrained anomaly-free spectrum. Introducing an extra light rolling scalar solely to solve dark energy would be constitutionally expensive unless that scalar was already required elsewhere. The minimal route is therefore to treat dark energy as the frozen residual of a completed geometric transition rather than as a new light field still rolling today.

Honesty is required here, and it is now sharper than a plausibility argument. The *sign* of any departure is a theorem: a ghost-free, minimally coupled vacuum sector cannot cross below $w = -1$, so the framework forbids phantom behavior outright (Theorem 29.2 below). The *magnitude* question is now also a theorem, conditional on the constitution:

Theorem 29.1 (Exact vacuum equation of state at the terminus (T-Cosmos-DE′), conditional)**.** *Assume:* ***(I)*** *the constitutional rigidity of the bulk profile (GG-A2-GG6 (Arisaka, 2026, A2-GG6); the axiomatic e^{-2r} warp, whose dynamical stabilization is the named upstream open item);* ***(II)*** *the KK spectral gap of GG-A5-GDOS (Arisaka, 2026, A5-GDOS) (T9.1): no massless modulus beyond the gauge zero modes, lightest massive mode at $M_{\rm lock} \simeq 3.85$ TeV;* ***(III)*** *Lemma DE-1 of GG-A6-GRIP (Arisaka, 2026, A6-GRIP) (vacuum-flux freezing at the Geometric Attractor, itself a consequence of GFF-ST6 and the BI-6 closure). Then the dark-energy equation of state on the late-time branch satisfies $w(a) = -1$ identically, with any sourced deviation bounded by*

$$|1 + w| \;\le\; O(\beta^2)\left(\frac{H_0}{M_{\rm lock}}\right)^2 \;\approx\; 10^{-91}$$

for any residual matter–modulus coupling $\beta \le 1$.

Proof. By Lemma DE-1 the vacuum export is individually conserved at the attractor, $\mathrm{d}\rho_\Lambda/\mathrm{d}\ln a = 0$; the continuity equation $\mathrm{d}\rho_\Lambda/\mathrm{d}\ln a = -3(1+w)\rho_\Lambda$ then forces $1 + w = 0$ (Step 1). Cosmological sources displace the state off the critical point by $\delta\phi \sim \beta\,\rho_m/(m_{\rm min}^2 M_{\rm Pl})$; with $m_{\rm min} \ge M_{\rm lock}$ from (II), the induced deviation is $1 + w \simeq \dot\phi^2/\rho_\Lambda = O(\beta^2)(H_0/M_{\rm lock})^2$ (Step 2). Combined with the unconditional no-phantom bound $w \ge -1$ (Theorem 29.2), the value is pinned two-sidedly (Step 3). ∎

Corollary (two-sided kill switch). Any confirmed deviation of w from -1, in either direction and at any late-time redshift — including a converged evolving-dark-energy (w_0–w_a) detection — refutes the fixed-point reading outright; the framework cannot absorb $w \ne -1$. *Provenance chain:* GG-A2-GG6 (rigid profile) → GG-A5-GDOS (T9.1 spectral gap) → GG-A6-GRIP (GFF-ST6 → Lemma DE-1) → the present theorem: every link is a constitutional axiom or a numbered theorem of the cohort, and the single dynamical residue — the profile-stabilization proof — is named in

hypothesis (I). Codebook consequence: row D16 conservatively remains D*, its named remaining lemma now being that stabilization proof rather than the exactness itself (Part IX). The current empirical determination, $w = -1.028 \pm 0.032$ from the Planck 2018 combination with supernovae and BAO (Planck Collaboration, 2018), is consistent with this expectation at the 0.9σ level (a $+2.7\%$ fractional deviation); the entry appears as an independent prediction in the parameter-audit Dev plot (Figure 31, Table 41).

Theorem 29.2 (No phantom: the terminus equation of state (T-Cosmos-DE))**.** *Let the late-time vacuum sector consist of a finite collection of canonically normalized scalar moduli $\{\varphi_i\}$ (the Σ_r scale-modulus, Green–Schwarz, and GFF directions), minimally coupled to gravity on the* FRW *background, and let the spectrum be ghost-free — as the anomaly-free, unitary construction of GG-A2-GG6/GG-A3-BI6 guarantees. Then the dark-energy equation-of-state parameter obeys*

$$w \;\geq\; -1 \qquad \textit{at all times,}$$

with equality if and only if every $\dot\varphi_i = 0$, i.e. at a stationary point of $V =$ GFF — the GGA*. Consequently a secure, model-independent measurement of $w < -1$ (a true phantom crossing) falsifies the construction.*

Proof. For canonical scalars the energy density and pressure on the homogeneous background are $\rho = \sum_i \tfrac{1}{2}\dot\varphi_i^2 + V$ and $p = \sum_i \tfrac{1}{2}\dot\varphi_i^2 - V$, so

$$\rho + p = \sum_i \dot\varphi_i^{\,2} \;\geq\; 0.$$

Hence $1 + w = (\rho + p)/\rho = \sum_i \dot\varphi_i^2 / \rho \geq 0$, since $\rho > 0$. Equality requires $\dot\varphi_i = 0$ for every i; the scalar field equation $\ddot\varphi_i + 3H\dot\varphi_i + \partial_i V = 0$ then forces $\partial_i V = 0$, a stationary configuration of GFF, which is the GGA. A value $w < -1$ would require $\rho + p < 0$, i.e. a negative kinetic term — a ghost — which the ghost-free spectrum excludes. ∎

The theorem fixes the *sign*; the *approach* to -1 is equally controlled. Near the GGA, write $m_\varphi^2 = V''(\varphi_\star) \geq 0$. In the friction-dominated regime ($\tfrac{1}{2}\dot\varphi^2 \ll V$, i.e. $m_\varphi \ll H$) the slow-roll relation $3H\dot\varphi \simeq -V'$ with $H^2 \simeq V/3M_{\rm Pl,red}^2$ gives

$$1 + w \;\simeq\; \tfrac{2}{3}\,\epsilon_V, \qquad \epsilon_V := \frac{M_{\rm Pl,red}^2}{2}\left(\frac{V'}{V}\right)^2 \geq 0, \tag{93}$$

and the decaying mode $\varphi - \varphi_\star \propto e^{-(m_\varphi^2/3H_\Lambda)t}$ drives $\epsilon_V \to 0$, so $w \to -1$ exponentially as the attractor is approached. Three statements therefore stand together: the sign $w \geq -1$ is a theorem; the slow-roll form $1 + w = \tfrac{2}{3}\epsilon_V$ is exact given a canonical kinetic term; and the present magnitude $1 + w_0 = \tfrac{2}{3}\epsilon_V(\varphi_0)$ is exponentially small once the system is inside the GGA basin. What remains open is the single D* item that already fences the cosmological principle — a proof that the realized late-time state lies in that basin — and the numerical value $\epsilon_V(\varphi_0)$, which needs the GFF curvature at its floor. Neither dents the sign theorem: the empirical $w_0 = -1.028 \pm 0.032$ sits on the non-phantom

edge (the 0.9σ entry noted above), and the framework predicts the central value should migrate to $w_0 \gtrsim -1$ as the data sharpen.

29.5 From flatness to the late-time density fractions

Once radiation becomes negligible, the first Friedmann equation may be written in the present epoch as

$$1 = \Omega_m + \Omega_\Lambda + \Omega_k, \qquad \Omega_m := \frac{\rho_m}{\rho_c}, \qquad \Omega_\Lambda := \frac{\rho_\Lambda}{\rho_c}, \tag{94}$$

with

$$\rho_c := \frac{3H_0^2}{8\pi G_N}. \tag{95}$$

If inflation drives the universe close to spatial flatness, then $\Omega_k \approx 0$, so

$$1 \simeq \Omega_m + \Omega_\Lambda. \tag{96}$$

Equation (96) is standard. The new question is whether GG-6 supplies a minimal late-time partition of this unit budget.

The working proposal sharpened in the cosmology notes of the project is a *π-locked late-time inventory*. The logic is the following.

1. The only continuous geometric quantity that survives as a primitive measure of the compact angle is the circle measure, and under the late-time lock slice the relevant minimal angular budget is the half-circle π rather than an arbitrary fraction.
2. MSMF forbids an arbitrary decimal occupancy of this compact budget.
3. The smallest nontrivial matter share compatible with one minimal occupancy of the locked slice is therefore taken to be $1/\pi$.

This yields the late-time matter fraction

$$\Omega_m = \frac{1}{\pi} \simeq 0.3183. \tag{97}$$

Flatness then forces the complementary vacuum fraction

$$\Omega_\Lambda = 1 - \frac{1}{\pi} \simeq 0.6817. \tag{98}$$

The baryonic share carries one further compact-winding suppression in the minimal OGN projection, giving

$$\Omega_b = \frac{1}{2\pi^2} \simeq 0.05066. \tag{99}$$

Therefore the dark-matter fraction in the same minimal inventory is

$$\Omega_{\rm DM} = \Omega_m - \Omega_b = \frac{1}{\pi} - \frac{1}{2\pi^2} \simeq 0.26765. \tag{100}$$

Taken together,

$$\Omega_b + \Omega_{\rm DM} + \Omega_\Lambda = \frac{1}{2\pi^2} + \left(\frac{1}{\pi} - \frac{1}{2\pi^2}\right) + \left(1 - \frac{1}{\pi}\right) = 1. \tag{101}$$

Equation (101) is mathematically trivial once the shares are chosen, but conceptually it is not trivial. It says that the present cosmic inventory is being treated as a geometric late-time partition of one unit budget, not as three unrelated fit parameters.

Here is the fence, stated once and plainly. The flatness sum rule is theorem-level once the standard FRW background is accepted. The precise π-locked partition is *derived-conditional*: it closes only up to Conjecture VIII, whose printed Step-4 cancellation is now *resolved by identification*: it sought an epoch-independent closure that provably does not exist, and the partition closes instead through the measure dictionary plus the declared epoch input (the residual uniformity lemma exact to one part in 10^{15}). The $1/\pi$ value is therefore carried as a structural ansatz of remarkable empirical precision — form derived, epoch declared — not as an unconditional theorem, and Ω_m, Ω_Λ, Ω_b are listed in the codebook as D* under that single named input. This is the one topological step on which the terminus depends, and it is fenced here so the rest of the descent can stand on its own.

29.6 The cosmic coincidence as an attractor-slice

With the late-time fractions in hand, the coincidence question becomes very concrete. In standard ΛCDM the present near-comparability of Ω_m and Ω_Λ looks accidental because matter dilutes as a^{-3} while vacuum energy does not dilute at all. Unless one introduces anthropic reasoning, it seems surprising that we happen to live close to the crossover epoch.

The GG-6/GDOS reinterpretation is different. If the present fractions are the geometric shares of a specific late-time attractor slice, then the current near-balance is not a coincidence inserted by random initial conditions. It is the natural signature that the system has already relaxed to the slice selected by the same long-term vacuum ledger — it is what the terminus of the descent looks like from inside.

One can quantify the crossover redshift without any additional freedom. The matter density evolves as

$$\rho_m(a) = \rho_{m0} a^{-3}, \tag{102}$$

while the vacuum density is constant. Equality occurs when

$$\rho_{m0} a_{\Lambda m}^{-3} = \rho_{\Lambda 0}, \tag{103}$$

so

$$a_{\Lambda m} = \left(\frac{\Omega_m}{\Omega_\Lambda}\right)^{1/3}. \tag{104}$$

Using Eqs. (97) and (98),

$$\frac{\Omega_\Lambda}{\Omega_m} = \pi - 1, \tag{105}$$

so

$$a_{\Lambda m} = (\pi - 1)^{-1/3}, \qquad z_{\Lambda m} = a_{\Lambda m}^{-1} - 1 = (\pi - 1)^{1/3} - 1 \simeq 0.289. \tag{106}$$

The standard coincidence question then changes form. Instead of asking "why do we happen to live near the equality epoch?", the framework asks "why did the late-time attractor slice selected by the locked geometry correspond to the present epoch?" That is still a deep question, but it is no longer the question of an apparently random numerical near-equality.

Theorem 29.3 (Minimal late-time π-locked attractor chart (T-Cosmos-5))**.** *Assume a spatially flat late-time geometric-attractor branch, a closed TGG ledger, the local seam admissibility of CCI, and the minimal compact-occupancy rule of the late GFF slice. Then the minimal late-time inventory is*

$$\Omega_m = \frac{1}{\pi}, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi}, \qquad \Omega_b = \frac{1}{2\pi^2}, \qquad \Omega_{\rm DM} = \frac{1}{\pi} - \frac{1}{2\pi^2}.$$

Proof. Flatness of the late attractor gives

$$1 \simeq \Omega_m + \Omega_\Lambda.$$

The framework then treats the matter share as the smallest nontrivial compact occupancy of the late branch, $x = 1/\pi$, rather than as an arbitrary decimal parameter. Hence $\Omega_m = x$ and $\Omega_\Lambda = 1 - x$. The baryonic share carries one further compact-winding suppression of $1/(2\pi)$ relative to total matter, so

$$\Omega_b = \frac{1}{2\pi}\Omega_m = \frac{1}{2\pi^2}.$$

The remaining matter share is then dark matter:

$$\Omega_{\rm DM} = \Omega_m - \Omega_b.$$

This is not a proof from classical FRW alone; it is a proof inside the late minimal-branch rules of the framework. ■

29.7 From the residual vacuum density to the present Hubble scale

It is useful to push the meV-scale discussion one step further and translate the residual vacuum density into a late-time Hubble scale. If

$$\rho_\Lambda \sim \zeta_\Lambda \left(\frac{M_{\rm lock}^2}{M_{\rm Pl,red}} \right)^4, \tag{107}$$

then the asymptotic de Sitter expansion rate is

$$H_\Lambda = \sqrt{\frac{\rho_\Lambda}{3M_{\rm Pl,red}^2}} \sim \sqrt{\frac{\zeta_\Lambda}{3}} \frac{M_{\rm lock}^4}{M_{\rm Pl,red}^3}. \tag{108}$$

Using the numerical lock scale from Eq. (2), one finds the characteristic magnitude

$$H_\Lambda \sim 10^{-33}\ {\rm eV} \qquad \text{for} \qquad \zeta_\Lambda = \mathcal{O}(10^{-2}), \tag{109}$$

which is precisely the order one expects for the present cosmic acceleration scale. The coefficient is not left at an order of magnitude: Section 29.8 selects $\zeta_\Lambda = \pi/LK = \pi/175$ from the ledger and closes the scale to the percent level (Eq. (112)). This is another place where the theory gains explanatory leverage by linking a cosmological infrared quantity directly to the already-derived particle hierarchy scale. Equation (108) is particularly satisfying because it compresses the whole late-time problem into one line: once the particle paper has fixed $M_{\rm lock}$, the vacuum seesaw automatically knows roughly how slow the present acceleration should be.

Proposition 29.4 (Late vacuum Hubble scaling (T-Cosmos-6))**.** *If the residual vacuum density is*

$$\rho_\Lambda \sim \zeta_\Lambda \left(\frac{M_{\rm lock}^2}{M_{\rm Pl,red}}\right)^4,$$

then the present Hubble scale is

$$H_0 \sim \left(\frac{M_{\rm lock}^2}{M_{\rm Pl,red}}\right)^2 \frac{1}{M_{\rm Pl,red}} \sqrt{\frac{\zeta_\Lambda}{3\Omega_\Lambda}},$$

which lies naturally in the observed late-time corridor for mild ζ_Λ.

29.8 The derivation of H_0: the two-route anchor of the present expansion rate

The scaling proposition (T-Cosmos-6) can be read as far more than a dimensional check: it says the present Hubble constant is not an independent input of the cosmology but a derived combination of quantities the particle papers already fixed. In fact the present rate is forced *twice over*, by two routes that share no ingredient except the geometry itself, and the order in which they are stated is the point. *Route I* is geometry alone: the theorem-grade lock scale, the vacuum-energy seesaw, and one codebook-selected ledger monomial — no telescope, no microwave sky, nothing measured off any astronomical object. *Route II* is the precision confirmation: the sky's two h-independent physical densities, divided by the derived π-locked fractions, land on the same rate. Their agreement at $H_0 \approx 67\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ is the robust content of this subsection; the Hubble companion GG-C2-Hubble (Arisaka, 2026, C2-Hubble) carries both routes at full depth, together with the selection's defense against a hostile reader.

Route I — the geometric determination: the vacuum-energy seesaw. The dark-energy sector delivers H_0 first, with no astronomical input anywhere in the chain. Collecting the seesaw density (90) and the flat late-time budget (96) into the first Friedmann equation $3M_{\rm Pl,red}^2 H_0^2 = \rho_\Lambda/\Omega_\Lambda$ gives the closed form

$$\boxed{H_0 = \frac{1}{M_{\rm Pl,red}} \left(\frac{M_{\rm lock}^2}{M_{\rm Pl,red}}\right)^2 \sqrt{\frac{\zeta_\Lambda}{3\,\Omega_\Lambda}}} \tag{110}$$

in which every symbol is inherited: $M_{\rm lock}$ from the lock (GG-A4-Euclid/GG-A2-GG6), $M_{\rm Pl,red}$ from the warp overlap (GG-A5-GDOS), $\Omega_\Lambda = 1 - 1/\pi$ from the late-time inventory (T-Cosmos-5), and ζ_Λ a single residual dimensionless coefficient. Numerically it reduces to $H_0 = 498.3\,\sqrt{\zeta_\Lambda}\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$.

The chain is stiff at every rung but the last. The lock scale is *not* an open factor — $M_{\rm lock} = \frac{1}{2}E_{\rm Pl}\,e^{-K} \approx 3.85$ TeV is a theorem of the scale chain (GG-A2-GG6 §5 with T-GG6-6, where the quantitative chain is owned; GG-A5-GDOS T8.7 restates it as an imported result in the warp setting) — and the skeptic's fear of a percent-level threshold correction (dangerous, since $H_0 \propto M_{\rm lock}^4$) is bounded in the laboratory, without cosmology: the same lock scale predicts the electroweak vacuum expectation value, $v = M_{\rm lock}\,(2/5)^3 = 246.33$ GeV against the measured 246.22 — a G_F pin at $+0.045\%$, which $H_0 \propto M_{\rm lock}^4$ maps to $\pm 0.18\%$ on the rate. Nor is ζ_Λ free: MSMF forces it to be a bounded-cost ledger monomial, and OGN forces every constituent of that monomial to be a named object on the stratum where ζ_Λ lives — the Wilson-line (Hosotani) zero mode on S^1_θ, on which exactly three objects carry names (the paragraph after the table says which). Enumerating every monomial in those three, with costs under GG-A9-SM's closure rule for K ($K = LJ$ is not an independent atom, so a K-factor costs two), gives the stratum cluster of Table 9.

Table 9: Every monomial in the three named objects of the Wilson-line stratum — the seam factor 2, the L-holonomy angle $2\pi/L$, and the dual-lock count K — at cost ≤ 5 whose implied Hubble constant lies in 40–100 $\mathrm{km\,s^{-1}\,Mpc^{-1}}$: fourteen monomials, five admitted by the residue rule and nine excluded by it, each exclusion with its reason. Costs are stated under GG-A9-SM's closure rule, under which K is not an independent atom and a K-factor costs two. The lower two blocks are the monomials the residue rule stated below the table excludes, printed with the clause each fails so the exclusion can be checked rather than taken. Of the monomials the rule admits, exactly one at cost ≤ 4 lands in the observed decade; the first admitted alternative costs one unit more and lands at 74.8. (An earlier version of this table was drawn from the unrestricted alphabet $\{2, N, L, J, \pi, e, K\}$, on which forty-nine monomials at cost ≤ 4 land between 60 and 80 — including $e/L^3 \to 73.5$, the ladder's value — and printed four of them; that table did not show what it was said to show, and is replaced by the stratum-restricted one, which does.)

cost	expression	ζ_Λ	$H_0 = 498.3\sqrt{\zeta_\Lambda}$
admitted by the residue rule			
2	$1/K$	0.0286	84.2
3	$1/(2K)$	0.0143	59.6
3	$(2\pi/L)/K$	0.0359	94.4
4	$\boldsymbol{(2\pi/L)/(2K) = \pi/(LK)}$	**0.017952**	**66.8**
5	$(2\pi/L)^2/(2K)$	0.02256	74.8
excluded: the angle inverted (a suppression where the residue takes an amplitude)			
3	$1/\big((2\pi/L)K\big)$	0.02274	75.1
4	$1/\big(2(2\pi/L)K\big)$	0.01137	53.1
4	$1/\big((2\pi/L)^2K\big)$	0.018093	67.0
5	$1/\big(2(2\pi/L)^2K\big)$	0.00905	47.4
5	$1/\big((2\pi/L)^3K\big)$	0.01440	59.8
5	$2/\big((2\pi/L)^2K\big)$	0.03619	94.8
excluded: the seam factor squared (the orbifold halves once)			
4	$1/(4K)$	0.00714	42.1
5	$(2\pi/L)/(4K)$	0.00898	47.2
excluded: no closure suppression (K absent)			
5	$1/32$	0.03125	88.1

The residue rule, stated before the table is read. ζ_Λ is the residue of the Wilson-line zero mode, and a residue is built from its stratum's objects in the roles those objects have there: the holonomy angle $2\pi/L$ enters as an *amplitude*, at a non-negative power, because it is the extent of the circle the mode winds and not a suppression; the seam factor 2 enters only as the orbifold halving, at power 0 or -1, the same $\frac{1}{2}$ that appears in $M_{\rm lock} = \frac{1}{2}E_{\rm Pl}e^{-K}$ and in $\Omega_b = 1/2\pi^2$; and K enters only as the closure suppression, at a negative power, as it does in e^{-K} and in $\bar{\theta}$. The rule is fixed by what the objects *are* on that stratum, not by which candidate it leaves standing, and it is stated here before the count that follows. It excludes four monomials at cost ≤ 4 in the window — $1/((2\pi/L)K)$ at cost 3, $1/(2(2\pi/L)K)$ and $1/((2\pi/L)^2K)$ at cost 4, all three inverting the angle, and $1/(4K)$ at cost 4, which halves twice where the orbifold halves once — and every excluded monomial is printed in the table with its reason, so a reader can judge the exclusion rather than accept it. The rule fixes which symbols may appear, not where the mode sits along the throat: GG-A9-SM computes that mode's radial profile on the canonical metric and returns an exactly flat mode whose norm is carried by the ultraviolet end (its PS-SM-22), and none of the three objects named here is an overlap integral or a normalization constant, so ζ_Λ does not move with the profile. What the profile does constrain is the completed computation (GG-A5-GDOS O6), which must now supply the $1/K$ from the winding count of the dual lock and not from the length of the interval. One reading of the residue is registered here because it exists in the cohort's open-item ledger, and it is registered as what it is: a conjecture this paper does not hold. The *twist-rate reading*, C-S3-1, reads $\zeta_\Lambda = \frac{1}{2}\langle A_{B,r}\rangle = \pi/(LK)$ as half the mean of the baryonic connection's radial component over the interval — an unweighted mean, which returns the $1/K$ from the length of the interval, exactly the normalization the rider above excludes, and which relocates the residue from the circle to the interval line. It is not adopted. It is carried with the two conditions it would have to meet and the falsifier it has: the interval-line period must be $\int_0^K A_{B,r}\,dr = 2\pi/L$ under the Dirichlet termination of GG-A3-BI6's clause (C9), which GG-A3-BI6 registers as its Open Problem O16 (Arisaka, 2026, A3-BI6); and the residue computed in the warped measure must return the same $1/K$ from the winding count and not from an unweighted normalization, which is this rider. Rows D5 and D6 do not move on that period in either direction; they move on the residue computation of GG-A5-GDOS O6 against the window below, and on nothing else.

Read that way, the cluster is a two-sided result: geometry has exactly one admitted reading at cost ≤ 4 in the observed decade, 66.8 $\mathrm{km\,s^{-1}\,Mpc^{-1}}$, with zero observational input of any kind; the first admitted alternative is dearer by one unit and lands at 74.8, 1.8% above the distance ladder's 73.5; and no admitted monomial at cost ≤ 5 returns the ladder's value ($\zeta_\Lambda = 0.0218$). That is weaker than "the geometry refuses the high reading" — a sentence an earlier version printed on the strength of a four-row table drawn from an unrestricted alphabet, which does not support it — and it is true. The enumeration is released as `zeta_table_v2.py`, complete in the stated window.

The selection of ζ_Λ: by the codebook, not by the data. Which candidate the ledger prefers is decided neither by cost alone — on the unrestricted alphabet, cost favors the orphan $1/J^2$, which the stratum rule excludes before the table is drawn — nor by the data. It is decided by GG-A9-SM's own selection machinery, applied without modification: every constituent of an admissible monomial

must have a *named structural role on the parameter's boundary stratum*, the discipline GG-A9-SM has already enforced against its own arithmetic (for $\sin^2\theta_{12}^{\rm PMNS}$ the bounded-cost search returned the *cheaper* orphan $1/\pi$, and GG-A9-SM rejected it as geometrically isolated). The stratum is fixed by what ζ_Λ *is*: the residue of the Wilson-line (Hosotani) zero mode on S^1_θ — the same kind of object as the strong-CP residue $\bar\theta \sim \sin(2\pi/L)\,2^{-J}e^{-K}$ of GG-A9-SM, a geometric closure-index suppression, not a radiative integral. On that stratum exactly three objects carry names: the L-holonomy angle $2\pi/L$ ($\bar\theta$'s own sine argument), the orbifold/dyadic seam factor 2 (the $\frac{1}{2}$ of $M_{\rm lock} = \frac{1}{2}E_{\rm Pl}e^{-K}$ and of $\Omega_b = 1/2\pi^2$), and the dual-lock count K (the e^{-K} suppression; the Kaxion's cap-loading ladder $K+\Sigma$). The three rejections are then one rule applied three times — a constituent with no bare name on the stratum: N nowhere (its triplicity roles live on matter strata); J only through the composite $K = LJ$, so the bare J^{-2} is an orphan, the same class as the rejected $1/\pi$ of θ_{12} and likewise cheaper than the winner; and π only through the composite $2\pi/L$, so the bare π downstairs of $2/(K\pi)$ is an orphan too. Exactly one candidate survives, and it decomposes entirely into the named objects:

$$\zeta_\Lambda \;=\; \frac{\pi}{LK} \;=\; \frac{1}{2}\times\frac{2\pi}{L}\times\frac{1}{K} \;=\; \frac{\pi}{175} \;=\; 0.017952, \tag{111}$$

the undrained vacuum residue of the lock episode: one orbifold-halved L-holonomy defect, shared over the K winding cells of the dual lock — every factor already load-bearing elsewhere in the cohort.

With ζ_Λ selected, the vacuum scale is no longer an order-of-magnitude statement. Evaluating (107) at the selected value gives

$$\rho_\Lambda^{1/4} \;=\; \zeta_\Lambda^{1/4}\,\frac{M_{\rm lock}^2}{M_{\rm Pl,red}} \;=\; 0.3660\times 6.083 \;=\; 2.23\ \text{meV}, \tag{112}$$

against 2.224 meV read off this paper's own derived cosmology ($H_0^{(0)} = 66.6$, the two-road reading of Route I's 66.8; $\Omega_\Lambda = 1 - 1/\pi$) and 2.240 meV off Planck 2018 — agreement at 0.1% and 0.6% in scale, 0.4% and 2.4% in density. The 10^{120} discrepancy is therefore not reduced to a factor of a few; on the selected coefficient it is closed to the percent level, and it is closed with nothing fitted — $M_{\rm lock}$ from the lock, K and L from the ledger, π from the geometry. The grade travels with the number: this is Derived (cond.) on the same stratum-assignment lemma as Route I, and it stands or falls with the same two-sided window, $\zeta_\Lambda \in [0.01784,\ 0.01806]$ — the $\pm 0.6\%$ the rate's own error implies, narrower than the $\pm 1\%$ the companion's registry first printed, because only the narrower window excludes the two nearest competitors.

With the selection made, Route I closes:

$$\boxed{H_0^{\rm geom} = 498.3\,\sqrt{\pi/175} \;=\; 66.8 \pm 0.2\ \text{(structural)}\ \text{km}\,\text{s}^{-1}\,\text{Mpc}^{-1}} \tag{113}$$

where the structural error carries the laboratory pin on the lock scale and the conditional status of the selection lemma. The grade is stated at once: *Derived (cond., MSMF/OGN)*, conditional on one named lemma — the *stratum assignment* (ζ_Λ is the Wilson-line zero-mode residue), the physical reading this section has always carried — with Derived (theorem) waiting on the Hosotani

computation itself (GG-A5-GDOS O6). And the selection is cashed by one pre-registered number: the completed computation must return $\zeta_\Lambda = \pi/175$ to within $\pm 0.6\%$, two-sided. A value outside $[0.01784, 0.01806]$ — the $\pm 0.6\%$ that the rate's own ± 0.2 km s^{-1} Mpc^{-1} implies, which excludes both the rejected $2/(K\pi) = 0.018189$ and the rule-excluded $1/((2\pi/L)^2 K) = 0.018093$ — refutes the selection lemma and demotes Route I from a point back to its bracket. The target is registered in the cohort falsifier ledger and printed as falsifier F-Hubble-11 of GG-C2-Hubble (Arisaka, 2026, C2-Hubble); it is printed before the lemma and cannot be renegotiated after it. This paper's window governs: GG-A5-GDOS, which owns the computation, carries this window in its Open Problem O6 and keeps no second one (Arisaka, 2026, A5-GDOS). (A field theorist's instinct to read ζ_Λ as a Coleman–Weinberg loop factor $N/16\pi^2 = 0.019$ is a numerical impostor: in the cost metric the loop form is expensive (cost seven), and GG-A9-SM already settled this category once, showing the Higgs quartic $\lambda_H = 2^{-N} = 1/8$ is "curvature stiffness of the selected attractor branch, not a one-loop residue." The quantitative form of the same point: a loop of the Kaluza–Klein tower does not produce the residue at all — it is about 10^{59} too large in scale and would need a coefficient of order 10^{-61} to land in the window, which is not a loop factor of anything.)

The closed form: the rate from $\{c, G_N, \hbar\}$ and the ledger alone. The whole of Route I now collapses into a single analytic expression, and it is worth printing, because the claim of this subsection — geometry first, the sky second — is nowhere more visible than in one equation whose every symbol can be pointed at. Substituting the lock theorem (in reduced units, $M_{\rm lock} = \sqrt{2\pi}\, M_{\rm Pl,red}\, e^{-K}$), the seesaw, and $\zeta_\Lambda = \pi/(LK)$ into the present-epoch Friedmann equation at $\Omega_\Lambda = 1 - 1/\pi$, every power of the Planck mass cancels except a single overall rate, the transcendentals collect into $2\pi^5/(\pi - 1)$, and Eq. (113) becomes

$$\boxed{H_0 \;=\; e^{-4K}\sqrt{\frac{2\,\pi^5\,c^5}{3\,L\,K\,(\pi-1)\,\hbar\,G_N}} \;=\; 66.7\ \text{km}\,\text{s}^{-1}\,\text{Mpc}^{-1},} \tag{114}$$

which a CODATA evaluation returns as 66.74. The provenance of every symbol fits in one sentence: c, G_N, and $\hbar$ carry the dimensions ($\sqrt{c^5/\hbar G_N}$ is the inverse Planck time — the GU-5 bridge and nothing else); the ledger integer $K = 35$ appears twice, as the depth of the descent (e^{-4K}) and in the selected monomial; $L = 5$ appears once in the monomial — and a second time in the laboratory, through the vev pin $v = M_{\rm lock}(2/5)^3$; π enters three ways (Einstein's 8π, the monomial's numerator, the inventory's $1 - 1/\pi$); e is the scale chain's exponential; and the structural rationals $\frac{1}{2}$ and 3 are the lock theorem's and Friedmann's bookkeeping. No astronomical quantity of any kind appears: the expansion rate of the universe, to the precision quoted, is a statement about three laboratory constants and two small integers. Two celebrated hierarchies fall out as corollaries. The age of the universe in Planck units is $H_0 t_{\rm Pl} = 0.738\, e^{-140} = 10^{-60.9}$ — the sixty-one orders between the Planck rate and today's expansion are $4K = 140$ e-folds of the derived descent, dressed by a geometric factor of order unity, so "why is the universe so old?" receives a one-integer answer. And $\rho_\Lambda/\rho_{\rm Pl,red} = \zeta_\Lambda (2\pi)^4 e^{-8K} = 10^{-120.2}$ — the famous one hundred and twenty orders of the cosmological-constant problem are $8K$ e-folds (121.6 orders) corrected by 1.4 orders of pure geometry: two powers of one derived hierarchy, not a cancellation. One honesty note travels with the box:

of the OGN-5 integers, M, N, and J do not appear in the final expression — they act upstream, selecting the geometry that makes the rungs theorems, with J surfacing in $z_* = e^7$ and the rejected $1/J^2$ — and the equation is exactly as conditional as its rungs (D^* inventory; selection lemma cashed by F-Hubble-11). The line adds exposure, not strength: a reader with a calculator and five constants can check it in under a minute. The fuller treatment — assembly, corollaries, and the audit of what a critic should aim at — is carried by the Hubble companion (§24.3 of Arisaka, 2026, C2-Hubble).

Route II — the precision confirmation: the π-locked inventory reads the same rate off the microwave sky. This route uses no new dynamics — only the π-locked inventory already fixed above ($\Omega_b = 1/2\pi^2$, $\Omega_m = 1/\pi$; Eqs. (97)–(99)) together with what the cosmic microwave background measures *model-independently*. The acoustic peaks do not measure H_0; they measure the two *physical* densities

$$\omega_b \equiv \Omega_b h^2 = 0.02237 \pm 0.00015, \qquad \omega_m \equiv \Omega_m h^2 = 0.1430 \pm 0.0011,$$

the first from the baryon-loading ratio of the second-to-first acoustic peak (and independently from the BBN deuterium abundance), the second from the peak-height envelope and the matter–radiation equality redshift.[2] Neither carries h. Since $\omega_i = \Omega_i h^2$, a GG-6 fraction and a measured physical density divide to give the rate with no freedom left,

$$h = \sqrt{\frac{\omega_i}{\Omega_i}} \implies \begin{cases} h = \sqrt{0.02237/0.05066} = 0.665 & \text{(baryon anchor, } H_0 = 66.5\text{)}, \\ h = \sqrt{0.1430/0.31831} = 0.670 & \text{(matter anchor, } H_0 = 67.0\text{)}, \end{cases} \tag{115}$$

two anchors on different observables that agree — and a third anchor removes the microwave sky from the confirmation entirely: the primordial deuterium abundance with the LUNA reaction rates gives $\omega_b^{\rm BBN} = 0.02233 \pm 0.00036$ with no CMB anywhere in the chain, returning $H_0^{\rm BBN} = 66.4 \pm 0.5$ — which is the Ω_b row read as a rate with BBN's ω_b in place of the CMB's, not a third road, and it should be read beside the other BBN-anchored, ladder-free inference in print, DESI's baryon-acoustic 68.51 ± 0.58 (§42), 2.8σ above it: the two share the BBN ω_b and differ in whether the sound horizon enters, and the framework predicts the first and not the second. Their combination is the precision reading of the derived rate,

$$\boxed{H_0^{(0)}\,\text{(density roads)} \;=\; 66.6 \pm 0.2,\ \text{on Route I's } 66.8 \pm 0.2 \;\approx\; 67\ \text{km s}^{-1}\,\text{Mpc}^{-1}\,.} \tag{116}$$

This confirmation is itself parameter-free — an inversion, not a fit. Standard ΛCDM *fits* Ω_m and Ω_b and reads off $H_0 = 67.4$; GG-6 instead *predicts* $\Omega_m = 1/\pi$ and $\Omega_b = 1/2\pi^2$, so H_0 is forced — and it lands on the independently derived Route-I value. A full Boltzmann solution completes the confrontation: fed GG-6's exact fractions at the carried $h = 0.6676$, the code (CLASS) reproduces the observed temperature spectrum — peak positions within 0.6% ($\ell = 221/539/816$ against the best

[2] ω_m is the total matter density including the neutrino-mass contribution (Planck TT,TE,EE+lowE+lensing); the baryon-plus-CDM sum alone is 0.1424.

fit's 220/536/813), the second-to-first peak-height ratio to 0.8%, and the diffusion-damping scale to 0.4%; the complete audit is Table 23. The π-locked inventory does not merely imply $H_0 \approx 67$ arithmetically; it reproduces the microwave sky at that value to the percent level, the acoustic angle 0.3% small ($\chi^2 = 187$ against the six-parameter best fit's 84; not a fit). The residual $\sim$1% below Planck's 67.4 is exactly GG-6's exact π-values sitting 1–3% above Planck's best-fit fractions, not a new tension.

The concordance, and the honest grade. The two routes meet: 66.8 ± 0.2 (structural; geometry alone) against 66.6 ± 0.2 (measured densities over derived fractions) — 0.25%, or 0.6σ of their combined errors — with the CMB-free BBN anchor at 66.4 ± 0.5 beside them. One route runs through TeV-scale physics and the vacuum, the other through the plasma of the young universe; they share no ingredient except the geometry itself, and they meet inside each other's error bars. This paper and the Hubble companion carry the structural 66.8 ± 0.2 as the rate; the two-road 66.6 ± 0.2 is the reading against the measured densities, and the 0.25% between the two is inside every comparison either makes. The grades follow the GG-A9-SM discipline: the inventory fractions $\Omega_m = 1/\pi$ and $\Omega_b = 1/2\pi^2$ are D* (closed up to Conjecture VIII and Theorem 74.1); the seesaw coefficient ζ_Λ is Derived (cond., MSMF/OGN) under the stratum-assignment lemma, one computation short of a theorem (GG-A5-GDOS O6, its target pre-registered above); the lock scale is a theorem, pinned in the laboratory through the electroweak vev. Honest bookkeeping records the history: an earlier reading of this section took the cost-two monomial $1/J^2$ as decisive and quoted $H_0^{(0)} = 71.3$; under GG-A9-SM's own admissibility rules the bare J^{-2} is an orphan on the Wilson-line stratum, and the microwave spectrum independently rejects it — GG-6's own fractions fit the acoustic peaks at ≈ 67, not 71. The early-versus-late *split* of this one background rate — the Hubble tension itself — is taken up in Section 42, where the derived background ≈ 67 is confronted with the local distance-ladder reading ≈ 73.5.

29.9 The de Sitter future and the entropy of the terminus

The long-term future is sharpened once $w \approx -1$ and $\Omega_\Lambda \approx 1 - 1/\pi$ are adopted. At late times the matter density continues to dilute while the vacuum term remains constant, so the Friedmann equation approaches

$$H^2 \to H_0^2 \Omega_\Lambda, \tag{117}$$

with asymptotic solution

$$a(t) \sim \exp\!\big(H_0 \sqrt{\Omega_\Lambda}\, t\big). \tag{118}$$

Therefore the minimal late-time fate is an asymptotic de Sitter-like expansion, not a recollapse and not a Big Rip. The reason a Big Rip is absent is simply that the equation-of-state parameter is not phantom-like:$w = -1$ rather than $w < -1$. The reason recollapse is absent is that the universe is already near flat and the vacuum term is positive. In plain English: the universe keeps expanding, galaxies outside the local bound structures disappear beyond the horizon, matter becomes increasingly dilute, and the future approaches a cold vacuum-dominated de Sitter phase. In the present framework this is not an independent assumption about the future. It is the direct

continuation of the same residual vacuum ledger that already explains late-time acceleration — the descent does not merely slow toward its terminus, it settles there forever.

The future event horizon of the minimal late-time solution is not merely a curiosity. It is the cleanest geometric signal that the universe is approaching a vacuum-dominated attractor. From Eq. (39), the asymptotic event-horizon radius is

$$R_e = \frac{c}{H_\Lambda}. \tag{119}$$

Because H_Λ is tiny, R_e is enormous on ordinary scales, but it is finite. This means that the future observable universe of a comoving observer does not grow without bound. There is a permanent causal screen in the late-time limit.

Associated with this horizon is the usual de Sitter entropy (Gibbons & Hawking, 1977),

$$S_{\rm dS} = \frac{\pi M_{\rm Pl,red}^2}{H_\Lambda^2} \quad \text{(up to the usual reduced vs unreduced Planck normalization convention).} \tag{120}$$

Within the present framework this entropy should be read as the asymptotic coarse-grained capacity of the late-time vacuum attractor. That is exactly the kind of statement for which the GFF ledger was invented. A late-time de Sitter phase is not only an equation-of-state statement. It is a boundary statement about the finite information capacity of the observable future patch — the entropy of the GGA itself.

This is also why the minimal dark-energy module is so tightly linked to the black-hole and information discussions later in the paper. A cosmology with a future event horizon inevitably raises questions about accessible information, coarse-grained entropy, and the difference between the full closed bulk and the finite records of a boundary observer. The late-time universe is therefore a gentle preview of the more extreme OGGEP issues that will return in the black-hole part.

29.10 Dark energy on the two fibers: the circle's share, the interval's scale

The terminus, too, reads cleanly on the campaign's anatomy (GG-W1-Symmetry; Arisaka, 2026, P5-CPT; Arisaka, 2026, C1-Kaxion), and the reading resolves a peculiarity of this Part that the derivation produced but did not name: dark energy's two numbers come from two different places. On the anatomy, they come from the two fibers — one each.

The share is the circle's. In the fiber-measure resolution of the π-locked inventory (§L.10), matter's share is the volume of the Hopf circle inside the electroweak S^3, $\Omega_m = \mathrm{Vol}(S^1)/\mathrm{Vol}(S^3) = 2\pi/2\pi^2 = 1/\pi$, and dark energy is *the uncovered remainder of the holonomy volume*, $\Omega_\Lambda = 1 - 1/\pi$. The largest single entry in the present energy budget is, on this reading, a statement about the compact fiber: the part of the holonomy volume that the matter fiber does not cover. The share is written in π because it is a ratio of circle-volumes; it is epoch-frozen because fiber-volumes do not dilute; and the competing candidates die by geometry — there is no fiber of volume $6\pi^2$ — rather than by fit. The coincidence redshift inherits the same character: $z_{\Lambda m} = (\pi - 1)^{1/3} - 1 \simeq 0.289$ is a pure circle number, read on the interval's clock — the epoch at which the diluting share crosses the

frozen remainder is fixed by π alone, because both shares are the circle's and only the dilution is the interval's.

The scale is the interval's. The lock-scale seesaw, written in e-fold form, is an exact identity of the published chain:

$$\frac{m_\Lambda}{M_{\rm Pl,red}} \;=\; \left(\tfrac{1}{2}\sqrt{8\pi}\right)^2 e^{-2K} \;=\; 6.28\, e^{-70} \;=\; 2.50\times 10^{-30}, \qquad m_\Lambda = 6.1\ \text{meV}, \tag{121}$$

— *twice the lock depth.* The meV is the descent's deepest export: the seesaw squares the one hierarchy the interval ever wrote (e^{-K}, the throat) and hands the result to the vacuum. The residual factor against observation ($\rho_\Lambda^{1/4} = 2.24$ meV) is $\zeta_\Lambda^{1/4} = 0.37$, the drainage inefficiency whose flagged home was recorded above. That factor is the circle's: it is the Wilson-line zero-mode residue selected in §29.8, and what the interval supplies here is the scale it trims. And the most notorious number in physics acquires an address: the cosmological-constant hierarchy is $-\ln(\rho_\Lambda/M_{\rm Pl,red}^4) = 276.6 = 8K - 3.4$ (within 1.2% of $8K = 280$), and the terminus's de Sitter entropy is the same number inverted, $\ln S_{\rm dS} = 278.9 = 8K - 1.1$ (within 0.4%). Honesty requires the immediate caveat, at the same register as the $4K$ compression of the age ledger: these are the age ledger *doubled* — one flagged observation wearing three faces, since $\rho_\Lambda \sim H_0^2 M_{\rm Pl}^2$ ties them to $\ln(t_0/t_{\rm Pl}) = 140.24 \approx 4K$ — and the e-fold budget campaign that would convert the $4K$ reading into a theorem would convert all three at once. On the anatomy the equivalence is no accident: *the terminus density is the descent's depth squared in the clock* — the 10^{-120} is the interval's odometer at the end of the drive.

$w = -1$ is the semigroup at rest. The exactness theorem (T-Cosmos-DE′) and the no-phantom theorem above acquire their one-sentence root: *arrival at a terminus is irreversible on the fiber that cannot revisit.* A rolling quintessence is a descent still in motion — interval-side and therefore transient by nature, obliged, like inflation, to *end*; $w = -1$ is the descent at rest; and $w < -1$ — phantom crossing — would be the semigroup running backward, un-arriving, which the fiber's own composition law forbids. The three behaviors the data could show are the three things a one-way descent can do at its floor: still rolling ($w > -1$, transient), at rest ($w = -1$, forever), climbing back out ($w < -1$, never). The framework's exposure is maximal, and its reason is topological. The de Sitter future is then the anatomy's last portrait: the interval exhausted — nothing left to forget — and the circle's shares frozen — nothing able to unwind — behind a permanent causal screen whose entropy, e^{8K} up to the bookkeeping above, stands as the terminus's ledger: the fiber that forgets, having forgotten everything it ever could, leaves the fiber that remembers holding a horizon-full of protected shares.

29.11 Bottom line of the terminus

The minimal dark-energy module of GG-A10-Cosmos is built from four connected claims.

(i) The late-time accelerating term is interpreted as a residual GFF/CEP vacuum export of the same throat transition that already governed inflation and the post-inflationary seam selection.

(ii) The meV scale arises naturally from the already-derived lock scale through the combination $M_{\rm lock}^2/M_{\rm Pl,red}$, so no new tiny fundamental scale needs to be inserted.

(iii) The minimal equation of state is $w = -1$, because the residual term is vacuum-like rather than an actively rolling new field.

(iv) The coincidence problem is reformulated as an attractor-slice problem, and in the stronger model-level implementation the present fractions are encoded by the π-locked inventory (97)–(100).

That is enough to give Part IV a coherent late-time conclusion: inflation and dark energy are not two disconnected sectors in the minimal GG-6 story; they are the early and late faces of one vacuum-selection history — the launch and the terminus of one descent.

Table 10: Minimal late-time GG-6 inventory. The flatness relation is theorem-level; the precise π-locked fractions are the minimal model-level implementation adopted here. The columns give the quantity, its closed-form expression, its numerical value, and its interpretation. The distinction the table exists to draw is between the flatness relation, which is theorem-level, and the individual pi-locked fractions, which are the snapshot calibrated to the sky, and a reader who takes both at the same grade will overstate the result.

Quantity	Expression	Numerical value	Interpretation
Total matter fraction	$\Omega_m = 1/\pi$	0.3183	Minimal non-vacuum share of the late-time locked slice.
Baryon fraction	$\Omega_b = 1/(2\pi^2)$	0.05066	Matter share with one additional compact-winding suppression.
Dark-matter fraction	$\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$	0.26765	Remaining matter share, to be interpreted later through the Kaxion window.
Dark-energy fraction	$\Omega_\Lambda = 1 - 1/\pi$	0.6817	Residual vacuum share of the same locked slice.
Equality redshift	$z_{\Lambda m} = (\pi - 1)^{1/3} - 1$	0.289	Epoch when matter and vacuum become equal in the minimal inventory.

30 Two crossings, three legacies

We have followed the descent from launch to terminus, but there is a unification point that deserves to be stated on its own, because it is one of the strongest structural claims of the paper. The descent breaks twice, and the two crossings share out three legacies. The *first* crossing — the graceful exit at the edge of the primary plateau — fixes the post-inflationary normalization and reheats the visible universe: everything luminous descends from it. The *second* crossing — the lock GSSB at the infrared seam, reached by the cooling bath at $t \simeq 1.6 \times 10^{-14}\,\mathrm{s}$ — stocks the dark side. Between them:

1. it fixes the post-inflationary physical normalization $\Omega = 1$;
2. it seeds the compact-angle misalignment that later appears as Kaxion dark matter; and
3. it leaves behind a small residual vacuum term that later becomes the dark-energy sector.

Inflation, dark matter, and dark energy are therefore three observational faces of one descent rather than three unrelated cosmological epicycles. Two geometric crossings, three legacies: the normalization and the visible sector from the first; the Kaxion relic (developed in Part V and closed quantitatively in GG-C1-Kaxion) and the residual vacuum that becomes the terminus from the second. The first saddle populates the visible universe; the second populates the dark one.

Figure 20 draws the same history in the descent's own variables: the scale dimension $r = \ln(E_{\rm GUT}/\mu)$ climbing from the primary plateau ($r_{\rm inf} \simeq 1$) through the prompt lock crossing ($r \simeq 29.4$) to the terminus ($r_{\rm GGA} \simeq 64.5$); the energy scale $\mu = \rho_{\rm tot}^{1/4}$ descending to the dark-energy floor; and the cosmic energy budget $\Omega_i(t)$, where the two crossings are visible as the two hand-offs — modulus to radiation at the first, and the late matter-to-Λ roll through the $\Omega_m = 1/\pi$ crossing that defines the present.

We state this with conviction. The three legacies are not a hopeful coincidence to be confirmed later; they are the structural content of one GRIP descent read at its two saddles, and the framework commits to it. The honest fence is narrow and specific: the dark-matter and dark-energy legacies inherit the same one named topological step (Conjecture VIII) that fixes the π-locked shares, and the launch — once merely asserted — is now derived in class and in observables, with the amplitude its one declared input. Within those named fences the linkage is not a maybe — it is the claim. The descent that lights the universe at its first break stocks it with the dark sector at its second.

The ordering, noted for what it is. One consequence of this history deserves its own sentence, because it will carry a Part of its own weight later: *the universe exploded before the electrons were made.* At the top of the descent there is no matter at all — only the vacuum on its plateau — so the inflationary stretch distributed no things, only the single causal patch itself; the visible sector was minted afterwards, everywhere at once, by one decay of one modulus, and the dark sector's one mark had been frozen while the sky was still one patch. Every agreement we now observe across the sky was therefore inherited, not negotiated. Part X (§72) develops this as the answer to a question the horizon-problem literature circled for fifty years but never posed in its own right: why the universe is everywhere made of the same things.

31 The same principle, from the electron to dark energy: the vacuum generations

Now the spine, made fully explicit. We have watched one descent run from the $E_{\rm GUT}$ edge to the de Sitter future. The claim of this Part — the claim of the whole program — is that this is not a cosmological story that happens to resemble particle physics. It is literally the same geometric principle, read at three scales.

The electron is the *local* ground-state attractor of the GRIP descent. In GG-P3-Electron, the electron's charge and mass are not inputs but the cost-minimum of the same GFF functional, reached by the same cascade: explore the admissible channels (DGI), rank them by the ledger, break the symmetric branch (GSSB), settle into the attractor. The electron is what the descent looks like when it terminates in a single particle's vacuum.

The Kaxion is a *frozen relic* attractor of the same descent, run in the compact angle. In GG-C1-Kaxion, the misalignment released at the lock crossing — the descent's second saddle — settles into a metastable attractor of the very same functional — not a global terminus but a relic that persists because the descent froze it in. It is the same cascade, halted and stored.

Dark energy is the *global, late-time terminus* attractor — the Ground Geometric Attractor, the GGA of GG-A11-Origin. A GFF landscape generally admits many stable basins, many Geometric Attractors; among them exactly one is the global minimum, and that is where the cosmic boundary relaxes when matter has diluted away. Dark energy is what the descent looks like when it reaches the bottom of the whole landscape and stays.

Read them side by side and the unification is mechanical, not poetic. The same GFF functional ranks all three. The same GRIP cascade — explore, rank, break, settle — produces all three. The same attractor language — GA for a transient or relic record, GGA for the persistent terminus — classifies all three. The electron's charge-and-mass and the cosmic vacuum are two readings of one geometric uniqueness statement: the boundary settles where the ledger is least, whether the boundary is one particle or the whole universe.

Here is the jaw-dropper, stated bold and fenced in the same breath: *the principle we have derived for the electron, we now show governs the vacuum of the universe — up to the single named topological step, Conjecture VIII, which the coming decade's data will test.* The electron is earned. The cosmic terminus is earned conditionally, hanging on one closed-packet identification now traced to a single declared input: the timing of the present epoch. That is the whole honest statement: one principle, two scales, one fence. If the fence holds, the same geometry that fixes a single electron fixes the energy of empty space.

31.1 The vacuum generations

The correspondence can be made exact, and it is worth a table in prose. The scale stratum Σ_r carries *three basins* — the primary plateau, the lock basin, the terminus — visited in order by one descent, with *two saddle crossings* between them. The family stratum carries three basins — τ, μ, e — visited in order by one cascade, with two decays between them. The dictionary is entry-for-entry. Three generations of matter; three generations of vacuum. Two decays; two crossings. The electron is the index-zero ground state whose stability is a theorem; dark energy is the GGA of Σ_r whose $w = -1$ is the same kind of theorem — *dark energy is the electron of the vacuum.* Each crossing radiates: the first fills the universe with the visible sector, as a decay fills the detector with its products; the second fills it with the dark sector; relic abundances are the decay kinematics of the vacuum cascade. Even the lifetimes rhyme: the middle basin is the muon of the story, real but brief — crossed, in the vacuum's case, promptly.

Two disciplined footnotes keep the correspondence honest. First, the rung sizes do not transfer:

the family ladder steps by the instanton unit π, the vacuum ladder by the warp depth K — the grammar is shared, each stratum supplies its own constant. Second, the counting claim — that the same 3π budget which forces $N \leq 3$ families also forbids a fourth vacuum epoch (no third inflation, no late entropy dump) — is stated here as the natural conjecture of the correspondence, supported by the prompt-crossing bound, and left for the constitutional analysis to close. If it holds, the universe has exactly as many vacuum generations as matter has family generations, and for the same ledger reason.

31.2 The cosmological triptych: rolling, ringing, rest

The spine just walked — electron, Kaxion, dark energy as local attractor, frozen relic, and global terminus of one cascade — gains a second reading from the campaign's anatomy, and the second reading is this paper's own summary of itself. The three dark episodes of cosmology are the bulk's two non-spatial coordinates and their joint terminus: *inflation is the interval's coordinate, rolling once (§28.15); dark matter is the circle's coordinate, ringing still (§35.8); dark energy is what remains when both coordinates have come to rest — its share written by the circle, its scale written by the interval (§29.10).* Figure 21 draws the three acts on one canvas, and Table 11 sets them side by side.

Ninety-five percent of the present energy budget — dark matter plus dark energy — is, on this reading, the bookkeeping of the two non-spatial fibers; the early universe's defining episode is the open coordinate's one traverse; and the three free mysteries of standard cosmology (what inflated, what the dark matter is, what the dark energy is) are answered by a census: the bulk has one fiber that must forget and one that cannot, and *the becoming and the being each showed up in the sky exactly once. The dark sector is the two fibers, seen in the sky.* No item in the table is new physics: every entry is a published result of this paper or its companions, sorted by fiber, with every grade inherited — the reading adds no claim, only the anatomy that makes the three episodes one story.

32 The cosmic coincidence, bounded: why the inventory is π-locked *now*

The π-locked chart asserts that today the universe is $\Omega_m = 1/\pi$ matter, $\Omega_\Lambda = 1 - 1/\pi$ dark energy, and $\Omega_b = 1/(2\pi^2)$ baryons. A careful reader raises the sharpest possible objection at once. The dark-energy density is a cosmological constant: ρ_Λ is fixed for all time. The matter density is diluted by the expansion as $\rho_m \propto a^{-3}$. The two are therefore equal at a single instant and at no other; before that instant matter overwhelms the vacuum, after it the vacuum overwhelms matter. To read $\Omega_m = 1/\pi$ *off the present sky* is to assert that we happen to live essentially *at* that crossover. Is the π-lock then a snapshot of a fleeting coincidence — a miracle of timing dressed in topological language? This section answers that objection in full: analytically, quantitatively, and without flinching from the part of it that remains open. The conclusion, stated in advance, is that the apparent miracle dissolves into three pieces — a permanent geometric fact, a removed fine-tuning, and a mild, bounded, anthropically expected timing — with exactly one honest residue, which we name.

Table 11: The cosmological triptych: the three dark episodes on the two fibers. Every entry is a published result of this paper or its companions (Arisaka, 2026, C1-Kaxion; Arisaka, 2026, P5-CPT), sorted by fiber; grades inherited. The columns are the three episodes and the rows the features they share; reading across a row is the point of the table, because the claim is that one geometry carries all three rather than that each is separately explained.

	Inflation	**Dark matter**	**Dark energy**
What it is	the interval's coordinate, rolling	the circle's coordinate, ringing	both coordinates at rest
Cascade stage	the opening (DGI explore)	the displaced angle's GRIP-In	the terminus (GGA)
Duration	transient: $N^* = 52.6$ e-folds, once	persistent: ringing since $T_{\rm osc} \simeq 1$ GeV	eternal: $w = -1$, de Sitter
Key numbers	$n_s = 0.9634$, $r = 0.0038$ exact ($1 - 2/N^* = 0.9620$, $12/N^{*2} = 0.0043$ at leading order)	8.1 μeV; 1.95 GHz; one line, alone	$\Omega_\Lambda = 1 - 1/\pi$; $m_\Lambda = M_{\rm lock}^2/M_{\rm Pl,red} \simeq 6$ meV
The circle's part	none at the exit (slow roll fails on the interval); its GSSB in Θ is the later lock crossing; the ledger untouched	everything it *is*: parents, protection, portals, one frozen angle	the share: the uncovered holonomy volume (Hopf)
The interval's part	everything: the roll, N^*, n_s, r, the input's address	everything it *has*: freeze, release, mass delivery, abundance	the scale: the e^{-2K} seesaw; the arrival itself
Why unique	one interval ⇒ one inflaton	one circle ⇒ one axion	one GGA ⇒ one vacuum
Kill criteria	r off 0.0038; n_s off 0.9634 in either direction at 3σ — *met on one of the two 2025 combinations*: the DESI-combined value sits 3.2σ above, the CMB-only value 1.5σ; the tilt is the framework's one currently failing test; supercooled GW background	axion-sector isocurvature; off the 1.59–2.13 GHz window	measured $w \neq -1$; phantom crossing

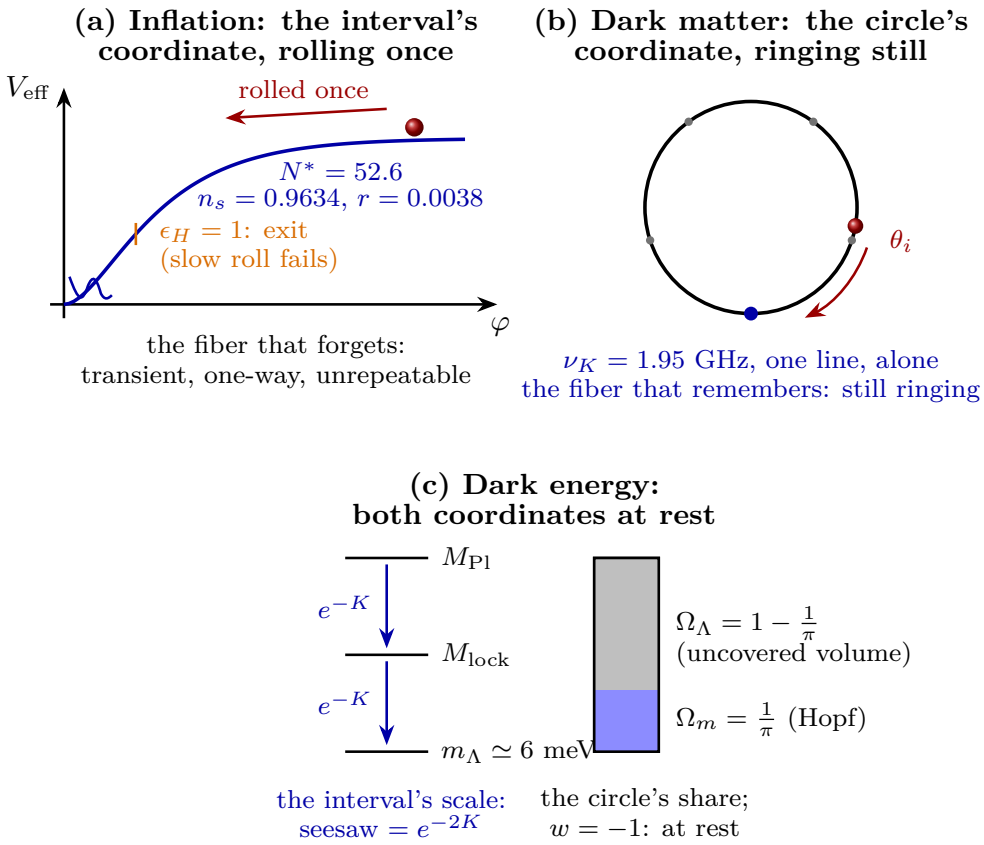

Figure 21: The trio on the two fibers: the three dark episodes of cosmology as the bulk's two non-spatial coordinates and their joint terminus. *(a)* Inflation: the interval's coordinate φ (the radial stratum Σ_r, the only non-conformal scalar the geometry possesses) rolls once along the warp-shaped plateau; the e-fold count $N^* = 52.6$ fixes $n_s = 0.9634$ and $r = 0.0038$; the exit at $\epsilon_H = 1$ is the failure of slow roll at the plateau's edge, and the later lock crossing is the GSSB in the holonomy variable Θ — the circle's first word, spoken after the interval's episode has ended. *(b)* Dark matter: the ledger's phase, held at θ_i (by conjecture), released once, ringing about the attractor floor at 1.95 GHz, standing alone in its window (Arisaka, 2026, C1-Kaxion). *(c)* Dark energy: the share is the circle's — the uncovered remainder of the holonomy volume, $\Omega_\Lambda = 1 - 1/\pi$ (Hopf reading) — and the scale is the interval's — the seesaw $m_\Lambda = M_{\text{lock}}^2/M_{\text{Pl,red}}$ squares the descent's one hierarchy, Eq. (121); $w = -1$ is the semigroup at rest. *The dark sector is the two fibers, seen in the sky.*

32.1 The exact background and the flatness sum rule

We work in the spatially flat Friedmann–Lemaître–Robertson–Walker background that inflation delivers and that the cohort inherits. With the scale factor normalized to $a_0 = 1$ today, the Friedmann equation reads

$$E^2(a) \equiv \frac{H^2(a)}{H_0^2} = \Omega_{r0}\, a^{-4} + \Omega_m\, a^{-3} + \Omega_\Lambda, \tag{122}$$

with the three dilution laws $\rho_r \propto a^{-4}$ (radiation, equation of state $w = 1/3$), $\rho_m \propto a^{-3}$ (pressureless matter, $w = 0$), and $\rho_\Lambda = \text{const}$ (vacuum, $w = -1$). Spatial flatness is the statement $\Omega_k = 0$, and it is preserved by the dynamics; it is the Friedmann constraint itself, not an initial condition that must be re-imposed. The time-dependent density parameters are

$$\Omega_m(a) = \frac{\Omega_m\, a^{-3}}{E^2(a)}, \qquad \Omega_\Lambda(a) = \frac{\Omega_\Lambda}{E^2(a)}, \qquad \Omega_r(a) = \frac{\Omega_{r0}\, a^{-4}}{E^2(a)}, \tag{123}$$

and they obey the exact sum rule

$$\boxed{\Omega_r(a) + \Omega_m(a) + \Omega_\Lambda(a) = 1 \quad \text{at every } a.} \tag{124}$$

This is the first thing the objection must absorb: the *total* never moves. The Ω's never leave the unit simplex; they only *redistribute* weight among radiation, matter, and vacuum as the universe expands. Inflation sets the budget to add to one and the budget stays at one forever. The present input ledger is the OGN-5-fixed $\Omega_m = 1/\pi$ and $\Omega_\Lambda = 1 - 1/\pi$, together with the present radiation content $\Omega_{r0} \simeq 9.2 \times 10^{-4}\%$ — numerically $\Omega_{r0} \approx 9.36 \times 10^{-5}$, the sum of the cosmic-microwave-background photons and the three cosmic neutrino species treated as relativistic. (The massive neutrinos of $\Sigma m_\nu \approx 59$ meV behave as radiation in the early universe and as a sub-percent matter component today; this transition shifts none of the conclusions below and is suppressed throughout for clarity.)

32.2 What is eternal and what is a snapshot: the decomposition

The decisive structural fact is hidden in plain sight in Eq. (123). Baryons and dark matter are both pressureless dust; both dilute as a^{-3}, exactly like the total matter. Their *ratios* therefore carry no factor of a at all:

$$\boxed{\frac{\Omega_b(a)}{\Omega_m(a)} = \frac{1}{2\pi} = 0.15915, \qquad \frac{\Omega_{\rm dm}(a)}{\Omega_m(a)} = 1 - \frac{1}{2\pi} = 0.84085, \qquad \frac{\Omega_b(a)}{\Omega_{\rm dm}(a)} = \frac{1}{2\pi - 1} = 0.18928,} \tag{125}$$

and these hold at $a = 10^{-9}$ and at $a = 10^{+9}$ identically. They are not present-epoch coincidences in any sense: the *internal composition of the matter sector is frozen for all cosmic time.* The relation $\Omega_b = \frac{1}{2\pi}\Omega_m$ is a permanent fingerprint of the high-energy ledger (baryogenesis and the compact-angle relic), not a statement about *when* we observe; and it agrees with the measured cosmic baryon fraction $\Omega_b/\Omega_m = 0.157$ at the percent level, against $1/(2\pi) = 0.159$.

The consequence is that the entire "coincidence" content of the π-locked chart is concentrated in

one number. Of the four headline fractions, three internal ratios are eternal; only the split between matter and vacuum,

$$\frac{\rho_m(a)}{\rho_\Lambda} = \frac{\Omega_m}{\Omega_\Lambda}\, a^{-3} = \frac{a^{-3}}{\pi - 1}, \qquad \text{equivalently} \qquad \frac{\rho_\Lambda}{\rho_m(a)} = (\pi - 1)\, a^3, \tag{126}$$

evolves with time. The "why now" problem is therefore *not* a problem about four miraculous numbers. It is a problem about exactly one time-dependent ratio, and the whole of the analysis below is the anatomy of that single quantity.

Figure 22: The cosmic energy budget through history, shown twice. The four density fractions (radiation, dark matter, baryons, dark energy) always sum to one (Eq. (124)); the universe redistributes weight among them as it expands, while the internal baryon:dark-matter seam stays fixed — the eternal ratio of Eq. (125) — and only the matter:vacuum seam moves. *Top:* against the log scale factor a, which stretches the early universe. *Bottom:* against linear cosmic time t, which stretches the late universe and renders the matter→ Λ handover as the gradual crossover it is. Each panel carries three coordinates — cosmic time t, scale factor $a = 1/(1+z)$, and the radial coordinate $r = r_{\rm now} + \ln a$ (Eq. (134)) — and the π-locked inventory is the snapshot at the $\Omega_m = 1/\pi$ crossing — under Planck central values 75 Myr before the present ($a_\pi = 0.995$); today ($t_0 \approx 13.9$ Gyr at GG-Theory's derived $H_0 \approx 66.6$, the two-road reading; Route I 66.8) the matter fractions sit fractionally lower, $\Omega_m = 0.315$: $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, with dashed guides at $\Omega_{\rm dm} = 0.268$ and $\Omega_m = 0.318$. Transitions are marked at matter–radiation equality ($z \simeq 3400$), recombination ($z_* = e^7$), acceleration onset ($q = 0$, $z = 0.62$), and matter–vacuum equality ($z = 0.29$). The inset in the lower panel resolves the two closing epochs: $\Omega_m(t)$ falls through $1/\pi$ at $t_0 - 75$ Myr (green) and reads 0.315 today (crimson), the separation far inside the ±160 Myr measurement blur.

The cosmic energy budget: the π-locked inventory is the $\Omega_m = 1/\pi$ crossing, 75 Myr before today

Bands always sum to 1 (flat universe). Only the matter:Λ split evolves; the baryon:dark-matter split is frozen for all time.

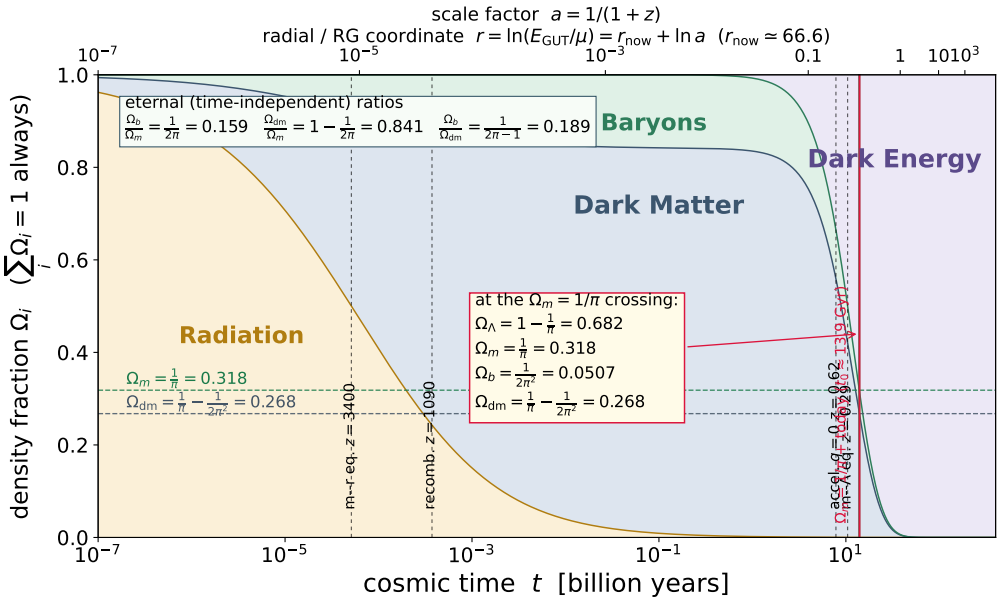

The same evolution in linear cosmic time: the matter → Λ handover, drawn gradually

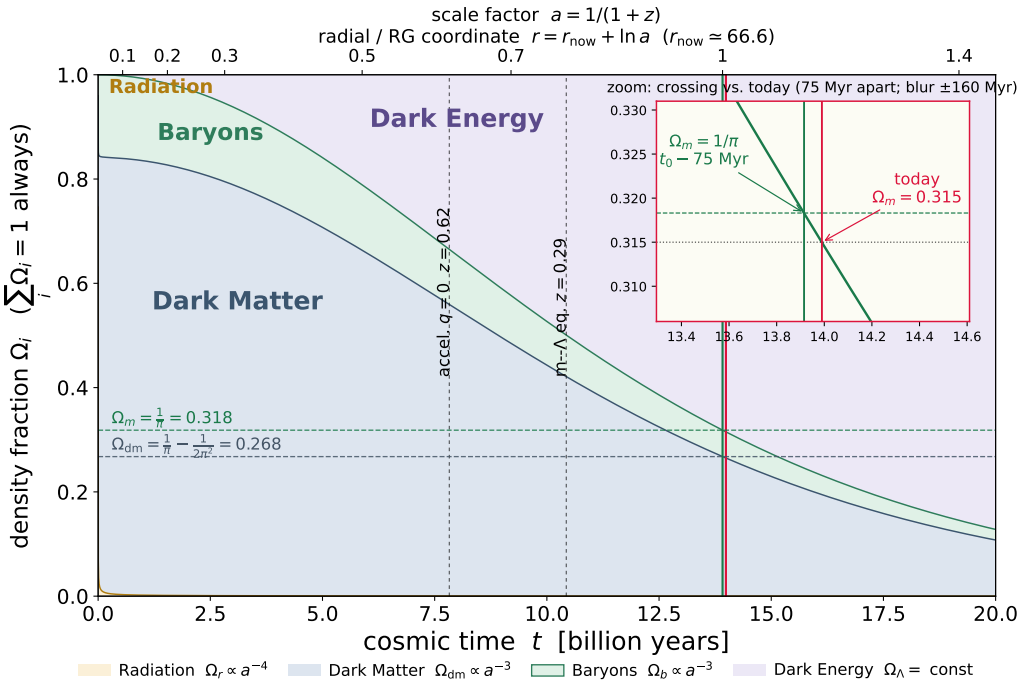

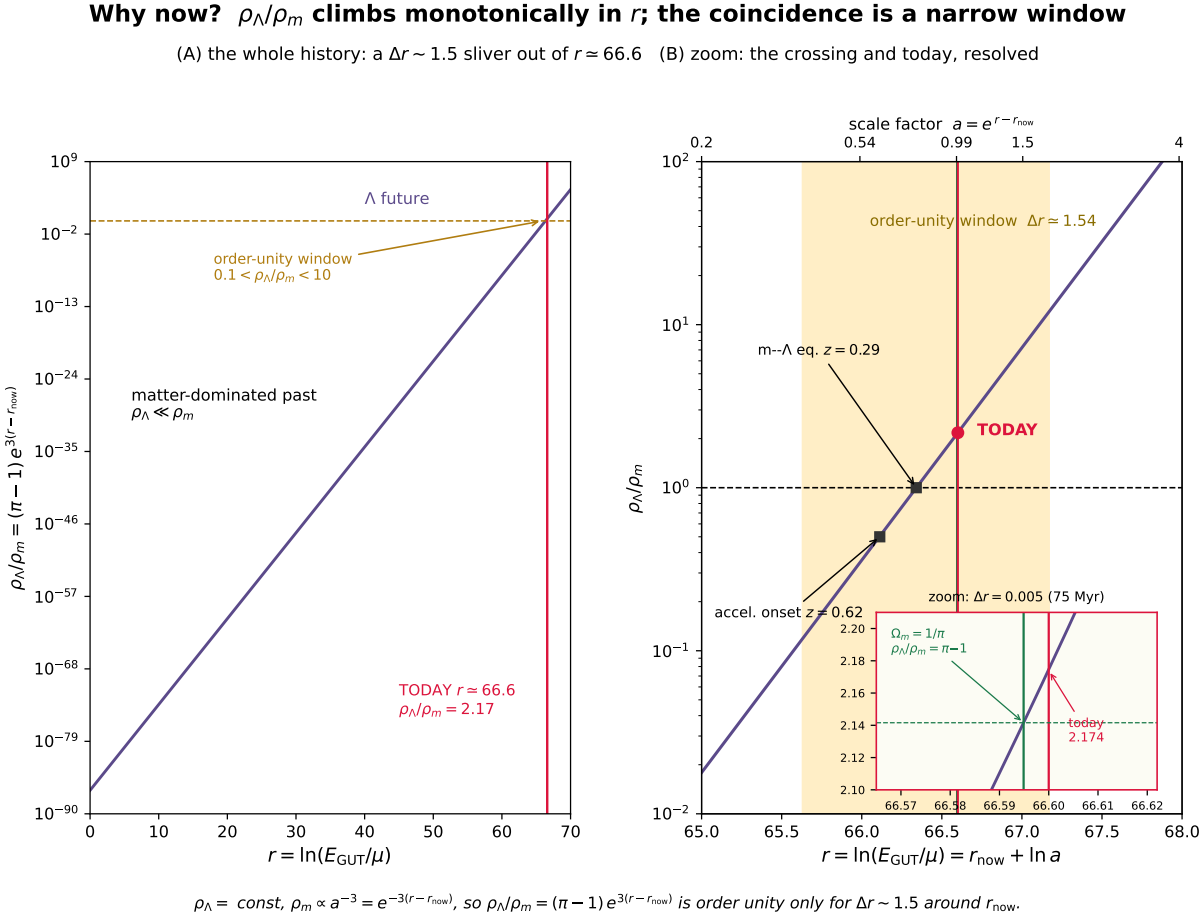

Figure 23: Why now, quantified. The ratio $\rho_\Lambda/\rho_m = (\pi - 1)\, e^{3(r-r_{\rm now})}$ (Eq. (135)) climbs monotonically through the radial coordinate. *Panel A:* the whole history — the ratio sweeps from $\sim 10^{-87}$ at the GUT edge to order unity now; the order-unity observation window is a thin sliver of width $\Delta r \simeq 1.5$ out of $r_{\rm now} \simeq 66.6$. *Panel B:* a zoom showing that acceleration onset ($z = 0.62$), matter–vacuum equality ($z = 0.29$), and today ($\rho_\Lambda/\rho_m = \pi - 1$) all fall inside the $\Delta r \simeq 1.54$ window, with the scale factor on the secondary axis; the magnifier inset resolves the $\Omega_m = 1/\pi$ crossing ($\rho_\Lambda/\rho_m = \pi - 1$, green) from today ($\rho_\Lambda/\rho_m = 2.17$, crimson), $\Delta r = 0.005$, i.e. 75 Myr. We observe inside the $\sim$2% of logarithmic cosmic history where matter and vacuum are comparable — bounded, and exactly where structure formation places observers.

32.3 The full evolution and the four transition epochs

Equations (122)–(123) integrate to the complete history shown as the stacked-band evolution of Figure 22. Reading the budget from the hot past to the cold future, the universe passes through four sharply defined transitions, each of which is a closed-form function of the π-locked inventory and is derived here without a free parameter.

(1) Matter–radiation equality. Setting $\Omega_m\, a^{-3} = \Omega_{r0}\, a^{-4}$ gives $a_{\rm eq} = \Omega_{r0}/\Omega_m = 9.36 \times 10^{-5} \times \pi \approx 2.94 \times 10^{-4}$, i.e. redshift $z_{\rm eq} \approx 3400$. Before this the budget is radiation; after it, matter. This is the only transition involving Ω_{r0} rather than the OGN-5 ledger directly.

(2) Acceleration onset. The deceleration parameter is $q = \frac{1}{2}(\Omega_m + 2\Omega_r - 2\Omega_\Lambda)$. It vanishes, and the expansion begins to accelerate, when $\Omega_m = 2\Omega_\Lambda$ (neglecting the tiny radiation term), i.e. at $\Omega_m = 2/3$. With the π-locked inventory this is

$$a_{\rm acc} = \left(\frac{\Omega_m}{2\Omega_\Lambda}\right)^{1/3} = \left(\frac{1}{2(\pi-1)}\right)^{1/3} = 0.616, \qquad z_{\rm acc} = 0.62. \tag{127}$$

The present deceleration parameter is then the pure number $q_0 = \frac{1}{2}(\Omega_m - 2\Omega_\Lambda) = \frac{1}{2}\left(\frac{1}{\pi} - 2 + \frac{2}{\pi}\right) = \frac{3}{2\pi} - 1 = -0.523$, in agreement with the measured $q_0 \approx -0.53$.

(3) Matter–vacuum equality. Setting $\rho_m = \rho_\Lambda$ in Eq. (126) gives

$$a_{m\Lambda} = \left(\frac{\Omega_m}{\Omega_\Lambda}\right)^{1/3} = \left(\frac{1}{\pi-1}\right)^{1/3} = 0.776, \qquad z_{m\Lambda} = 0.29. \tag{128}$$

This is the geometric heart of the coincidence: the unique instant at which the two dominant components carry equal energy density. Note that acceleration begins *before* equality (at $\Omega_m = 2/3$, while matter still dominates), because acceleration requires only $\rho_\Lambda > \rho_m/2$, a consequence of the $(1 + 3w)$ weighting of pressure in the second Friedmann equation.

(4) The $\Omega_m = 1/\pi$ crossing, and today. The inventory reads exactly $\Omega_m = 1/\pi = 0.318$, $\Omega_\Lambda = 1 - 1/\pi = 0.682$, $\rho_\Lambda/\rho_m = \pi - 1 = 2.14$ at the crossing epoch — under Planck central values 75 Myr before the present ($a_\pi = 0.995$; timing blur ±160 Myr, §32.11). Because every matter fraction (Ω_m, Ω_b, $\Omega_{\rm dm}$) decreases with time while Ω_Λ grows, today's central values sit fractionally below their π-forms ($\Omega_m = 0.315$, $\rho_\Lambda/\rho_m = 2.17$). We are past both acceleration onset and matter–vacuum equality, in the early vacuum-dominated era. The complete numerical history is tabulated in Table 12; note in particular that $\Omega_m = 1/\pi$ coincides with *no* sharp dynamical marker — it is neither equality ($\Omega_m = 1/2$) nor acceleration onset ($\Omega_m = 2/3$) — which is the precise sense in which "today" is a chosen reading rather than a dynamical instant. We return to this in §32.12.

The figure makes the structural point unmistakable. The radiation band reigns until $z \simeq 3400$; the matter band then carries essentially the whole budget ($\Omega_m \approx 0.99$) for the long middle of cosmic history; and the dark-energy band opens only in the last factor of a few in a. The seam between

Table 12: The cosmic energy budget through history in the π-locked inventory. Each Ω_i follows from Eqs. (122)–(123); the internal split $\Omega_b = \Omega_m/(2\pi)$, $\Omega_{\rm dm} = (1 - 1/2\pi)\Omega_m$ is fixed at every epoch by Eq. (125). Radiation entries below $z = 1$ are negligible and omitted. The *age* column is the flat-FRW age integral over the inventory at the derived $H_0 = 66.6$ (two-road reading; Route I 66.8; $H_0 t_0 = 0.948$). The $\Omega_m = 1/\pi$ crossing and today are separated by 75 Myr under Planck central values (timing blur ±160 Myr, §32.11): all matter fractions (Ω_m, Ω_b, $\Omega_{\rm dm}$) decrease with time while Ω_Λ grows, so today's central values sit fractionally below their π-forms.

epoch	z	a	**age**	Ω_m	Ω_Λ	Ω_r	ρ_Λ/ρ_m
deep radiation era	10^9	10^{-9}	∼s	$\to 0$	$\sim 10^{-27}$	≈ 1	2×10^{-27}
matter–radiation eq.	3400	2.9×10^{-4}	51 kyr	0.500	$\sim 10^{-11}$	0.500	5×10^{-11}
corridor switch / decoupling ($z_{\rm sw}$=e^7; atomic $z_{\rm dec}$)	1090	9.2×10^{-4}	0.4 Myr	0.760	1.3×10^{-9}	0.240	1.6×10^{-9}
matter era	10	0.091	0.47 Gyr	0.995	0.0016	0.003	0.0016
acceleration onset $q = 0$	0.62	0.616	7.8 Gyr	0.667	0.333	—	0.50
matter–vacuum equality	0.29	0.776	10.4 Gyr	0.500	0.500	—	1.00
$\Omega_m = 1/\pi$ **crossing**	0.005	0.995	13.8 Gyr	$\mathbf{1/\pi = 0.318}$	$\mathbf{1 - 1/\pi = 0.682}$	9.4×10^{-5}	$\pi - 1 = 2.14$
today (Planck ctr.)	0	1.000	13.9 Gyr	0.315	0.685	9.4×10^{-5}	2.17
near future	−0.5	2.0	25 Gyr	0.055	0.945	—	17
far future	−0.99	100	95 Gyr	$\to 0$	$\to 1$	$\to 0$	2×10^6

baryons and dark matter inside the matter band is a fixed horizontal line — the eternal ratio of Eq. (125) — while the seam between matter and vacuum is the single moving boundary that the coincidence is about.

32.4 The true age of the universe

Of all the numbers in this paper, one is meant to be carried out of it and remembered: *the universe is* 13.9 *billion years old today — and it crossed the exact π-locked reading, $\Omega_m = 1/\pi$, at* 13.8*, seventy-five million years before.* What GG-Theory adds is not (only) a refined value but a new status: both numbers are here *direct derivations* from the geometry, not figures read back from the sky.

The age in two forced factors. The age factorizes cleanly. The flat-FRW age integral over the π-locked inventory returns a pure number, independent of any expansion rate,

$$H_0\, t_0 \;=\; \int_0^1 \frac{\mathrm{d}a}{a\,E(a)} \;=\; 0.948 \qquad (\Omega_m = 1/\pi,\ \Omega_\Lambda = 1 - 1/\pi), \tag{129}$$

fixed once and for all by the matter fraction $\Omega_m = 1/\pi$. To convert it to years one multiplies by the Hubble time $1/H_0$, and here GG-Theory supplies its *own* derived rate — the background value $H_0^{(0)} \approx 67\ \mathrm{km\,s^{-1}Mpc^{-1}}$ derived twice in Section 29.8: by the geometry alone (Route I) and by the π-locked inventory against the sky's densities (Route II). The two forced inputs multiply to the forced age,

$$\boxed{t_0 \;=\; \frac{0.948}{H_0^{(0)}} \;\approx\; 13.9 \text{ billion years} \quad (\text{the } \Omega_m{=}1/\pi \text{ crossing at } 13.8)} \tag{130}$$

No step is tuned to the sky: the dimensionless 0.948 is fixed by $\Omega_m = 1/\pi$, and the rate ≈ 67 is itself a theorem of the inventory (at the D*/Asserted level of the named closures). So 13.9 is here a prediction the geometry hands down for today — and the familiar 13.8 of the textbooks is recovered as the age at which the universe wore exactly the π-locked inventory, 75 Myr before the present.

Why the textbooks say 13.8 — and where that number now lives. The familiar 13.8 is recent: it is essentially the *Planck* value, fixed around 2013 by reading the sound waves frozen into the microwave background through standard ΛCDM, which returns $H_0 \approx 67.4$ and hence $0.948/67.4 = 13.8$ Gyr. That number is careful but *indirect*: it inherits whatever the microwave pipeline assumes. What GG-Theory supplies is a reason the same figure follows from first principles — because the inventory fractions GG-6 derives ($\Omega_m = 1/\pi$, $\Omega_b = 1/2\pi^2$) are exactly the ones that place the background rate at ≈ 67. The agreement of the derived age with the measured one is not a coincidence to be explained away; it is the point.

The age carries a mild tension, because H_0 does. The one wrinkle is the same one that afflicts H_0. The local distance ladder — nearby supernovae calibrated by Cepheid stars — reads

$H_0 \approx 73$, some 9% above the background, and a faster *locally measured* rate would name a younger age. The age therefore inherits the ladder's displacement,

$$\underbrace{66.8}_{\text{derived background (two-road 66.6)}} \rightarrow 13.9 \text{ Gyr (today; } \Omega_m{=}1/\pi \text{ at } 13.8), \qquad \underbrace{73.0\text{–}73.5}_{\text{local ladder}} \rightarrow 12.7 \text{ Gyr.} \tag{131}$$

GG-Theory does not sit on the fence: it *derives* the background rate ≈ 67, so its age today is 13.9 — with the familiar 13.8 as the crossing age — and it reads the ladder's 73 as a measurement carrying a calibration systematic the framework predicts (Section 42; companion GG-C2-Hubble), not the cosmic background rate. The conventional 13.8 therefore survives, relocated with a sharper meaning: it is the age at which the cosmos read exactly $\Omega_m = 1/\pi$, while today, 75 Myr on, is 13.9 — both now derived rather than measured.

A falsifiable bet, settled this decade. This is a prediction, not a slogan, and it has a clean arbiter. The model-independent gravitational-wave standard siren reads the background expansion without either a sound horizon or a distance ladder; GG-A10-Cosmos predicts it will converge near $H_0 \approx 67$, giving $t_0 \approx 13.9$ today. If the sirens instead land near 73, the background is younger and the inventory derivation is wrong. The difference is only a few percent — the universe is "thirteen-and-a-half to fourteen billion years old" either way — but it is a definite, decidable few percent, and GG-Theory has placed its bet on the derived ≈ 67.

The line to remember is short: 13.9 *billion years today, and* $\Omega_m = 1/\pi$ *crossed at* 13.8 — the familiar age relocated by one stellar heartbeat, *derived* from the same ledger that fixes the electron's mass and the dark fraction of the cosmos, now read on the axis of time.

32.5 The future: the de Sitter fate and the exact $a(t)$

For $a \gg a_{\rm eq}$ the radiation term is negligible and the matter–vacuum Friedmann equation integrates in closed form,

$$a(t) = \left(\frac{\Omega_m}{\Omega_\Lambda}\right)^{1/3} \left[\sinh\!\left(\tfrac{3}{2}\sqrt{\Omega_\Lambda}\, H_0\, t\right)\right]^{2/3} = 0.776 \left[\sinh(1.24\, H_0 t)\right]^{2/3}. \tag{132}$$

Setting $a = 1$ inverts to the present age, $H_0 t_0 = 0.948$; converted to years with GG-Theory's own derived rate $H_0^{(0)} \approx 67$ this is the headline value $t_0 \approx 13.9$ Gyr (today; the $\Omega_m{=}1/\pi$ crossing at 13.8), the subject of the dedicated §32.4 above. As $t \to \infty$, $\sinh \to \frac{1}{2}e^{1.24 H_0 t}$ and the universe asymptotes to exact de Sitter expansion,

$$a(t) \propto e^{H_{\rm dS} t}, \qquad H_{\rm dS} = H_0\sqrt{1 - \tfrac{1}{\pi}} = 0.826\, H_0 \;\simeq\; 55.2 \text{ km s}^{-1}\text{Mpc}^{-1}. \tag{133}$$

The fate is an empty, exponentially expanding vacuum with $\Omega_\Lambda \to 1$: matter and radiation dilute to nothing, and the π-locked chart, true only at $a = 1$, is false in the limit. The π-lock is, in this precise sense, a *single photograph* of a budget caught in transit.

32.6 The coincidence quantified: the radial observation window

It is most illuminating to phrase the coincidence in the cohort's own radial coordinate $r = \ln(E_{\rm GUT}/\mu)$, the renormalization-group and expansion axis introduced in the chronology Part. Because the photon temperature obeys $T \propto 1/a$, the radial coordinate and the scale factor are linearly related,

$$r = \ln\frac{E_{\rm GUT}}{\mu} = \underbrace{\ln\frac{E_{\rm GUT}}{T_0}}_{r_{\rm now}\simeq 66.6} + \ln a, \qquad \text{so} \qquad a = e^{\,r-r_{\rm now}}, \tag{134}$$

and the matter–vacuum ratio of Eq. (126) becomes a clean exponential in r,

$$\frac{\rho_\Lambda}{\rho_m} = (\pi-1)\, e^{\,3(r-r_{\rm now})}. \tag{135}$$

This is the content of Figure 23. On a logarithmic vertical axis the ratio is a straight line in r: it climbs from $\sim 10^{-87}$ at the GUT edge ($r = 0$) to order unity now ($r = r_{\rm now}$) and onward to infinity in the de Sitter future. The matter–vacuum crossover sits at $r_{m\Lambda} = r_{\rm now} - \frac{1}{3}\ln(\pi-1) = r_{\rm now} - 0.25$. The interval over which ρ_Λ/ρ_m is of order unity — say between 0.1 and 10, the window in which an observer would call matter and vacuum "comparable" — is

$$\Delta r = \frac{\ln 10 - \ln 0.1}{3} = \frac{2\ln 10}{3} = 1.535, \qquad \Longleftrightarrow \qquad a \in [0.36,\, 1.67], \quad z \in [1.78,\, -0.40]. \tag{136}$$

We observe at $a = 1$, squarely inside this window. The coincidence, stated with full rigor, is therefore this: out of the $r_{\rm now} \simeq 66.6$ units of radial flow from the GUT edge to the present — or, counted from the end of inflation, out of roughly 69 e-folds of cosmic expansion — we find ourselves in a window of width $\Delta r \simeq 1.5$, about 2% of the logarithmic history. That is the whole of the "why now" coincidence, measured: a factor-of-a-few timing, not a tuning of a fundamental constant. The remaining subsections explain why a 2% window is exactly where an observer should expect to find itself, and why the dangerous part of the coincidence has already been removed.

32.7 Physics intuition: why these magnitudes, and why this window

Three physical pictures make the numbers feel inevitable rather than miraculous.

Why matter thins and vacuum does not. Matter is "stuff": a fixed number of particles spread through a volume that grows as a^3, so its density falls as a^{-3}. The cosmological constant is the energy of empty space itself; doubling the volume doubles the amount of space and so doubles the total vacuum energy, leaving its *density* unchanged. A falling function and a constant cross exactly once. There is no choice in the matter: any universe with both dust and a cosmological constant has a unique matter–vacuum crossover, and "now" is simply near it.

Why the internal ratios are eternal but the matter–vacuum ratio is not. Baryons and dark matter are both cold and pressureless; the expansion cannot tell them apart, so it dilutes them in lockstep and their ratio is whatever the high-energy ledger stamped on it — here the geometric $1/(2\pi)$. The matter–vacuum ratio, by contrast, compares a diluting component to a non-diluting one, so it must run with the expansion. The composition of matter is an ultraviolet fingerprint, frozen; the

matter–vacuum balance is an infrared, time-dependent affair. The π-locked chart mixes one of each, and only the second can ever be a "coincidence."

Why a 2% window is where observers live. The growth of cosmic structure — the assembly of the galaxies, stars, and planets that observers require — proceeds while matter dominates and effectively freezes once the vacuum takes over, because accelerated expansion dilutes density perturbations faster than gravity can grow them. The last generation of bound structures is therefore laid down around the matter–vacuum crossover, and observers who arise in them (after the ∼Gyr of stellar and chemical evolution that the cohort's particle ledger makes possible) naturally open their eyes a few billion years later — which is to say, a fraction of an e-fold past the crossover. The matter–vacuum equality time $t_{m\Lambda} \approx 10.4$ Gyr and the present $t_0 \approx 13.9$ Gyr differ by about one stellar lifetime. This is the quantitative core of the anthropic answer first made precise by Weinberg: a randomly chosen observer, weighted by where and when observers can exist, is overwhelmingly likely to find ρ_Λ/ρ_m of order unity, exactly the $\Delta r \sim 1.5$ window of Eq. (136). The picture is the hiker on a long staircase: descending 66 steps from the GUT edge, photographed on the one-and-a-half steps where two slopes cross — not because the staircase was built around the photograph, but because that is the only stretch with anything to see.

32.8 The growth-suppression integral: structure formation is freezing out now

The anthropic picture of §32.7 is not a hand-wave; it is governed by a single integral. In linear perturbation theory the matter density contrast grows as the scale-independent factor

$$D(a) \propto E(a) \int_0^a \frac{\mathrm{d}a'}{[a'\,E(a')]^3}, \qquad E(a) = \sqrt{\Omega_m\, a^{-3} + \Omega_\Lambda}, \tag{137}$$

normalized so that $D \to a$ deep in the matter era. Its logarithmic derivative is the growth rate $f(a) \equiv \mathrm{d}\ln D/\mathrm{d}\ln a$, captured to better than a percent by $f \simeq \Omega_m(a)^{0.55}$: it is unity while matter dominates (perturbations grow in step with the expansion) and falls below unity once the vacuum takes over, because accelerated expansion stretches perturbations faster than self-gravity can amplify them. Evaluated on the π-locked background,

$$f(z{=}1) = 0.88, \quad f(z_{\rm acc}{=}0.62) = 0.80, \quad f(z_{m\Lambda}{=}0.29) = 0.68, \quad f(\text{today}) = \Omega_m^{0.55} = 0.53, \tag{138}$$

and it continues monotonically to zero in the future. Two numbers make the freeze-out concrete. The present growth has been suppressed to $D_0/a_0 = 0.79$ of its matter-era extrapolation; and the *total remaining* growth from today to the infinite future is only

$$\frac{D(\infty)}{D(\text{today})} = 1.41, \tag{139}$$

so cosmic structure will grow by a mere 41% more — ever — before the de Sitter future freezes it permanently. We live in the final fraction of an e-fold of structure formation: the last bound systems large enough to host observers are being assembled essentially now, which is exactly why an observer finds the matter and vacuum densities comparable. The coincidence of Eq. (136) and

the growth freeze-out of Eq. (137) are two faces of one fact.

32.9 The anthropic weight, quantified

The statement "observers arise near the crossover" can be turned into an explicit probability. Weight the observation epoch by the cosmic star-formation rate — the Madau–Dickinson history $\psi(z) \propto (1+z)^{2.7}/[1+((1+z)/2.9)^{5.6}]$, which peaks near $z \simeq 1.9$ — convolved with a delay τ for the stellar, planetary, and biological evolution an observer requires. The probability that such an observer measures ρ_Λ/ρ_m inside the order-unity window of Eq. (136) is then

$$P\Big(0.1 < \tfrac{\rho_\Lambda}{\rho_m} < 10\Big) = \frac{\displaystyle\int_{\text{window}} \psi(t-\tau)\,\mathrm{d}t}{\displaystyle\int_0^\infty \psi(t-\tau)\,\mathrm{d}t} = 0.90\text{–}0.96 \quad \text{for} \quad \tau = 3\text{–}8\ \text{Gyr}, \tag{140}$$

with a median observation time of 8–13 Gyr (redshift $z \simeq 0.3$–0.5, $\rho_\Lambda/\rho_m \simeq 0.6$–$1.0$) — essentially *at* matter–vacuum equality. The result is robust to the delay: across the plausible range of τ, about nine observers in ten find matter and vacuum comparable. This is the quantitative form of Weinberg's anthropic argument on the π-locked background, and it converts "why now" from a puzzle into an expectation — the window is not where the universe was tuned to place us, it is where any star-formation-weighted observer is overwhelmingly likely to be. As always with anthropic estimates the precise number depends on the assumed star-formation history and delay distribution; what is robust, and all the argument requires, is that the probability is of order unity rather than exponentially small.

32.10 The resolution in three movements

The apparent miracle now decomposes cleanly.

First movement — most of the chart is not a coincidence at all. By Eq. (125), the baryon, dark-matter, and total-matter ratios are time-independent: they are exact at every epoch and carry no "why now" content whatsoever. Three of the four π-locked numbers are permanent geometry. The coincidence problem shrinks, rigorously, from four numbers to one: the present value of $\rho_\Lambda/\rho_m = \pi - 1$.

Second movement — the lethal fine-tuning has been removed. The genuinely dangerous quantity in standard cosmology is not the *ratio* but the *absolute* vacuum density, $\rho_\Lambda^{1/4} \sim 2$ meV, some 10^{-120} in Planck units — the worst fine-tuning in physics. GG-Theory does not tune it: the meV scale is a seesaw descendant of the lock seam, $m_\Lambda^{\text{geom}} = M_{\text{lock}}^2/M_{\text{Pl,red}} \sim 6$ meV (the Derived vacuum-scale entry of the codebook), the same infrared residue that the inventory derivation already uses. Once ρ_Λ is a geometric output rather than a tuned input, the 120-order miracle is gone, and what remains of "why now" is only the mild order-unity timing of the first movement. The geometry has converted a catastrophic fine-tuning into a shallow, expected coincidence.

Third movement — the residual timing is anthropically natural and bounded. As §§32.8–32.9 make quantitative, the surviving statement — that we observe within $\Delta r \simeq 1.5$ of the crossover — is precisely where structure-formation physics places observers: cosmic growth has only 41% of its

amplitude left before it freezes (Eq. (137)), and a star-formation-weighted observer sees matter and vacuum comparable with probability 0.90–0.96 (Eq. (140)). It is not eliminated, but it is bounded, printed, and expected. Nothing here is tuned after the fact; the single residual coincidence is the one the universe forces on any observer in any universe with dust and a small cosmological constant.

32.11 What is predicted, and what is not: form, scale, and epoch

It is worth stating with care which features of the inventory the geometry actually fixes, because the headline fractions can otherwise be read as claiming more than is meant. A cosmological inventory carries three logically independent freedoms, and GG-Theory removes two of them but not the third.

Form (removed). The fractions are forced to be *simple, coupled powers of* π: one closed-packet topology fixes $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, and $\Omega_b = 1/(2\pi^2)$ *together* (Conjecture VIII, Theorem 74.1). A curve fit can match one number with one dial, but it cannot manufacture three *coupled* π-powers from a single topological normalization. This coupled form is the genuine, falsifiable structural prediction.

Scale (removed). The absolute vacuum energy is fixed — to order of magnitude — as the geometric seesaw residue $m_\Lambda^{\rm geom} = M_{\rm lock}^2/M_{\rm Pl,red} \sim$ meV, not tuned. This is the step that defeats the 10^{120} cosmological-constant fine-tuning: the vacuum density is no longer an input chosen to one part in 10^{120} but an output of the lock chain.

Epoch (not removed). Here is the freedom that remains, and it is the crux. The matter fraction $\Omega_m(a)$ is *monotonic* — it sweeps continuously from 1 in the matter era to 0 in the de Sitter future (Figure 22) — so the equality $\Omega_m = 1/\pi$ holds at *one instant only*, and which instant we occupy is set by the present temperature T_0, the single observational input named in §32.12. The three absolute fractions are therefore *present-epoch readings*, not epoch-independent constants. The crossing is sharply datable, and its precision deserves stating in full. Under the Planck central values the exact crossing $\Omega_m(a) = 1/\pi$ lies slightly in the *past*: $a_\pi = 0.9949$ ($z_\pi = +0.005$), about 75 Myr before the present at the canonical H_0 — the last 0.5% of cosmic history. But the measurement uncertainty ± 0.007 on Ω_m translates, through $\mathrm{d}\Omega_m/\mathrm{d}\ln a \approx -0.65$, into a timing blur of ± 160 Myr, so the central-value offset is only 0.5σ from zero: *the data are fully consistent with the crossing being now.* The precise statement is therefore $t_0 = t_\pi$ *within measurement precision*, with the central-value reading placing the crossing ≈ 75 Myr in the past inside a ± 160 Myr window; equivalently, a 1% offset in $\pi\,\Omega_m$ corresponds to ≈ 72 Myr of timing.

What "$\Omega_m = 1/\pi$" therefore asserts — and what it does not. Because $\Omega_m(a)$ is monotonic, the only genuinely epoch-independent content of the inventory is the *ratio* $\Omega_b/\Omega_m = 1/(2\pi)$ (together with the vacuum scale); the three absolute fractions $1/\pi$, $1 - 1/\pi$, and $1/(2\pi^2)$ are the present-epoch values of those ingredients. Said most sharply: given the geometric ρ_Λ and the matter-per-photon content the framework fixes, the statement "$\Omega_{m,0} = 1/\pi$" is mathematically equivalent to a specific

present temperature T_0 — which the framework takes from observation. So we do *not* claim that the number 0.318 is an epoch-independent constant on the footing of $\alpha^{-1} = 137$; we claim that the present fractions take a forced, coupled π-form, read at the one epoch we calibrate to the sky. The codebook tags the inventory entries D* under Conjecture VIII for their *form*; this subsection records that they additionally inherit the *epoch* as the L4 calibration input. We make the point explicit precisely to forestall the natural misreading that the absolute value is derived independently of when we look.

The residual coincidence: a few percent, not 10^{120}. This makes the honest bottom line sharp, and it is worth stating without varnish: *a "why now" coincidence genuinely remains.* The standard cosmological-constant problem is in truth two coincidences at once — a fine-tuning of the vacuum energy by one part in 10^{120}, and a timing coincidence that we observe near matter–vacuum equality. The geometry dissolves the first outright (the meV scale is forced, not tuned) and leaves only the second, now bounded to the *few-percent* level: we sit inside a $\Delta r \simeq 1.5$ window out of $r_{\rm now} \simeq 66.6$ — about 2% of cosmic logarithmic history — with anthropic probability 0.90–0.96 (Eq. (140)). The "why now" problem is therefore *not eliminated*, and we do not pretend it is; but its severity has been demoted by roughly 120 orders of magnitude, from a tuning of the vacuum *scale* to a timing of the observation *epoch*. That demotion, not a false claim of no coincidence at all, is the achievement.

GG-Theory removes the *form* freedom (the fractions are coupled π-powers) and the *scale* freedom (ρ_Λ is the forced meV residue), but *not* the *epoch* freedom: $\Omega_m(a)$ is monotonic, so $\Omega_m = 1/\pi$ is a present-epoch reading, not an epoch-independent constant. A "why now" coincidence therefore survives — but at the *few-percent* level (a $\sim$2% timing window), the 10^{120} vacuum-scale tuning having been removed. The coincidence is demoted by $\sim$120 orders of magnitude, not abolished. (141)

32.12 The honest residue: GG-Theory fixes the form and scale, not yet the epoch

Calibrated honesty, the discipline of §85.1, requires naming exactly what is *not* resolved. GG-Theory fixes the *form* of the inventory (the coupled π-powers, Derived-conditional under Conjecture VIII and Theorem 74.1) and the *scale* of the vacuum (the meV seesaw residue). What it does *not* yet do is derive the *epoch*: it does not predict $r_{\rm now}$ — equivalently the present temperature T_0, age t_0, or Hubble rate H_0 — from the geometry, so that "now" would be forced to be the instant at which $\Omega_m = 1/\pi$. Indeed, as Table 12 shows, $\Omega_m = 1/\pi$ lands on no sharp dynamical marker, so the present reading is a present-epoch labeling rather than a geometric instant. This is precisely the *Tier-3* content of the can/cannot ladder (§93.2): the geometry computes the inventory and geometrizes the chronology, but the dynamical selection of the observation epoch — like the precision C_ℓ spectrum — is the irreducible work the framework frames rather than replaces. If a future development of the cohort were to fix $r_{\rm now}$ geometrically, the coincidence would close completely; until then it is honestly carried as bounded and anthropic, not derived.

What fixing the epoch would require. To close the residue is to predict $r_{\rm now} = \ln(E_{\rm GUT}/T_0)$ — equivalently the present photon temperature T_0, the age t_0, or the Hubble rate H_0 — from the geometry alone. It is worth being precise that this is *exactly* the single observational-calibration input the framework already declares: in the input ledger of the cost-of-being-right accounting (Section 94), the present CMB temperature and today's Hubble scale are the one empirical datum admitted, used only to place the boundary history on physical units. The "why now" residue is therefore not a new admission but that same calibration input viewed from the cosmological side: the geometry fixes every *ratio* in the inventory and the *scale* of the vacuum, and takes exactly one number — where "now" sits on the radial axis — from observation. Three routes could in principle close even that: a ground-state GGA selection fixing a preferred number of e-folds of boundary flow; a holographic relation tying the de Sitter horizon entropy $S_{\rm dS} \sim (H_0^2 G)^{-1}$ to the same lock ledger that sets ρ_Λ, which would link H_0 to the vacuum scale; or a structure-formation completeness condition that fixes the epoch at which the growth integral (137) has effectively saturated. None is established here; each is named as the concrete work that would promote the epoch from calibration to derivation.

The π-lock is not a miracle. Its internal matter ratios are *eternal* geometry ($\Omega_b : \Omega_{\rm dm} : \Omega_m = \frac{1}{2\pi} : (1 - \frac{1}{2\pi}) : 1$, true at every epoch); the only time-dependent statement, $\rho_\Lambda/\rho_m = (\pi - 1)a^3$, is the ordinary coincidence problem with its lethal 10^{120} tuning already removed by the geometric origin of ρ_Λ; and what survives is a bounded $\Delta r \simeq 1.5$ ($\sim$2% of cosmic logarithmic history) observation window where structure formation places observers. The one honest residue — a first-principles fixing of the epoch $r_{\rm now}$ itself — is named, conceded, and left open. (142)

32.13 Falsifiable corollaries

The reading above is not merely interpretive; it carries commitments. First, because the internal matter ratio is eternal, the baryon fraction of matter must be $\Omega_b/\Omega_m = 1/(2\pi)$ at *every* redshift at which it can be measured — in the cosmic-microwave-background acoustic pattern ($z \simeq 1090$), in the Lyman-α forest ($z \sim$ 2-4), and in clusters today — with no evolution beyond the standard astrophysical baryon budget; a measured drift of Ω_b/Ω_m with redshift away from $1/(2\pi)$ would falsify the eternal-ratio claim. Second, the present deceleration parameter is the pure number $q_0 = \frac{3}{2\pi} - 1 = -0.523$; a secure determination away from this value falsifies the π-locked split independently of the inventory's normalization. Third, the matter–vacuum equality redshift is fixed at $z_{m\Lambda} = (\pi - 1)^{1/3} - 1 = 0.29$ and the acceleration-onset redshift at $z_{\rm acc} = (2(\pi - 1))^{1/3} - 1 = 0.62$; both are zero-parameter postdictions of the same ledger, and both are now within reach of supernova and baryon-acoustic-oscillation surveys. None of these is tuned; each is a forward-checkable consequence of reading the cosmic budget as one geometry caught, today, in transit.

33 Falsification and bottom line

A unification this strong is worth nothing if it cannot be killed, so we end by exposing it. Three live stands carry the Part:

(i) $w = -1$ **exactly.** The terminus is a completed transition, not a rolling field, so the framework predicts a vacuum-like equation of state pinned at -1 with no significant late-time clustering. This is a clean falsifiable bet against the recent DESI hint of evolving dark energy: if the measured $w(z)$ settles decisively away from -1, the terminus reading fails.

(ii) $r = 0.0038$**, the tensor target.** The derived plateau forces Eq. (73): a firm central value, a factor nine below the present bound and several times above the Simons-Observatory/LiteBIRD floor. A tensor measurement inconsistent with it beyond the launch's spread 0.0037–0.0040, in either direction, kills the derived launch; so would a tilt excluding 0.9634 at three standard deviations in either direction (a value below 0.955 would in addition demand a diluted window) or a loud supercooled-transition gravitational-wave background (a camped lock). The switches are two-sided and cannot be repaired without abandoning the derived class itself.

(iii) **The de Sitter fate.** $w = -1$ and positive Ω_Λ on a near-flat universe force an asymptotic de Sitter future — no Big Rip, no recollapse. Evidence for phantom behavior ($w < -1$) or for a closing universe would contradict the terminus directly.

To these one should add the inventory test: if the measured shares settle stably away from $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, then Conjecture VIII is false and the closed-packet derivation is wrong. That is a single-module correction — the late-inventory interpretation is retired, but the launch and the dark-matter legacy do not automatically fall with it.

33.1 The significance: two unknowns, one hierarchy

It is worth standing back, at the close of the descent, to say what has and has not been done — because the two sectors this Part treats are exactly the two places where the concordance model falls silent. ΛCDM does not explain inflation; it posits an inflaton with a potential chosen from a wide menu. It does not explain dark energy; it inserts a cosmological constant whose magnitude is the worst quantitative mismatch in the history of physics. These are the model's two great unknowns, and they are normally treated as unrelated.

The structural content of Part IV is that they are not unrelated: they are the two ends of one descent. The launch is the derived plateau at the top — class forced by the boundary law, height the one declared input, sitting one e-fold below the GUT edge with $H_1 \simeq 1.5 \times 10^{13}\,\mathrm{GeV}$ — and it computes its own observables: $n_s = 0.9634$, $r = 0.0038$, $N^* = 52.6$, $T_{\rm rh} \simeq 6 \times 10^{10}\,\mathrm{GeV}$. The terminus is held at the meV scale by the inherited lock, $m_\Lambda = M_{\rm lock}^2/M_{\rm Pl,red} \approx 6\,\mathrm{meV}$, which with the ledger coefficient $\zeta_\Lambda = \pi/LK$ takes the notorious 10^{120} vacuum-energy discrepancy to $\rho_\Lambda^{1/4} = 2.23\,\mathrm{meV}$ against a measured 2.24, landing the dark-energy density on its observed value rather than merely in its corridor; and the lock's own basin — the middle of the descent, crossed promptly — is the very place whose curvature scale *is* that seesaw, $H_2 = M_{\rm lock}^2/\sqrt{3}M_{\rm Pl}$. The lock

no longer sets the launch; it sets the middle crossing that stocks the dark sector and the terminus that ends the story. To derive the top of cosmic history in class and in numbers, and to reduce its bottom to a single, independently fixed scale — analytically, with the magnitudes falling out rather than being fitted — is the distinctive contribution of this Part.

The honesty of the claim lives in the grading, and it bears repeating without softening. What is *forced and quantitative* is now a long list: the plateau class (GDOS), the tilt $n_s = 0.9634$, the tensor target $r = 0.0038$, the e-fold count and reheating temperature, the promptness of the lock crossing, the meV floor of the terminus, and the vacuum equation of state $w = -1$ — sharp, falsifiable, and now two-sided where they used to be one-sided. What remains *open* is named, not hidden: the scalar amplitude A_s — and through it the plateau height — is the sector's one declared input, its warp candidate (75%) flagged in the codebook, with the GDOS heat-kernel evaluation the one route still standing; the seesaw coefficient ζ_Λ is not yet computed from its Wilson line (GG-A5-GDOS O6), its codebook-selected value $\pi/175$ pre-registered as a two-sided target; and the π-locked fractions $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$ are derived-conditional under the single closed-packet conjecture (Conjecture VIII). The boldness and the honesty are once more the same statement read at two resolutions: the top of the descent is now derived, the bottom reduced to one fixed scale, and the ledger of what that has and has not yet closed is handed to the reader in full.

In plain words, this Part has told one story three times. Inflation is the descent's launch; dark matter is the relic it froze at its middle crossing; dark energy is the ground it finally reached — *the electron of the vacuum.* Where the standard account carries three unrelated mechanisms with three sets of dials, this framework carries one geometric descent and no dials at all — which is why it can be wrong, and why it is worth the reader's attention.

The tiered close is therefore short. The electron ledger is derived; the Kaxion relic rides on one declared calibration and the single-patch conjecture. The inflationary tilt and tensor ratio are derived-conditional, with the amplitude the launch sector's one declared input. The π-locked inventory and dark energy are derived-conditional under Conjecture VIII. Within those fences, Part IV has watched one open-system descent run from the lawful beginning of Part III to the de Sitter terminus — a launch, a break with three legacies, and a terminus — governed throughout by the same GRIP geometry that fixes the electron. The next Part picks up the first of those three legacies: the origin of matter, and the Kaxion that the graceful-exit event seeded.

34 Part IV in plain English: the two ends of a single descent

The organizing idea of this Part is that inflation and dark energy are not two phenomena that happen to bracket cosmic history but the opening and the terminus of one relaxation. The same runtime that settles an electron into its ground state settles the universe into its vacuum, and the two epochs are the first and the last legs of that descent.

If this is right, the cosmic coincidence changes character. The near-equality of the matter and vacuum densities in our own epoch stops being a coincidence that needs a selection effect to explain and becomes a statement about where in the descent an observer necessarily sits. The paper is careful about how much this buys: the argument bounds the coincidence rather than eliminating it, and a bound is weaker than a derivation.

The π-locked inventory is the sharpest thing in the Part and it should be read with its grade attached. The coupled closed form for the three fractions is derived conditional on a named conjecture, and the individual present-day values are a snapshot calibrated to the sky rather than epoch-free predictions. Two of those statements are much stronger than the third, and the table that carries them says which is which. A reader who takes the agreement at three percent as confirmation of the whole structure has read more into it than the Part claims.

Part V
Origin of matter: Kaxion dark matter, baryogenesis, and BBN

Kaxion cold dark matter from the compact-angle sector with the π-locked relic abundance $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$, baryogenesis through the inherited neutrino-and-CP ledger, and Big Bang nucleosynthesis with the lithium-7 audit

---

The second cluster of the open cosmological problems is the origin-of-matter cluster: what the dark matter is, why baryons dominate over antibaryons, and why the early-universe nuclear audit is simultaneously so successful and yet not completely closed. Within the architecture of the present paper this cluster is the most direct bridge between GG-A9-SM and cosmology. The earlier particle paper already fixed the narrowest pieces of the ledger that this part needs: the scale chain $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock}$, the exact $\mathrm{U}(1)_B$ protection of proton stability, the compact-angle ontology on S^1_θ, the minimal Majorana/Weinberg channel for the light neutrinos, the leptonic CP structure, and the general rule that source sectors should first be sought inside the inherited geometry before new freedom is introduced.

For that reason Part V is organized as a deliberately tight chain rather than three unrelated mini-models. First, the Green–Schwarz tensor forced by anomaly freedom and its compact-angle dual suggest that dark matter should be sought in the same protected compact sector that already entered the particle paper through holonomy, strong-CP suppression, and anomaly inflow. That leads to the Kaxion window. Second, once the neutrino block is no longer arbitrary but derived, the matter–antimatter problem is naturally driven toward leptogenesis rather than direct baryon violation, because the exact local $\mathrm{U}(1)_B$ symmetry of GG-6 strongly disfavors any proton-destabilizing route. Third, Big Bang nucleosynthesis becomes the first precision audit of that same matter ledger. BBN is where one cannot hide behind qualitative rhetoric: the expansion history, baryon density, neutrino sector, and any nonstandard ports or residual interactions must all survive comparison with light-element abundances.

A useful way to summarize the logical order is

$$\begin{aligned} &\text{compact angle} + \text{BI-6} \Longrightarrow \text{Kaxion window} \Longrightarrow \Omega_{\rm DM}, \\ &\text{Majorana neutrino block} + \text{CP structure} \Longrightarrow \text{leptogenesis backbone} \Longrightarrow \eta_B, \end{aligned} \tag{143}$$

followed by the first high-precision audit,

$$(\Omega_b,\, \eta_B,\, \text{dark-sector history, radiation ledger}) \Longrightarrow \text{BBN} \Longrightarrow (Y_p,\, D/H,\, {}^3\mathrm{He}/H,\, {}^7\mathrm{Li}/H). \tag{144}$$

This is the correct spirit in which to read the remainder of the part. The aim is not to present three disconnected solutions. The aim is to show that the same inherited GG-6/GDOS ledger naturally

feeds all three problems.

Before entering the individual derivations it is helpful to record the specific data imported from GG-A9-SM, because Part V depends on them repeatedly.

Table 13: Inherited particle-sector inputs from GG-A9-SM that are used repeatedly in Part V. No measured dark-sector or baryogenesis parameter is inserted by hand; only the already-derived particle and geometric ledger is used. The columns give the quantity, a representative inherited value, and why it matters here. Every entry is quoted at the grade its owning paper gives it; where that paper hedges, this one hedges in the same place.

Quantity	**Representative inherited value**	**Why it matters in Part V**
Lock scale	$M_{\rm lock} = E_{\rm Pl}/(2e^{35}) \simeq 3.85\,\text{TeV}$	Sets the natural scale for heavy-throat relics and leptogenesis
GUT edge	$E_{\rm GUT} = E_{\rm Pl}e^{-2\pi} \simeq 2.28 \times 10^{16}\,\text{GeV}$	Sets the ultraviolet seam for the Majorana/Weinberg operator and early source ledger
OGN-5 integers	$(M, N, L; J, K) = (1, 3, 5; 7, 35)$	Fix the minimal discrete arithmetic behind many ratios used below
Exact baryonic gauge protection	local $\mathrm{U}(1)_B$	Explains why proton decay is forbidden while baryogenesis must proceed indirectly
Compact angle and Green–Schwarz tensor	S^1_θ holonomies (lifted) and the tensor's dual	The tensor's dual is the Kaxion; the holonomies carry the phase-driven CP structure
Light neutrino masses	$(m_1, m_2, m_3)_{\rm QD} = (0.0653, 8.5923, 50.5131)\,\text{meV}$	Fix the boundary neutrino block that a later leptogenesis module must descend to
Neutrino mass sum	$\Sigma m_\nu \simeq 59.17\,\text{meV}$	Provides the cosmology-facing neutrino benchmark
PMNS Dirac phase	$\delta^{\rm QD}_{\rm PMNS} \simeq 196.06^\circ$	Supplies the order-one leptonic CP throughput used by leptogenesis
Majorana phases	$(\alpha_{21}, \alpha_{31}) = (0, \pi)$	Supply the minimal parity structure of the neutrino block
Jarlskog invariant	$J_{\rm PMNS} \simeq -9.2 \times 10^{-3}$	The rephasing-invariant measure of leptonic CP that controls ϵ
Baryonic ledger entry	$-2/5$ of the SU(5) anomaly entry (GG-A3-BI6)	Fixes the Kaxion's baryonic bookkeeping; no light channel below $M_{\rm lock}$ under the boundary-condition termination

The scientific posture of Part V should also be stated explicitly. Dark matter, baryogenesis, and BBN are all mature subjects with vast literatures. The present section therefore does not claim that the standard literature becomes irrelevant. What it claims is narrower and sharper: once the inherited ledger is accepted, the order in which possible explanations are explored is no longer arbitrary. The first candidate for dark matter is the compact-angle descendant of BI-6. The first

candidate for baryogenesis is the leptogenesis channel of the already-derived Majorana neutrino block. The first place where the framework can fail quantitatively is BBN. This ordering is the central methodological gain of Part V.

35 Kaxion dark matter

35.1 Historical background

The axion arose as a solution to the strong-CP problem of QCD. Peccei and Quinn in 1977 (Peccei and Quinn, 1977) proposed a new global U(1) symmetry whose spontaneous breaking would dynamically relax the QCD vacuum angle to zero. Weinberg (Weinberg, 1978) and Wilczek (Wilczek, 1978) pointed out independently in 1978 that this mechanism necessarily produces a light pseudo-Nambu-Goldstone boson — the axion. The cosmological role of the axion as cold dark matter via the misalignment mechanism, and the more general compact-angle descendants in higher-dimensional theories, follow from the same idea. The Kaxion of the present paper is a compact-angle pseudo-scalar inheriting the Peccei–Quinn template through the anomaly ledger of GG-6 — the boundary dual of the Green–Schwarz tensor that anomaly freedom forces — rather than through a postulated low-energy U(1).

The Kaxion misalignment dark-matter candidate is situated within the broader history and candidate landscape of dark matter reviewed in (Bertone & Hooper, 2018).

Theorem 35.1 (Compact-angle dark-matter theorem (T-Cosmos-7))**.** *If the light compact-angle mode surviving the BI-6/TGG/CCI descent is the dominant cold relic of the compact-angle sector, then its abundance must be matched not to an arbitrary Ω_c but to the late attractor target*

$$\Omega_{\mathrm{DM}} = \frac{1}{\pi} - \frac{1}{2\pi^2},$$

and its effective phenomenology is axion-like but logically narrower than a free Peccei–Quinn axion sector.

Proof. BI-6 and the compact-angle ontology already guarantee that a light axion-like direction is available in the theory without importing a separate global symmetry: the boundary dual of the Green–Schwarz two-form is a compact angle by the large gauge transformations of the tensor. TGG places that direction inside the same total cosmic packet as inflationary exit and late vacuum residue. CCI makes its four-dimensional response lawful on the seam. GFF ranks the misalignment branch after reheating. GRIP then turns the primary descent, the release of the angle at the lock crossing, and the late cold relic export into one event chain. The resulting relic abundance is therefore constrained by the same late attractor chart that fixes the total matter share. ■

35.2 Why dark matter should be sought first in the compact-angle sector

The dark-matter problem is old and concrete at once. Across galaxies, groups, clusters, the CMB, and structure formation, ordinary baryonic matter does not supply enough gravitating mass

to account for the observed dynamics or the measured growth history. At the level of the boundary ΛCDM bookkeeping one writes

$$\Omega_m = \Omega_b + \Omega_{\rm DM}, \qquad \Omega_{\rm DM} \gg \Omega_b, \tag{145}$$

with baryons a minority of the present matter density. The bottom-up literature then chooses among a long menu — WIMPs, sterile neutrinos, axions, fuzzy scalars, primordial black holes, self-interacting hidden sectors. That freedom is useful, but in the present program it is the first thing to be reduced.

The reason is direct. GG-A9-SM and its predecessors already contain a compact angular degree of freedom, already carry extra Abelian structure through $\mathrm{U}(1)_B \times \mathrm{U}(1)_{Qi}$, and already use holonomy and Wilson–Hosotani logic for symmetry breaking and strong-CP suppression; GG-A3-BI6 forces a Green–Schwarz tensor on the same bulk to cancel the anomaly. The first dark sector to investigate is therefore not a new unrelated particle inserted by hand. It is the protected low-energy remnant of the same anomaly ledger — the Kaxion.

The name "Kaxion" marks that the candidate is axion-like but not a relabeled Peccei–Quinn axion. It descends from the Green–Schwarz sector of GG-6, so its existence is motivated by geometry already in play rather than by an auxiliary global symmetry added for cosmological convenience. That matters physically, not just philosophically: it narrows the admissible mass and coupling windows.

A compact-angle relic has four immediate virtues. It is naturally neutral and colorless; naturally long-lived, because its leading couplings are derivative or anomaly-descended rather than unsuppressed renormalizable decays; naturally cold, because a coherently displaced compact-angle field behaves as a pressureless condensate once oscillations begin; and naturally tied to the strong-CP and holonomy sectors already present in the particle paper. That combination is exactly what one wants before any relic calculation.

What is *not* claimed is that all alternatives are ruled out. The claim is the disciplined rule that, within the inherited GG-6/GDOS ledger, the compact-angle descendant is the first candidate to explore, because it is the one already written into the geometry before any new freedom is introduced.

35.3 From BI-6 and the compact angle to an axion-like 4D field

The shortest route from the six-dimensional bulk to a four-dimensional axion-like field uses two structures already present in the earlier papers:

1. the Green–Schwarz tensor sector, with a two-form B_2 and field strength H_3, and
2. the compact angular coordinate $\theta \in S^1_\theta$, whose holonomies already control part of the symmetry-breaking logic.

The Green–Schwarz tensor admits the usual duality

$$H_{\mu\nu\rho} = f_{\rm GS}^{-1}\, \epsilon_{\mu\nu\rho\sigma}\, \partial^\sigma a_{\rm GS}, \tag{146}$$

where a_{GS} is a pseudoscalar dual to the two-form and periodic under the tensor's large gauge transformations, while the compact sector contributes the Wilson-line holonomies of the rigid circle,

$$\vartheta_a(x) = \oint_{S^1_\theta} A^a_\theta(x,\theta)\,\mathrm{d}\theta, \qquad a = 1,\dots,6, \tag{147}$$

one for each Cartan direction of the rank-six bulk gauge group. The two kinds of angle have different fates, and the difference is the whole content of the census. The holonomies are the angles of a gauge field on a circle of unit metric: their potentials are generated at the circle's own scale, and GG-A9-SM's coset zero mode registers that lift near 2×10^{15} GeV (its PS-SM-22; none of the six lifts is computed, and the companion GG-C1-Kaxion carries them as *expected*, not derived). The tensor's dual has no such potential: on the intended background class no scalar–scalar bilinear couples it to the holonomies, and the only potential it acquires below the bulk scale is the QCD-generated cosine of the gluonic portal. The compact sector therefore hands the boundary exactly one light pseudoscalar — the tensor remnant — and that is the Kaxion. The earlier reading of this sector, in which the tensor's dual and one holonomy were two light angles geared together into a doublet, is withdrawn by the mode analysis of GG-A3-BI6 (its section on the self-duality of the Green–Schwarz tensor, open items O13/O14); GG-C1-Kaxion records the withdrawal in full.

After descent to the boundary and coarse-graining of heavy modes, the low-energy quadratic Lagrangian of the light field consistent with the inherited symmetries is

$$\mathcal{L}_{\mathrm{ax}} = \frac{1}{2}\,\partial_\mu a_K \partial^\mu a_K - \chi_{\mathrm{QCD}}\left[1-\cos\!\left(\frac{a_K}{f_K}\right)\right] + \frac{\alpha}{8\pi}\,C_\gamma\,\frac{a_K}{f_K}F_{\mu\nu}\widetilde{F}^{\mu\nu} + \cdots, \tag{148}$$

with f_K the physical decay constant of the canonically normalized boundary eigenstate and the ellipsis denoting higher operators and the $\mathrm{U}(1)_B$ and $\mathrm{U}(1)_{Qi}$ entries of the anomaly ledger, which have no light vector to couple to (below). There is exactly one physical Kaxion, with one decay constant, one mass, and one laboratory line.

The next step is standard axion physics. Once QCD generates a topological susceptibility χ_{QCD}, the light eigenstate acquires a mass

$$m_K^2 f_K^2 \simeq \chi_{\mathrm{QCD}}, \tag{149}$$

or, in the familiar phenomenological form,

$$m_K \simeq 5.7\,\mu\mathrm{eV}\left(\frac{10^{12}\,\mathrm{GeV}}{f_K}\right), \tag{150}$$

up to the quark-mass sensitivity the companion prices through the lattice susceptibility ($\pm1.3\%$ in the mass at fixed f_K). Within ordinary axion phenomenology Eq. (150) is simply the QCD mass–decay-constant relation; within GG-6 it becomes informative, because f_K is not scanned. It is calibrated once, and after that calibration the geometry's integers leave it a discrete ladder rather than a line.

35.4 Mass and decay-constant benchmark: one calibration, then a ladder

The companion paper Arisaka-GG-C1-Kaxion (Arisaka, 2026, C1-Kaxion) states the convention this paper adopts, and it should be stated here in the same words, because an earlier reading of both papers said something stronger. The physical decay constant is the bulk Green–Schwarz normalization reduced by an integer cap loading, $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$. The bulk normalization is *not* derived: the orientation convention of the anti-self-dual sector that would fix it is an open item of GG-A3-BI6 (its O14), and neither face-value chiral norm of the tensor reproduces the calibrated value. GG-C1-Kaxion therefore *calibrates* $f_{\text{bulk}} = 3.04 \times 10^{13}\,\text{GeV}$ through the relic closure — the requirement that, at the benchmark loading $\Xi_{\text{cap}} = K + N + L = 43$, a generic misalignment angle saturates the π-locked share — and declares the calibration as such. Everything after it is arithmetic: $f_K = 7.1 \times 10^{11}\,\text{GeV}$, the mass by Eq. (149), the frequency by the standard conversion. The cap loading itself is read from the ledger by an ansatz, $\Xi_{\text{cap}} = K + \Sigma$ with Σ a subset sum of $\{N, L, J\}$ in which at most a pair appears — a restriction the ansatz stipulates, not derives — giving the seven rungs $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$; the ansatz is graded as one in GG-C1-Kaxion (its Ansatz on the cap-loading ladder), and the benchmark row is the marked rung, not a derived point.

Table 14: The single physical Kaxion benchmark, as the companion paper Arisaka-GG-C1-Kaxion (Arisaka, 2026, C1-Kaxion) states it. One scale is calibrated through the observed abundance and declared; the numbers after it are arithmetic in that scale, the five integers and one lattice input. The bulk normalization ($3.04 \times 10^{13}\,\text{GeV}$) is a bookkeeping quantity, $\Xi_{\text{cap}} f_K$, not a propagating line.

Quantity	Value	Status
Bulk normalization f_{bulk}	$3.04 \times 10^{13}\,\text{GeV}$	calibrated through the relic closure at $\Xi_{\text{cap}} = 43$; declared (GG-A3-BI6 O14 open)
Physical decay constant f_K	$7.1 \times 10^{11}\,\text{GeV}$	$f_{\text{bulk}}/43$; the ladder $\Xi_{\text{cap}} \in \{35, \ldots, 47\}$ is an ansatz
Mass m_K	$8.1\,\mu\text{eV}$ ($\pm 1.3\%$ at fixed f_K)	the QCD image of f_K, $m_K f_K = \sqrt{\chi_{\text{QCD}}}$; the uncertainty is the lattice susceptibility's
Frequency ν_K	$1.95\,\text{GHz}$ ($\pm 25\,\text{MHz}$); window 1.59–$2.13\,\text{GHz}$	benchmark rung; the window is the ladder's residual, $-19\%/+9\%$ about the benchmark

These follow from Eq. (150) and the standard conversion

$$\nu_K\,[\text{MHz}] \simeq 241.8 \times m_K\,[\mu\text{eV}], \tag{151}$$

which gives $8.06\,\mu\text{eV} \Leftrightarrow \nu_K = 1.95\,\text{GHz}$, with the ladder spanning $[1.59, 2.13]\,\text{GHz}$. The window is discrete: its rungs are 45–136 MHz apart, wider than the $\pm 25\,\text{MHz}$ the lattice susceptibility leaves at fixed f_K, so a line on a rung would read Ξ_{cap} under the ansatz, and a line between rungs would retire the ansatz — not the Kaxion. One qualification travels with that, from the companion's own

accounting: the closest pair of rungs, $\Xi_{\rm cap} = 42$ and 43, are separated at only 1.8σ on the $\pm 25\,\mathrm{MHz}$ line width, so "resolved" holds of the ladder as a whole and not of every adjacent pair in it (Arisaka, 2026, C1-Kaxion). A laboratory apparatus never measures a basis; it measures *eigenstates*, and the canonically normalized light eigenstate is the one physical field every boundary portal operator couples to. This is a real narrowing of the dark-matter landscape — not a corridor of decades, but a window of half a gigahertz with a marked rung — and it is a narrowing bought with one declared calibration, which this paper will not describe as a derivation.

The OGN integers enter in one place that is a theorem and one that is not. The lock scale is fixed by

$$M_{\rm lock} = \frac{E_{\rm Pl}}{2e^{K}} = \frac{E_{\rm Pl}}{2e^{35}} \simeq 3.85\,\mathrm{TeV}, \tag{152}$$

and it is at the crossing of this scale that the compact angle is released (below). The same integers enter the window through the ladder $K + \Sigma$, and that entry is the ansatz just named. An earlier reading of this sector put the integer K to a third use — a twin-peak splitting $\Delta\nu/\nu = 1/K$ of the haloscope line — which the count of GG-C1-Kaxion withdraws: the second light angle that the doublet required was never supplied by the anomaly companion, and the prediction is now one line, standing alone, with the doublet's absence a test.

35.5 Relic abundance from misalignment and the π-locked inventory

The real cosmological test is not only whether the field exists but whether it accounts for the observed dark-matter abundance without re-importing uncontrolled freedom. The Kaxion is attractive because its preferred production channel is the standard coherent misalignment mechanism: simple, cold, robust.

At very early times, when $3H \gg m_K$, the compact-angle field is overdamped and sits at an initial displacement θ_i from its minimum. (In runtime terms this is GRIP-In on the compact-angle stratum: Hubble friction is the within-basin flow, and the misalignment energy is the height above the basin floor that the descent converts into the relic.) When the Hubble rate falls to

$$3H(T_{\rm osc}) \simeq m_K, \tag{153}$$

the field begins coherent oscillation and thereafter redshifts as matter,

$$\rho_K \propto a^{-3}, \tag{154}$$

a natural cold-dark-matter component.

For a QCD-anchored axion-like field in the single-patch misalignment calculation the abundance is

$$\Omega_K h^2 \simeq 0.135\,\theta_i^2 \left(\frac{f_K}{10^{12}\,\mathrm{GeV}}\right)^{1.165} \qquad \text{(harmonic)}, \tag{155}$$

where the exponent encodes the temperature dependence of the QCD topological susceptibility near oscillation onset, and the prefactor is the direct integration of the equation of motion through onset that GG-C1-Kaxion performs. The familiar instantaneous-onset estimate, which sets $3H = m$ at a

single moment and gives 0.090, understates the relic by a factor 1.5; an earlier version of this paper printed the estimate as if it were the integrated value, and Appendix M.5 now carries both with the correction stated. For an order-one angle the full cosine is integrated rather than its harmonic approximation, and that is the number this paper uses.

The logic runs in a direction that must be stated exactly, because it is easy to overstate. Equation (155) has two unknowns, f_K and θ_i, once the bulk normalization is recognized as open. The convention of GG-C1-Kaxion, adopted here, closes it in one step: the benchmark loading $\Xi_{\rm cap} = 43$ is fixed, $f_{\rm bulk}$ is calibrated so that the benchmark row saturates the π-locked share at a *generic* angle, and the closure then returns that angle,

$$\theta_i \;=\; 1.18 \pm 0.04 \text{ rad} \qquad (\text{tabulated } \chi_{\rm QCD}(T),\ \text{full cosine}), \tag{156}$$

the generic $\mathcal{O}(1)$ misalignment angle every standard axion-dark-matter scenario assumes — neither the unbroken point 0 nor the hilltop π. The angle is read back, not derived and not tuned; the calibration is declared once; and after it nothing about the Kaxion is adjustable. An earlier version of this paper carried the angle as 1.4 rad, the value the instantaneous-onset estimate returns; the integrated pipeline supersedes it.

The target itself is the π-locked cosmic inventory of the present paper. Within it the matter and baryon shares are

$$\Omega_m = \frac{1}{\pi} \simeq 0.3183, \qquad \Omega_b = \frac{1}{2\pi^2} \simeq 0.0507, \tag{157}$$

so the cold-dark-matter share is

$$\Omega_{\rm DM} = \Omega_m - \Omega_b = \frac{1}{\pi} - \frac{1}{2\pi^2} \simeq 0.2676, \tag{158}$$

against the Planck 2018 value $\Omega_m - \Omega_b = 0.2657 \pm 0.0070$ (Planck's $\Omega_m = 0.315 \pm 0.007$ less $\Omega_b = 0.0493$; the cold component alone, $\Omega_c = 0.2645$, excludes the neutrino share, and the companion compares against it). This is the target the Kaxion relic calculation must saturate, not a substitute for it. One convention must be printed with it: converting the share to the comoving density $\Omega_{\rm DM}h^2 = 0.1215$ that Eq. (155) uses takes $h = 0.674$, the Planck-fit value that the companion adopts; at this paper's carried rate $H_0 = 66.8$ the same share is 0.1194 (angle 1.17 rad), and at the two-road reading 66.6 it is 0.1187 (1.17 rad), both on the companion's tabulated full-cosine integration. The difference is a convention, not an error, and it moves nothing else.

The cosmological point stands, and it is the point this Part is entitled to: no unrelated dark sector had to be imported to obtain a cold relic, and one declared calibration, not a fit, makes the inherited geometry carry the observed cold-dark-matter abundance at a generic angle.

35.5.1 Oscillation onset, free streaming, and why the Kaxion behaves as cold dark matter

The onset estimate connects the compact field to structure-formation language directly. During radiation domination

$$H(T) \simeq 1.66\, g_*^{1/2} \frac{T^2}{M_{\rm Pl}}, \tag{159}$$

so $3H(T_{\rm osc}) \simeq m_K$ gives

$$T_{\rm osc} \simeq \left(\frac{m_K M_{\rm Pl}}{4.98\, g_*^{1/2}}\right)^{1/2} \propto m_K^{1/2}. \tag{160}$$

For $m_K = 8.1\,\mu$eV this naive onset lies in the radiation era (the temperature-dependent susceptibility lowers it to $\simeq 1\,$GeV, $t \simeq 0.3\,\mu$s, in the full computation of Arisaka-GG-C1-Kaxion), far later than the GUT-scale history. Either way the field starts oscillating deep in the radiation era and redshifts as nonrelativistic matter — the only behavior required for cold dark matter.

The smallness of m_K should not be read as "hot" or "warm." The criterion is whether the relic is a thermal relativistic gas or a coherent classical condensate; the Kaxion is the second type. Its momentum dispersion after onset is set by the Hubble scale, not a bath temperature, so its velocity dispersion redshifts fast ($\sim 4 \times 10^{-14}$ at equality) and its de Broglie scale is microscopic; the free-streaming length is $\sim 10^{-6}\,$pc and the Jeans length $\sim 3 \times 10^{-4}\,$pc. At $8.1\,\mu$eV the Kaxion is some sixteen orders heavier than the ultralight fuzzy regime $m \sim 10^{-22}\,$eV, so it behaves as an ordinary cold component in linear and nonlinear structure formation, not a wave-supported suppression field. Part VI tests the growth history against S_8, the matter power spectrum, and CMB lensing; within the precision reach relevant to CMB, BAO, and large-scale growth, the Kaxion is effectively cold and collisionless.

35.5.2 The history of the angle: released at the lock crossing, single-patch by conjecture

The inflationary fork that an ordinary axion faces is closed here on one side and left open, by name, on the other. A Peccei–Quinn axion may follow the pre- or post-inflationary history because its angle is the phase of a complex scalar whose radial mode can be thermally restored. The Kaxion has no such scalar: its angle is the dual of a tensor on a rigid geometry, so there is no order parameter to restore, no thermal epoch in which separate patches could acquire separate dials by that route, and no defect network seeded by restoration. That closes thermal randomization. It does not close de Sitter randomization, which needs no radial mode: during the primary (GUT-adjacent) inflation of Part IV, with $H_1 \simeq 1.5 \times 10^{13}\,$GeV, a light angle would random-walk by $H_1/(2\pi f_K) \simeq 3.4\,$rad per e-fold. An earlier version of this paragraph called the result of that random walk the *post-inflationary class*, and drew from it the consolation that a uniformly distributed angle has the same coarse-grained density in every patch and so imprints no microwave-background isocurvature. That reading is withdrawn, and the reason it fails is the reason this Part's conjecture is load-bearing. The post-inflationary class is the case in which the angle is randomized *after* inflation, on comoving scales that are small today; it requires a symmetry restoration, and the compact angle has none — it is a geometric datum carried by the cycle at every epoch, and no sector of the ledger restores it. Randomization by de Sitter diffusion *during* inflation is the opposite case: at $3.4\,$rad per e-fold the angle equilibrates on the circle in about a quarter of an e-fold, while the observable microwave sky spans about seven e-folds of comoving separation, so every observable scale carries an independent draw. A uniformly distributed angle has fractional density variance $\langle\theta^4\rangle/\langle\theta^2\rangle^2 - 1 = 4/5$, an order-unity cold-dark-matter isocurvature against a measured adiabatic power of 2.1×10^{-9}: the isocurvature power exceeds the power the *Planck* bound $\beta_{\rm iso} < 0.038$ allows by ten orders of magnitude. And no

decay constant this geometry supplies evades it, though the argument changes form as f grows: at the bulk scale $f_{\rm bulk} = 3.04 \times 10^{13}\,\mathrm{GeV}$ the step is 0.08 rad per e-fold and the angle is not randomized, but an unpinned angle then carries the ordinary pre-inflationary fluctuation $\delta\theta = H_1/(2\pi f)$, whose isocurvature $(2\delta\theta/\theta_i)^2$ exceeds the bound by eight orders of magnitude. Holding the angle unpinned and quiet enough requires $H_1/(2\pi f) \lesssim 5 \times 10^{-6}\,\mathrm{rad}$ at the misalignment angle $\theta_i = 1.1$ (the bound moves by less than a factor of three over $\theta_i \in [0.5, 3]$), that is, $f \gtrsim 5 \times 10^{17}\,\mathrm{GeV}$, four orders above the largest scale the compact sector carries. What an unpinned angle would still fix, for the record, is the abundance: circle-averaged it closes the inventory only at $f_K \lesssim 2.3 \times 10^{11}\,\mathrm{GeV}$, a line at $\gtrsim 6\,\mathrm{GHz}$ from misalignment alone, above the window, with an order-unity minicluster spectrum at the onset horizon. But that reading is not an alternative history the sector may adopt; it is a history the isocurvature measurement has already excluded.

The single-patch history that this Part's relic calculation presumes therefore rests on a conjecture — and, by the isocurvature argument just given, on a *necessary* one: this is not a choice between two histories with different signatures, but the only history compatible with the measured isocurvature bound. The present paper carries the conjecture under the companion's name; the dark-matter companion rejects the same branch on its own grounds (a line at or above 6 GHz, miniclusters) and carries the pinning as C-Kaxion-d, and the isocurvature argument above is this paper's addition:

> ***Conjecture C-Kaxion-d (inflation-era pinning; GG-C1-Kaxion).*** *During the observable e-folds of primary inflation the compact angle carries an effective mass* $m_{\rm pin} \gg H_1$ *— a pinning potential* $V_{\rm pin} \gg H_1^2 f_K^2 \simeq 10^{50}\,\mathrm{GeV}^4$*; the conjecture is two-valued, held or not held — whose minimum sits displaced by the misalignment angle,* $\Delta\theta \simeq 1.1\,\mathrm{rad}$*, from the CP-conserving point of the QCD potential (so it is not the strong-CP mechanism of GG-A9-SM); which releases the angle at the lock crossing identically in every horizon patch; and which is absent below the lock crossing — switched off by some* 10^{64} *— to the level* $V_{\rm res} \lesssim 4 \times 10^{-15}\,\mathrm{GeV}^4$ *that the neutron bound* $|\bar\theta| < 10^{-10}$ *allows. No sector of the ledger is presently known that could supply such a potential.*

An earlier version of this paper stated the pinning as forced — "the same descent-built suppression that protects strong CP today, only mildly wound down near the GUT edge, pins the angle at its geometric minimum" — and cited the companion for it. That sentence is withdrawn on two grounds. The companion carries the pinning as a conjecture, not a result, so the citation was circular; and the minimum named was the wrong one: a potential that pins the angle at the CP-conserving point delivers $\theta_i = 0$ and no relic, whereas a 1.1 rad deposit requires a minimum displaced from that point by 1.1 rad and switched off afterward. The computation that would establish or kill the conjecture — the compact angle's effective potential on this paper's own primary-inflation background — is the sharpest open item this Part hands to the companion, and it is registered in the codebook under Lemma I (§68.2.3) as the condition that lemma presupposes.

What follows from the conjecture, and what would kill it, are both plain. Under it the angle is released once, at the lock crossing (T of order a few TeV as the bath crosses $M_{\rm lock}$; the radiation Hubble rate equals the lock basin's curvature $H_2 \simeq 3.5 \times 10^{-12}\,\mathrm{GeV}$ at $T \simeq 1.6\,\mathrm{TeV}$), identically in every patch; the crossing is prompt ($N_2 = 0$) and imprints no super-horizon perturbation of any kind,

so $\beta_{\rm iso}$ vanishes identically rather than merely lying below the Planck bound — the isocurvature budget is decided by primary inflation, not by the lock epoch, and an earlier version of this paper that quoted a tiny $\beta_{\rm iso}$ from the lock epoch's H_2 was answering at the wrong epoch. A confirmed Kaxion minicluster population, or a confirmed line at the randomized-angle mass ($\gtrsim 6\,\mathrm{GHz}$), kills the conjecture — and, because the unpinned history is excluded by the isocurvature bound whatever decay constant is chosen, kills with it this paper's compact-angle route to the dark-matter share, which would then have to be carried by something else. A CMB isocurvature detection attributable to the axion sector would kill the conjecture outright — the angle was not held — and the paper does not carry a partially pinned branch to retreat to. This is the sharpest exposure in Part V, and it is stated here rather than left for a reader to assemble.

35.5.3 Why the Kaxion is not merely a rebranded Peccei–Quinn axion

Four differences matter structurally. First, the Kaxion is not introduced by postulating and then breaking a new global Peccei–Quinn symmetry at a free scale; its origin is geometric, descending from the Green–Schwarz sector that anomaly freedom already forces. Second, it shares its ancestry with the holonomy and strong-CP sectors already used in GG-A9-SM, so one geometric structure participates in symmetry breaking, CP suppression, and dark matter at once. Third, it carries a signature no generic axion-like particle carries: a pre-declared window of seven rungs around one calibrated benchmark, and a *count* — one line, standing alone across the window, the count conditional on the parents' registered light modes (GG-C1-Kaxion's T-Kaxion-5) — so that a second in-window line at the same grade would refute the count and name which registered mode is light. Fourth, the framework already links the Kaxion to the cosmic-inventory target of Eq. (158): it is the first and most natural candidate expected to saturate the nonbaryonic share left once baryons are fixed, which turns it from an optional embellishment into a central bookkeeping object of the matter budget.

In one sentence:

> The Kaxion is axion-like in effective low-energy phenomenology, but its logical status is entirely different: it is a geometric descendant of the anomaly ledger's Green–Schwarz sector, not an auxiliary repair field introduced only after the cosmological problem is posed.

That difference of status is why it deserves to be the primary dark-matter candidate of GG-A10-Cosmos.

35.5.4 What the direct-search program would actually look for

Translated into resonant-search language, a haloscope is not asked to scan all axion-like masses with uniform prior weight. It is asked to examine one window and to ask three questions of it:

1. a single narrow line in the window $\nu_K \in [1.59,\, 2.13]\,\mathrm{GHz}$ of Table 14, the benchmark rung at $1.95\,\mathrm{GHz}$ first;
2. the *address*: the line on the DFSZ coupling line, $E/N = 8/3$, at $g_{K\gamma\gamma} = 1.2 \times 10^{-15}\,\mathrm{GeV}^{-1}$;

3. the *count*: the line alone across the window — a deliberate test, since a single-mode cavity suppresses a detuned second line by some 74 dB and must be retuned, or a second mode instrumented, to ask it;

together with cosmological consistency of the same benchmark with the relic-density target. A candidate satisfying all three would reproduce not one number but the whole fingerprint of the compact-angle sector as the companion now states it. The registered kill conditions are the companion's, with one verdict sharpened here (the companion reads a line at the randomized-angle mass as killing the pinning alone, F-Kaxion-1): a window swept at DFSZ depth and silent kills the calibrated placement and its ladder (F-Kaxion-1c, at $g_{90} \leq 0.8\, g_{\rm DFSZ}$); a confirmed compact-angle-like line outside the window does the same (F-Kaxion-1d), and at the randomized-angle mass kills the pinning conjecture and with it this paper's compact-angle route to the dark-matter share; a confirmed in-window line whose coupling sits off the DFSZ locus at $> 5\sigma$ refutes the address (this paper's addition, registered here as F-Kaxion-1e for the companion to adopt); a second in-window line refutes the count.

35.6 Couplings, lifetime, and the single-line haloscope signature

A dark-matter candidate is useful only if its interactions are weak enough to evade bounds, strong enough to be observable in principle, and constrained enough to define a sharp search. The Kaxion performs on all three. The effective low-energy couplings are inherited from the anomaly ledger:

$$\mathcal{L}_{K,\rm int} \supset \frac{\alpha_s}{8\pi f_K} C_g\, a_K\, G\widetilde{G} + \frac{\alpha}{8\pi f_K} C_\gamma\, a_K\, F\widetilde{F} + \cdots , \tag{161}$$

where the ellipsis carries the ledger's $\mathrm{U}(1)_B$ and $\mathrm{U}(1)_{Qi}$ entries. The baryonic entry is exact at the anomaly level — the coefficient vector of GG-A3-BI6 gives it as $-2/5$ of the SU(5) entry —

$$C_B : C_g = -\tfrac{2}{5} : 1, \tag{162}$$

but it has no laboratory channel: under the boundary-condition termination of GG-A3-BI6 the gauged $\mathrm{U}(1)_B$ has no vector below the lock scale — its Z'_B sits at $M_{\rm lock}$ as a Dirichlet-terminated mode that endpoint-localized quarks do not produce at tree level (GG-P1-KK; conditional on GG-A9-SM's cascade-depth lemma) — so the ledger entry survives as bookkeeping and as a chiral-level nucleon coupling out of any spin port's reach at this f_K. An earlier reading of this sector inverted the entry into a light-channel power ratio of 5:2 between baryonic and electromagnetic readouts; that inversion rested on a light vector the termination does not supply, and GG-C1-Kaxion withdraws it. Using

$$g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K}\left(\frac{E}{N} - 1.92(4)\right), \tag{163}$$

with the anomaly ratio fixed by the companion paper's anomaly-matching theorem at $E/N = 8/3$ — placing the Kaxion precisely *on the DFSZ line* — one finds

$$g_{K\gamma\gamma} = (1.2 \pm 0.07) \times 10^{-15}\,\mathrm{GeV}^{-1} \qquad (\pm 5.5\% \text{ at fixed } f_K), \tag{164}$$

the value any haloscope sees (Arisaka, 2026, C1-Kaxion). This is small enough to evade ordinary detection yet large enough that DFSZ-depth resonant cavities of the late 2020s reach it directly. The lifetime is correspondingly enormous: for $a_K \to \gamma\gamma$,

$$\tau_{K\to\gamma\gamma} \sim \frac{64\pi}{g^2_{K\gamma\gamma} m_K^3}, \tag{165}$$

giving

$$\tau_{K\to\gamma\gamma} \sim 4 \times 10^{32}\, t_U, \tag{166}$$

with t_U the age of the universe, so cosmological stability is automatic.

The direct-search signature is a single line, and the paper says so because an earlier version said the opposite. The compact sector contains one light zero-mode, not two: the holonomy that would have paired with the tensor's dual is lifted, so there is no doublet, no locked spacing, and no second channel. What replaces the doublet as the identification test is the count — one line, alone, across a window narrow enough to interrogate completely — and the count is the sharper test, because it is scale-free and because its failure would not merely disappoint but would name a registered mode of the parents as light. The signal logic is summarized in Table 15.

Table 15: Signal logic of the Kaxion window. The columns give the observable, what the framework expects of it, and what makes that expectation distinctive rather than generic. The third column is the one that matters: an expectation shared with every axion-like model tests the class, not this framework, and the rows are ordered so that the distinguishing signatures come first.

Observable	**Framework expectation**	**Why it is distinctive**
Mass / frequency	Window $\nu_K \in$ [1.59, 2.13] GHz, seven rungs, benchmark 1.95 GHz	One calibrated scale and an integer ladder (an ansatz), not a scan over decades; the single laboratory eigenstate of the compact-angle sector
Count	One line, standing alone across the window	Derived-conditional on the parents' registered light modes; a second in-window line names one of them
Address	On the DFSZ line, $E/N = 8/3$	Exact anomaly matching; the coupling carries $\pm 5.5\%$ at fixed f_K
Lifetime	cosmologically stable by many orders of magnitude	Automatic for a compact-angle pseudoscalar in the inferred mass-coupling window
Relic behavior	cold, pressureless after oscillation onset	Standard misalignment logic with no need for a thermal WIMP freeze-out sector

The practical message is clear. Dark matter is not solved by declaring "there exists an axion-like field." It is addressed by a narrow window with a marked rung, an address, and a count. If the

window is exhausted at DFSZ depth without a signal, the calibrated placement and its ladder fail — the correct scientific standard.

35.7 Why the Kaxion is the first dark-matter candidate of GG-A10-Cosmos

The Green–Schwarz descendant of GG-6 makes an axion-like cold relic not merely possible but structurally expected: naturally stable, naturally cold, naturally connected to sectors already present in the particle paper, and naturally associated with a narrow direct-search window. That is sufficient reason to make it the primary dark-matter sector of GG-A10-Cosmos. The strength of the GG-2026 sequence is not that it reproduces the entire dark-sector literature; it is that, given the inherited ledger and one declared calibration, it can say where one should look first. If the compact-angle sector supplies the dominant matter relic, the baryonic component must be explained separately and precisely — which leads directly to the matter–antimatter asymmetry.

35.8 The Kaxion on the two fibers: the anomaly ledger's own phase, ringing still

The dark-matter companion develops the full portrait (Arisaka, 2026, C1-Kaxion); this subsection carries its essence at the cosmological register, so that the middle act of the trio (§28.15, §29.10) stands beside its siblings in this paper's own pages. The campaign's triptych sentence places it: *the electron is a holonomy class wearing a mass; the proton is two holonomy ledgers wearing a weld; the Kaxion is the anomaly ledger's own phase, wearing an occupation.* Dark matter, in GG-Theory, is the sixth dimension relaxing into its ground attractor after the fifth has carried the universe below the barrier.

Everything the Kaxion *is* is circle-side: it is the boundary dual of the tensor the anomaly ledger forces — the closed fiber's bookkeeping made field — and its would-be siblings, the circle's own holonomies, are the fiber's other creatures, lifted and absent from the light spectrum; its lightness is the circle's memory (a periodic direction admits no polynomial potential — the Kaxion is light because the circle cannot forget); its portals are the BI-6 ledger's, transported unchanged — it couples only through topology because it is *made* of topology, which is the entire origin of its darkness; and its universality is one angle — one value for the whole observable universe, the condensate analog of the electron's sameness, which this paper carries as the pinning conjecture rather than as a theorem. Everything the Kaxion *has* is interval-side: the freeze is the Hubble friction — the semigroup itself gripping the angle while $3H \gg m_K$; the release at $T_{\rm osc} \simeq 1$ GeV is the moment the angle's clock outruns the interval's, crossed exactly once because the descent is one-way; the mass $m_K = \sqrt{\chi_{\rm QCD}}/f_K$ is delivered by the descent as it cools through the QCD scale; and the relic abundance is the memory of the crossing, redshifted — the misalignment condensate is primordial, unique, and unreplenished as a corollary of the fiber anatomy, given the single-patch history. The pendulum equation of the misalignment mechanism is the anatomy in one line: the field is the ledger's phase, the $3H\dot\theta$ friction is the interval acting on it, and the stiffness is QCD's, switched on by the boundary at a temperature the one-way history fixes.

Two features make this the cleanest attractor exhibit in the cohort, and both belong in the cosmological narrative. First, *the Hessian is the observable*: the dark-matter mass is the curvature

of the QCD-generated cosine at the attractor floor, read through the one calibrated scale, so the 1.95 GHz line is the runtime's termination certificate, heard as a tone — a haloscope detection would be the first direct measurement of a GGA-class Hessian. Second, *the cascade is still finishing*: the electron's and proton's cascades terminated within the first second; the Kaxion's coherent oscillation is the final GRIP-In of the observable universe, caught mid-ring. Between the interval's opening act and the joint terminus, the circle's act is the one still playing — which is precisely why it is the one the laboratory can catch.

36 Baryogenesis via the inherited neutrino and CP ledger

36.1 Historical background

The necessary conditions for dynamically generating a baryon-antibaryon asymmetry from a symmetric initial state were stated by Sakharov in 1967 (Sakharov, 1967): baryon-number violation, C and CP violation, and departure from thermal equilibrium. Every viable mechanism in the past sixty years realizes those three conditions. The present section asks which realization is admissible once the neutrino sector and the CP ledger have been fixed by the upstream particle paper, rather than chosen freely.

The electroweak sphaleron transport and leptogenesis mechanisms that tie the baryon asymmetry to the inherited neutrino and CP structure are established in (Fukugita & Yanagida, 1986; Kuzmin *et al.*, 1985).

Proposition 36.1 (One-ledger baryogenesis statement (T-Cosmos-8))**.** *The baryogenesis module is not an independent CP-violating add-on. It is the first thermal transport block in which the inherited Majorana neutrino and CP data of GG-A9-SM are forced through the cosmological ladder and then matched to the late baryon share* $\Omega_b = 1/(2\pi^2)$.

$$\dot{n}_{N_i} + 3Hn_{N_i} = -\Gamma_i\left(n_{N_i} - n_{N_i}^{\mathrm{eq}}\right), \qquad Y_B \sim c_{\mathrm{sph}}\,\kappa \sum_i \epsilon_i\, Y_{N_i},$$

the point being that the ϵ_i, Γ_i, and the allowed heavy-sector structure are fixed by the inherited particle ledger rather than left phenomenologically free.

It is worth recording which constitutional pillar carries each of Sakharov's three conditions, because none of the three is imported. *Baryon-number violation*: the bulk baryonic ledger is exactly conserved (gauged $\mathrm{U}(1)_B$); what evolves is the projected subsystem's budget, and the $\Delta L = 2$ step occurs only at the dial-free door of Proposition 36.2. *C and CP violation*: the CP-odd input is a real, lattice-locked angle of the inherited flavor geometry — in the real presentation a loop holonomy of the coupling texture, the flavor analog of an Aharonov–Bohm angle (Arisaka, 2026, A9-SM; Arisaka, 2026, A8-QM). *Departure from equilibrium*: constitutional — the boundary law is a forward semigroup, openness being the default (Arisaka, 2026, A5-GDOS); the transport equation above is a one-way relaxation; and the resonance condition that lifts the Yukawa suppression is itself an openness condition, met where the width matches the lock splitting. In one line: GRIP owns the angle, GDOS owns its observability, OGGEP owns the ledger.

36.2 Why the matter–antimatter problem becomes sharper after GG-A9-SM

The observed universe is overwhelmingly baryonic. The comoving baryon asymmetry,

$$\eta_B \equiv \frac{n_B - n_{\bar{B}}}{n_\gamma}, \tag{167}$$

is nonzero and small rather than zero. The late-universe benchmark is

$$\eta_B^{\rm obs} = (6.12 \pm 0.04) \times 10^{-10}, \tag{168}$$

the CMB-inferred value from the Planck 2018 baryon density (Planck Collaboration, 2018). In ordinary model building this is a target to be generated by some out-of-equilibrium mechanism satisfying the Sakharov conditions. Here the question is narrower and harder. Because GG-A9-SM has already fixed a neutrino sector and a nontrivial leptonic CP structure, the matter–antimatter problem is no longer an invitation to add arbitrary phases or heavy states. It is the question of whether the already-derived neutrino and CP ledger admits a viable leptogenesis descendant.

This sharpening matters because exact local $\mathrm{U}(1)_B$ is a structural claim of GG-6: baryon number is not something the framework may casually violate at low energy, and exact proton stability is one of its cleanest falsifiers. Any viable mechanism must therefore (i) leave the low-energy proton absolutely stable while (ii) the early universe nonetheless generates a nonzero net baryon abundance. This pushes the framework almost automatically toward leptogenesis: a lepton asymmetry created first and partially converted to baryon number by electroweak sphalerons, with no explicit proton-destabilizing operators in the late universe. The choice is not arbitrary. The earlier papers already select the minimal Majorana/Weinberg route for the light neutrinos; once a Majorana channel exists, lepton number is not exactly gauged as baryon number is, so the lepton sector is the path of least resistance for an early asymmetry.

36.3 Sakharov conditions in GG-6/GDOS language

(i) Effective B generation without late-time proton decay. In the GG-6 bulk local $\mathrm{U}(1)_B$ is exact, so the total six-dimensional baryonic ledger is conserved. But the four-dimensional boundary universe is an open subsystem, and in any open-system description a quantity can flow across the subsystem boundary while remaining conserved in the full system. Electroweak sphalerons supply the ordinary Standard-Model conversion,

$$\Delta(B + L) \neq 0, \qquad \Delta(B - L) = 0, \tag{169}$$

while the exact bulk baryonic charge stays globally protected. The framework can thus generate an *effective* boundary baryon excess without a late-time proton-decay instability.

(ii) CP violation as an inherited geometric phase. GG-A9-SM already derived a nontrivial leptonic Dirac phase $\delta_{\rm PMNS} \simeq 196^\circ$, the protected Majorana phases $(\alpha_{21}, \alpha_{31}) = (0, \pi)$, and the rephasing-invariant Jarlskog measure $J_{\rm PMNS} \simeq -9.2 \times 10^{-3}$. The compact-angle logic associates

CP-odd structure with holonomy and Wilson-line phase selection. The CP source need not be inserted by hand; it is inherited.

(iii) Out-of-equilibrium dynamics as a GDOS transition. A successful asymmetry cannot be generated in exact equilibrium. Within GDOS this is the ordinary fact that cosmological subsystems evolve under time-dependent ports, cooling, and changing attractors. The relevant stage is the period after reheating when the heavy neutrino sector ceases to track the bath and decays with a CP-odd bias — precisely what GDOS was built to describe.

The Sakharov conditions are therefore not external checkboxes but latent in the inherited ledger: $\mathrm{U}(1)_B$ forces indirect generation, the neutrino and compact-angle sectors supply the CP structure, and cosmological cooling supplies the open-system departure from equilibrium.

36.4 The inherited light-neutrino block and the heavy completion

The light-neutrino data of GG-A9-SM were already tabulated, but for baryogenesis it is worth displaying their implications. The completed particle paper yields

$$(m_1, m_2, m_3)_{\mathrm{QD}} = (0.0653,\ 8.5923,\ 50.5131)\ \mathrm{meV}, \qquad \Sigma m_\nu \simeq 59.17\ \mathrm{meV}, \tag{170}$$

with

$$\delta_{\mathrm{PMNS}} \simeq 196.06^\circ, \qquad (\alpha_{21}, \alpha_{31}) = (0, \pi), \qquad J_{\mathrm{PMNS}} \simeq -9.2 \times 10^{-3}. \tag{171}$$

These are normally ordered, Majorana-like, with order-one CP structure, and they descend from the same compact-angle and OGN arithmetic that organizes the rest of the paper. (The J_{PMNS} value is the first-order-angle layer; the NuFit-angle evaluation, the convention carried by GG-P5-CPT, gives -9.4×10^{-3}.) Crucially, in GG-A9-SM these light masses arise from the Weinberg dimension-five operator at the GUT seam,

$$\mathcal{L}_5 = \frac{(LH)(LH)}{\Lambda_{\mathrm{W}}}, \qquad m_\star = \frac{\pi v^2}{E_{\mathrm{GUT}}}, \tag{172}$$

with $v \simeq 246\,\mathrm{GeV}$, the convention at which GG-A9-SM evaluates $m_\star \simeq 8.4$ meV. GG-A9-SM fixes the light sector through Eq. (172) but *explicitly leaves the ultraviolet completion — the heavy Majorana states and the leptogenesis they drive — to GG-A10-Cosmos*. That completion is the subject of this section.

Before the completion is written down, the constitutional question must be answered: *is* a heavy Majorana singlet an assumption — the model-builder's standard import — or is it forced? In this geometry it is forced, by a one-line piece of fiber arithmetic that deserves theorem status, since the entire module rests on it.

Proposition 36.2 (The dial-free door: a Majorana mass exists iff the fiber weight vanishes (T-Cosmos-18 / R10))**.** *Let a fermion multiplet on the GG-6 bulk carry weight n under the compact phase direction S^1_θ (spinor and family indices commute with the fiber action). Then a Lorentz-invariant, gauge-invariant Majorana-type self-pairing — equivalently, a $\Delta L = 2$ bilinear — exists if and only if $n = 0$; and on every $n \neq 0$ sector the unique invariant symmetric self-pairing is the*

charge-conserving Dirac pairing $\psi\bar{\psi}$. Moreover the trilinear LHN is the unique fiber-neutral portal that touches the weight-zero sector linearly.

Proof. In the real presentation the weight-n sector is, pointwise, the real pair (a, b) on which the circle acts by the rotation $R(n\alpha)$. The symmetric square is three-dimensional, spanned by $a^2 + b^2$, $a^2 - b^2$, and $2ab$. The first is rotation-invariant for every n; it is $\psi\bar{\psi}$, the Dirac pairing. The remaining pair $(a^2 - b^2,\, 2ab)$ is exactly $(\mathrm{Re}\,\psi^2, \mathrm{Im}\,\psi^2)$ — the real form of the Majorana bilinear — and transforms as a weight-$2n$ doublet: under $R(n\alpha)$ it rotates by the angle $2n\alpha$. For $n \neq 0$, invariance under all α forces it to vanish; for $n = 0$ the action is trivial, the pairing is unobstructed, and the reality condition of the weight-zero sector is precisely the Majorana condition. The spinor factor contributes the usual charge-conjugation pairing, inert under the fiber, and cannot compensate a non-zero weight. For the portal: a trilinear is fiber-neutral iff its weights sum to zero; with L at weight $-\frac{1}{2}$ and H at $+\frac{1}{2}$ in the unified-U(1) normalization, the singlet N at weight 0 closes the unique such trilinear linear in N. ■

The proposition does three jobs at once. First, it geometrizes lepton number: charge conservation on the boundary is winding conservation on the circle, so L is exact throughout the dial-carrying ledger, and the sole legal $\Delta L = 2$ step in the entire matter sector is the Majorana mass of the one field the circle does not move. Second, it promotes the seesaw from choice to consequence: the type-I completion below is not one ultraviolet option among several but the unique fiber-neutral door between the conserved winding ledger and the weight-zero sector. Third, it closes the last imported ingredient of this paper's matter story: with the heavy sector now forced — its Majorana character by this proposition, its scale by the throat integer $K = 35$ — the chain from the constitution to the observed baryon share runs with no assumed particle content. The statement extends the real-bulk series of (Arisaka, 2026, A8-QM) (as its entry R10), where the weight-zero sector's real, dial-free character is developed in full; the standalone, full-length development of the series is the Foundations Branch member paper GG-P4-Real (Arisaka, 2026, P4-Real).

The natural completion of a Weinberg operator is the type-I seesaw: heavy Majorana neutrinos N_i whose exchange generates Eq. (172) on integrating out. In the present geometry they have only one natural home — the lock scale,

$$M_{N_i} = M_{\mathrm{lock}}\,(1 + \delta_i), \qquad \delta_i = O\!\left(\frac{1}{K}\right), \qquad K = 35, \tag{173}$$

so the triplet sits in the few-TeV corridor with small, geometrically controlled splittings. A minimal benchmark set is

$$(M_{N_1}, M_{N_2}, M_{N_3}) \sim M_{\mathrm{lock}}\,(1, 1 + K^{-1}, 1 + 2K^{-1}), \tag{174}$$

i.e. percent-level splittings around $M_{\mathrm{lock}} \simeq 3.85\,\mathrm{TeV}$ — the near-degenerate regime that resonant enhancement requires. The register of that last clause matters, so it is stated at GG-P5-CPT's own grade: the $O(1/K)$ splittings motivate *proximity* to the resonant regime, but at vacuum numbers the system sits far off the Breit–Wigner peak — the splitting $d \sim 1/K \simeq 3 \times 10^{-2}$ against the vacuum width $\Gamma_j/M_j \sim 2 \times 10^{-14}$ (Appendix I) — so the on-resonance condition is the chain's named conjecture, not its natural consequence, and "natural rather than tuned" would overstate

what vacuum numbers support.

36.5 Seesaw matching: the geometry fixes the Yukawas

The decisive quantitative step is that the seesaw fixes the neutrino Yukawa magnitude with no new freedom. Integrating out the lock-scale triplet gives the light mass matrix

$$m_\nu = \frac{Y_\nu^2\, v^2}{M_N}, \tag{175}$$

which must reproduce the GG-A9-SM Weinberg scale $m_\star = \pi v^2/E_{\rm GUT}$ of Eq. (172). Matching the two expressions cancels v^2 and fixes the Yukawa magnitude entirely in terms of inherited scales,

$$Y_\nu^2 \simeq \pi\,\frac{M_N}{E_{\rm GUT}} \simeq \pi\,\frac{3.85\times10^{3}}{2.28\times10^{16}} \simeq 5.3\times10^{-13}, \qquad Y_\nu \simeq 7.3\times10^{-7}. \tag{176}$$

The Yukawa is tiny because the heavy scale is low: a TeV-scale M_N with sub-eV light masses *requires* $Y_\nu \sim 10^{-6}$. This is not a problem to be tuned away but the direct consequence of placing the seesaw at the geometry's lock scale rather than at an arbitrary intermediate scale. The tree-level decay width that follows is correspondingly small,

$$\Gamma_N = \frac{Y_\nu^2}{8\pi} M_N \simeq 8\times10^{-11}\,\text{GeV}, \qquad \frac{\Gamma_N}{M_N} \simeq 2\times10^{-14}. \tag{177}$$

A tiny width would ordinarily doom leptogenesis. The resonant regime is what rescues it.

36.6 Resonant CP asymmetry from the inherited phases

The heavy Majorana neutrinos decay through

$$N_i \to \ell H, \qquad N_i \to \bar{\ell} H^\dagger, \tag{178}$$

with the CP-odd asymmetry

$$\epsilon_i \equiv \frac{\Gamma(N_i\to\ell H) - \Gamma(N_i\to\bar{\ell}H^\dagger)}{\Gamma(N_i\to\ell H) + \Gamma(N_i\to\bar{\ell}H^\dagger)}. \tag{179}$$

The standard self-energy (resonant) formula is

$$\epsilon_i = \frac{\Im\big[(Y_\nu Y_\nu^\dagger)^2_{ij}\big]}{(Y_\nu Y_\nu^\dagger)_{ii}(Y_\nu Y_\nu^\dagger)_{jj}}\,\frac{(M_i^2-M_j^2)M_i\Gamma_j}{(M_i^2-M_j^2)^2+M_i^2\Gamma_j^2}, \qquad i\neq j. \tag{180}$$

The two factors are a CP-odd invariant and a resonant Breit–Wigner enhancement. Both are narrowed by the geometry. The enhancement is maximal at the resonance condition $|M_i - M_j| \simeq \Gamma_j/2$, where the Breit–Wigner factor reaches its peak value $\frac{1}{2}$ and, decisively, the tiny Yukawa magnitude *cancels between numerator and denominator of the prefactor*. At resonance the asymmetry is

therefore controlled not by Y_ν but by the geometric phase alone,

$$\epsilon^{\rm res} \simeq \frac{1}{2}\sin(2\,\phi_{\rm geo}), \quad \phi_{\rm geo} \text{ from the flavored projection of the inherited phases (Appendix I).} \tag{181}$$

This is why a TeV-scale, tiny-Yukawa seesaw can still produce a viable asymmetry: the resonance lifts the would-be Y_ν^2 suppression, leaving a phase-controlled number. At the TeV lock scale the chain is fully flavored (the flavor obligation L-Cosmos-F of Lemma I.1, Appendix I), and the operative CP object is the flavored asymmetry $\epsilon_{i\alpha}$, computed on the inherited phases in Appendix I: the driver is the Majorana parity $\alpha_{31} = \pi$, which sets the sign of the matter excess, and the chain carries

$$\boxed{\epsilon^{\rm res} \simeq 2.5 \times 10^{-6}} \tag{182}$$

as the target value that the partial-resonance hypothesis H-$\phi_{\rm geo}$-1 ($\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$; Appendix I) must deliver — derived-conditional, not a derived output: on the exact Breit–Wigner peak the flavored computation overshoots the required throughput by $\sim 3 \times 10^4$, and at vacuum splittings it undershoots by $\sim 4 \times 10^6$. The flavored computation and the hypothesis are worked in Appendix I, where the hypothesis now carries an identified candidate mechanism at candidate grade — the flavor-measurement regulator, with the zero-dial closed form $\Delta M/\Gamma^{\rm eff} = 1/(35\,c_\tau y_\tau^2)$ — whose conditions C1–C6 remain open.

36.7 Washout and the strong-washout efficiency

Whether this asymmetry survives is decided by washout. The relevant control parameter is the effective neutrino mass that the decaying state couples to,

$$\widetilde{m}_i = \frac{(Y_\nu Y_\nu^\dagger)_{ii}\, v^2}{M_i}, \tag{183}$$

compared with the equilibrium scale

$$m_\star = \frac{8\pi v^2}{M_i}\,\frac{H(M_i)}{M_i} \simeq 1.1 \times 10^{-3}\,\text{eV}. \tag{184}$$

The inherited light spectrum makes the relevant effective mass of order the heaviest light neutrino, $\widetilde{m} \sim m_3 \simeq 5 \times 10^{-2}\,\text{eV}$, so

$$K_{\rm wash} \equiv \frac{\widetilde{m}}{m_\star} \simeq 50 \gg 1, \tag{185}$$

i.e. strong washout. In the strong-washout regime the surviving fraction is the standard suppressed efficiency,

$$\kappa \simeq \frac{1}{K_{\rm wash}} \text{ (up to an } \mathcal{O}(1) \text{ factor)} \simeq 1 \times 10^{-2}. \tag{186}$$

Strong washout is, paradoxically, a stabilizing feature: it makes the final asymmetry largely independent of initial conditions, which is precisely the regime in which an inherited (rather than fitted) calculation is trustworthy. The Boltzmann estimate behind Eq. (186) is given in Appendix I. The washout here runs at the bath scale $\widetilde{m} \sim m_3$, while the resonance width of the seesaw completion

runs at the seam-matched $\widetilde{m} \simeq 8.4$ meV — two deliberate conventions whose factor-$\simeq 12$ spread in Y_ν^2 is declared, accounted for (the single-convention ratio $2\,m_3/\widetilde{m}_{\rm seam} \simeq 12$, the factor of two a normalization bookkeeping), and priced within the $O(10)$ register in the convention note at Eq. (611) of Appendix I.

36.8 Assembling η_B: the mechanism behind $e^{-K/2}$

The net CP-odd throughput that reaches the boundary is the product of the resonant asymmetry and the washout efficiency,

$$\epsilon_{\rm geo}^{\rm eff} \equiv \epsilon^{\rm res}\,\kappa \simeq (2.5\times10^{-6})\times(1.0\times10^{-2}) \simeq 2.51\times10^{-8}. \tag{187}$$

The arresting fact is that this product reconstructs, to the displayed precision, the throat-depth exponential that the geometry independently predicts,

$$\epsilon_{\rm geo}^{\rm eff} \simeq e^{-K/2} = e^{-17.5} \simeq 2.51\times10^{-8}. \tag{188}$$

This is the central result of the section. The compressed ledger value $e^{-K/2}$ is not a separate assertion: it is the *coarse-grained output of the resonant-leptogenesis transport*, the product of a phase-controlled resonant asymmetry ($\sim 10^{-6}$) and a strong-washout efficiency ($\sim 10^{-2}$). The transport mechanism is the physics behind the number. (The exponent $K/2$ rather than K or $K/4$ is the compression statement the Resonant-leptogenesis-normalization conjecture (Conjecture VI) must ultimately fix from first principles; what is new here is that the independently derived washout efficiency and the candidate regulator's closed form bracket the required throughput — the identity at displayed precision uses the $\mathcal{O}(1)$ freedom in κ and the target status of $\epsilon^{\rm res}$.)

The remainder is standard transport bookkeeping, and one bookkeeping factor deserves display rather than absorption: the asymmetry is generated on the heavy sector's *equilibrium abundance*. A relativistic Majorana fermion with $g_N = 2$ spin states has number density $n_N^{\rm eq} = \frac{3}{4}(\zeta(3)/\pi^2)\,g_N T^3$, the bath's entropy density is $s = (2\pi^2/45)\,g_{*s}T^3$, and their ratio is the equilibrium yield

$$Y_N^{\rm eq} \equiv \frac{n_N^{\rm eq}}{s} = \frac{135\,\zeta(3)}{8\pi^4}\,\frac{g_N}{g_{*s}} \simeq \frac{0.42}{g_*} \ \text{ per species}, \qquad \sum_{i:\,m_i>0} Y_{N_i}^{\rm eq} \simeq \frac{0.833}{g_*} \tag{189}$$

for the two sourcing states — physically, a species in equilibrium carries its fair share of the bath's heat, two spin states' worth per member. The spectrum has three members, $(m_1, m_2, m_3) = (0, 1, 2N)\,m_*$ with $R = 1$ (GG-A9-SM), but the first carries no Yukawa, hence no width and no asymmetry: it is in the spectrum and not in the source, and the sum runs over the two that are. Each net decay deposits the CP-odd fraction $\epsilon_{\rm geo}^{\rm eff} = \epsilon^{\rm res}\kappa$ of that abundance into lepton number, so the comoving lepton asymmetry per entropy is

$$Y_L \simeq \Big(\sum_i Y_{N_i}^{\rm eq}\Big)\,\epsilon_{\rm geo}^{\rm eff} = \frac{0.833\,\epsilon_{\rm geo}^{\rm eff}}{g_*}, \qquad g_* \simeq 106.75, \tag{190}$$

for the Standard-Model-like bath at the TeV scale. (The frequently quoted shorthand $Y_L \sim \epsilon/g_*$

absorbs the equilibrium-yield factor silently; displayed, it is worth a quarter of the final answer.) Electroweak sphalerons convert part of this to baryon number with the one-Higgs-doublet coefficient

$$c_{\rm sph} = \frac{28}{79} \simeq 0.3544, \tag{191}$$

and the entropy-normalized yield converts to the baryon-to-photon ratio through the standard late-time factor 7.04:

$$\eta_B \simeq 7.04\, c_{\rm sph}\, Y_L \simeq 7.04 \left(\frac{28}{79}\right) \frac{0.833\, \epsilon^{\rm res}\, \kappa}{g_*} = 7.04 \left(\frac{28}{79}\right) \frac{0.833\, e^{-K/2}}{g_*}. \tag{192}$$

Inserting the inherited numbers,

$$\eta_B \simeq 7.04 \times 0.354430 \times 0.833 \times \frac{2.511 \times 10^{-8}}{106.75} \simeq 4.89 \times 10^{-10}. \tag{193}$$

This sits within 20% of the observed baryon asymmetry of Eq. (168), well inside the declared order-of-magnitude register. Nothing in the chain was freely fitted: the light sector and phases come from GG-A9-SM, the heavy scale $M_N \simeq M_{\rm lock}$ and the near-degeneracy $\delta \sim 1/K$ come from the geometry, and the seesaw matching of Eq. (176) fixes the Yukawas. The one named input is the *partial*-resonance assumption $\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$ at finite temperature — not the exact-peak condition $|M_i - M_j| \simeq \Gamma_j/2$, which overshoots the required throughput by that same factor and is *not* what is assumed here. The OGN splitting $\delta_i \sim 1/K$ makes proximity to the regime plausible and does not by itself enforce it: at vacuum numbers GG-A9-SM puts $\Delta M/\Gamma_N \simeq 6.8 \times 10^{11}$, so the assumption is a conjecture (Lemma VI) and the finite-temperature width that would meet it is the open computation named below. The full transport calculation, including the explicit seesaw matching, the resonant CP-asymmetry projection, and the flavored Boltzmann integration, is carried out in Appendix I; the reader who wishes to see every factor is encouraged to consult it.

Against the CMB-inferred $\eta_B^{\rm obs} = (6.12 \pm 0.04) \times 10^{-10}$ (Planck Collaboration, 2018) the derived value 4.89×10^{-10} sits at -20%. The honest tier is order-of-magnitude with a controlled mechanism: the resonant condition and the $\mathcal{O}(1)$ factor inside the strong-washout efficiency are not yet computed to percent precision, so the entry carries the factor-$\sim$3 theory band shown in the parameter-audit Dev plot (Figure 31, Table 41). Displaying the equilibrium-yield factor of Eq. (189) matters for honesty in both directions: absorbing it silently into the $\mathcal{O}(1)$ register would place the central value within a few percent of the observation — a closeness that would be an artifact of the undisplayed factor. Closeness of central values at this tier is never a precision success; the genuine claim is that an explicit, inherited resonant-leptogenesis chain reconstructs the geometric $e^{-K/2}$ ledger value and lands within its declared band of the sky.

36.9 Boltzmann transport and why the geometry compresses the calculation

It is worth showing how Eq. (193) sits inside the familiar Boltzmann language. With the yields

$$Y_{N_i} \equiv \frac{n_{N_i}}{s}, \qquad Y_{B-L} \equiv \frac{n_{B-L}}{s}, \qquad z \equiv \frac{M_{\rm lock}}{T}, \tag{194}$$

the resonant-leptogenesis system is

$$\frac{\mathrm{d}Y_{N_i}}{\mathrm{d}z} = -D_i(z)(Y_{N_i} - Y_{N_i}^{\mathrm{eq}}), \tag{195}$$

$$\frac{\mathrm{d}Y_{B-L}}{\mathrm{d}z} = \sum_i \epsilon_i D_i(z)(Y_{N_i} - Y_{N_i}^{\mathrm{eq}}) - W(z)Y_{B-L}, \tag{196}$$

with D_i the decay terms and W the washout term. In an unconstrained model both the ϵ_i and the washout profile $W(z)$ range over a large Yukawa-and-mass space. Here the heavy masses are tied to the lock scale, the splittings to the OGN integer, the Yukawas to the seesaw matching, and the CP data to the inherited phases. The system therefore collapses to one effective dimensionless throughput. A useful factorization is

$$\epsilon_{\mathrm{geo}}^{\mathrm{eff}} = \mathcal{R}_{\mathrm{res}}\, \mathcal{W}_{\mathrm{eff}}\, \sin\delta_{\mathrm{geo}} = \epsilon^{\mathrm{res}}\, \kappa, \tag{197}$$

where $\mathcal{R}_{\mathrm{res}}$ is the resonant enhancement, $\mathcal{W}_{\mathrm{eff}} = \kappa$ the survival after washout, and $\sin\delta_{\mathrm{geo}}$ the holonomy-controlled phase — the same three factors made numerical in Eqs. (182)–(187). The geometry locks their product to $e^{-K/2}$.

The timing is consistent. At $T \sim M_{\mathrm{lock}}$ the age of the universe is

$$t \simeq \frac{0.30\, M_{\mathrm{Pl}}}{g_*^{1/2} T^2} \simeq 1.6 \times 10^{-14}\,\mathrm{s} \quad (g_* = 106.75), \tag{198}$$

late enough for a valid boundary FRW description but early enough that the kinetics are primordial. A few-TeV channel needs no extreme ultraviolet prehistory, and the resonant regime — rather than a far-higher-scale seesaw — is what the tiny inherited light masses point to.

36.9.1 A numerical audit table for the baryogenesis chain

Table 16: Resonant-leptogenesis chain used in Part V. Each row is fixed by an inherited input; the one named assumption is the on-resonance condition, in its partial-resonance sharp form (Appendix I). The columns give each ingredient, its value, and where in the framework it comes from. The third column is what makes the chain checkable: an entry sourced to an upstream paper carries that paper's grade, and the single named assumption is marked as one rather than absorbed into the others.

Ingredient	Value	Origin in the framework
Lock scale / heavy mass	$M_N \simeq M_{\rm lock} \simeq 3.85\,\text{TeV}$	OGN-5 throat-depth reduction from $E_{\rm Pl}$
Seesaw Yukawa	$Y_\nu \simeq 7.3 \times 10^{-7}$	Matching to GG-A9-SM Weinberg scale, Eq. (176)
Heavy-neutrino splittings	$\delta_i = O(1/K)$	Quasi-degenerate lock-scale triplet
Resonant CP asymmetry	$\epsilon^{\rm res} \simeq 2.5 \times 10^{-6}$	Flavored projection of inherited phases; H-$\phi_{\rm geo}$-1 target value, Eq. (182)
Strong-washout efficiency	$\kappa \simeq 1 \times 10^{-2}$	$K_{\rm wash} \simeq 50$, Eq. (186)
Effective CP throughput	$\epsilon^{\rm res}\kappa \simeq 2.51 \times 10^{-8} = e^{-K/2}$	Product reconstructs the geometric ledger value
Relativistic d.o.f.	$g_* \simeq 106.75$	Standard SM-like bath at the TeV scale
Sphaleron factor	28/79	Standard one-Higgs-doublet conversion
Equilibrium yield	$\sum_i Y_{N_i}^{\rm eq} \simeq 0.833/g_*$	Two sourcing Majorana states of the three (the massless first has no Yukawa), Eq. (189)
Entropy-to-photon factor	7.04	Standard late-time yield-to-η_B conversion
Final asymmetry	$\eta_B \simeq 4.89 \times 10^{-10}$	Eq. (193)
Present baryon fraction	$\Omega_b = 1/(2\pi^2) = 0.0507$	π-locked inventory target

Once the lock scale and throat-depth integer are inherited, the seesaw matching fixes the Yukawas, and the route is fixed to leptogenesis, the asymmetry estimate becomes rigid: there is very little hidden freedom left.

36.10 From baryon asymmetry to the present baryon density fraction

Cosmology also summarizes the late baryon sector through the density fraction. Within the π-locked inventory,

$$\Omega_b = \frac{1}{2\pi^2} \simeq 0.0507. \tag{199}$$

This compact number gives the BBN and CMB sectors a sharp benchmark and makes clear that the baryonic share is a structured sub-share of the matter share, not comparable to it. The relation

between the early asymmetry and the late density runs through

$$\eta_B \Longrightarrow n_B(a) \Longrightarrow \rho_b(a) = m_p n_B(a) \Longrightarrow \Omega_b = \frac{\rho_b(t_0)}{\rho_{\rm crit}(t_0)}, \tag{200}$$

ordinary evolution once the asymmetry is fixed. Equation (193) says the initial asymmetry sits at the correct level.

36.11 Why exact proton stability and successful baryogenesis are compatible

If $\mathrm{U}(1)_B$ is exact, how can the universe acquire a baryon asymmetry? Exact $\mathrm{U}(1)_B$ refers to the fundamental conserved baryonic ledger of the full GG-6 system, whereas cosmological baryogenesis is a statement about the projected four-dimensional open subsystem. The relevant processes are not local proton-decay operators but boundary processes — a lepton asymmetry, sphaleron conversion, open-system transport. The full bulk ledger remains consistent while the observable subsystem relaxes into a state with nonzero net baryon density. This is exactly the situation OGGEP and GDOS were introduced for: a conserved quantity in the full system can present as an evolving charge budget in an open subsystem when ports and projection are present. There is no contradiction. Proton decay remains a sharp falsifier; baryogenesis remains available through the constrained leptogenesis channel rather than through late-time baryon-violating operators — more economical than GUT scenarios in which the baryon-creating sector also threatens the proton.

36.12 What the baryogenesis block predicts and how it can fail

1. The heavy-neutrino sector sits near the lock scale, $M_N \simeq M_{\rm lock}$, not at an arbitrary intermediate scale.

2. The seesaw matching to the GG-A9-SM Weinberg scale fixes $Y_\nu \simeq 7 \times 10^{-7}$; the light block and its CP structure are compatible with a viable resonant-leptogenesis descendant.

3. The net throughput is $\epsilon^{\rm res}\kappa \simeq e^{-K/2}$ with $K = 35$: a phase-controlled resonant asymmetry times a strong-washout efficiency.

4. The resulting asymmetry is $\eta_B \approx 5 \times 10^{-10}$ (central 4.89; band $\times 3^{\pm 1}$: (1.6–14.7) $\times 10^{-10}$; the time-resolved integration lands a factor of about ten below the ledger road's 6.18×10^{-10}, §I.6) and the late baryon fraction is $1/(2\pi^2)$.

Every element can fail. If a full transport calculation finds no finite-temperature regulator in the partial-resonance decade $\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$ near the lock scale — the exact-peak regime is the one the flavored recomputation excludes — the chain fails. If cosmological data force a neutrino mass sum or CP structure incompatible with the particle paper, the backbone fails. If proton decay is observed, the whole baryonic logic fails. The section is acutely exposed to observation.

36.13 Exact neutrality: why the books balance and the electrons match the protons

The chain just assembled produces asymmetries; this subsection records why it can never produce the wrong kind. The statement is a theorem of the inherited structure: *every gauged charge has exactly zero density at every epoch; asymmetry is possible only in global numbers.*

Four links. *First, the start is neutral.* The edge data are geometric and the geometry is a gauge singlet: a nonzero gauge charge on $\Sigma^{\mathrm{Bdy}}_{\mathrm{GUT}}$ would be an extra initial datum that the inheritance clause of §18 does not supply. Independently, at the edge the boundary group is the unbroken unified group, and a generator of a simple group is traceless on every multiplet (Arisaka, 2026, A9-SM): the same trace identity that makes the hydrogen atom neutral makes any democratic population of complete multiplets neutral. One identity, two epochs — the neutrality of the first plasma and of the last atom are the same traceless matrix, read at $r = 0$ and at recombination.

Second, the zero is conserved. The covariant conservation identity holds the boundary Ward currents conserved as one bulk identity down the entire descent (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A2-GG6), and its microscopic face is that every vertex conserves charge: charged quanta enter the world only in zero-net-charge combinations, the electron–positron pair being the minimal case.

Third, the zero is enforced, not merely remembered. A gauged charge sources its own long-range field, so it cannot hide. An exactly homogeneous and isotropic background admits no field strength, and the Maxwell equation then forces the charge density to vanish: a charged universe cannot be FRW — as was recognized the first time a charged cosmology was seriously proposed (Lyttleton & Bondi, 1959). Dynamically the enforcement is faster than any cosmological clock: a local excess relaxes at the plasma frequency, some twenty orders of magnitude above the expansion rate during nucleosynthesis. And even a maximal primordial excess would have been diluted by $e^{-3N^*} \sim 10^{-69}$ during the launch — the theorem is protected by symmetry and, independently, by dilution.

Fourth, the zero decides who survives. $B-L$ is global; no field polices it; the lock-scale block may load it, and does. The sphaleron bookkeeping that converts it runs on the plasma's chemical equilibrium, in which the hypercharge-neutrality condition — so often imported silently — is simply the third link applied to the gauged hypercharge above the crossover. After the electroweak epoch the surviving charged species are the baryons and the electrons, and the constraint $\sum_i Z_i n_i = 0$, enforced continuously, pins the electron density to the proton charge as the positrons annihilate away:

$$n_e \;=\; n_p + 2\, n_{\mathrm{He}} + \cdots \tag{201}$$

exactly — the relation every Saha and Peebles equation of the recombination chain silently uses, here no longer silent. The electron excess over the positrons is not separately arranged; it is the image of the baryon excess under the neutrality constraint. Empirically the exactness is audited at one part in 10^{21} by the neutrality of bulk matter (Bressi et al., 2011).

The register of the result matters: neutral in every charge that is gauged, asymmetric only in charges that are global — neutral where a field would betray the imbalance, asymmetric only where nothing watches. Both halves are forced, and the third horizon problem of Part X (§73) is the same coin seen from the sky.

37 Big Bang nucleosynthesis and the lithium-7 audit

37.1 Historical background

The hot-Big-Bang nucleosynthesis program was inaugurated in 1948 by Alpher, Bethe, and Gamow (Alpher et al., 1948), who proposed that the light elements are produced by neutron-proton interactions in a cooling, expanding plasma. Modern BBN predicts the primordial abundances of D, ^{3}He, ^{4}He, and ^{7}Li from a small handful of inputs — the baryon-to-photon ratio, the effective neutrino number, and the standard thermal history — and agrees with observation to better than ten percent for the first three. The residual lithium-7 discrepancy is the only persistent anomaly. The present section asks whether the inherited ledger modifies any BBN input in a way that addresses lithium-7 without disturbing the other three.

37.2 Why BBN is the first precision audit of Part V

Cosmology does not permit stopping at attractive narratives. At temperatures of order the MeV and times of order seconds to minutes, the universe records its matter and radiation ledger in a way still observable today. BBN is therefore the first hard quantitative audit of the whole origin-of-matter story. Its dependence on the input set is broad yet clean: the expansion rate, the baryon density, the neutrino/radiation content, the weak freeze-out history, and the nuclear rates. It is close enough to fundamental physics to test the inherited ledger directly, but late enough that the boundary FRW description already works well. The empirical situation is instructive: standard BBN agrees remarkably with primordial ^{4}He and deuterium, while the long-standing ^{7}Li problem persists. A successful framework must preserve what works and change only what is under pressure.

37.3 Standard BBN baseline: freeze-out, bottleneck, and synthesis

Above about an MeV, neutrons and protons interconvert rapidly,

$$n + \nu_e \leftrightarrow p + e^-, \qquad n + e^+ \leftrightarrow p + \bar{\nu}_e, \qquad n \leftrightarrow p + e^- + \bar{\nu}_e, \tag{202}$$

so the neutron-to-proton ratio is approximately thermal,

$$\frac{n}{p} \simeq e^{-\Delta m/T}, \qquad \Delta m \equiv m_n - m_p \simeq 1.293\,\mathrm{MeV}. \tag{203}$$

As the weak rate falls below Hubble, the ratio freezes out near one neutron to several protons; after a little free decay these neutrons are incorporated into ^{4}He once the network ignites. The delay is the deuterium bottleneck: the intense photon bath photodissociates deuterons until

$$T_{\mathrm{D}} \sim 0.07\text{–}0.09\,\mathrm{MeV}, \tag{204}$$

after which the chain

$$p + n \to D + \gamma, \quad D + p \to {}^3\mathrm{He} + \gamma, \quad D + n \to {}^3\mathrm{H} + \gamma, \quad {}^3\mathrm{He} + {}^3\mathrm{H} \to {}^4\mathrm{He} + p + n, \tag{205}$$

processes most neutrons into ^{4}He. The standard yields are

$$Y_p \equiv \frac{4n_{^4\mathrm{He}}}{n_b} \sim 0.245, \qquad D/H \sim 2.5\times10^{-5}, \qquad {}^7\mathrm{Li}/H \sim \text{few}\times10^{-10}. \tag{206}$$

The general evolution

$$\dot{Y}_i = \sum_r N_{ir}\,\lambda_r(T)\prod_j Y_j^{\nu_{jr}}, \tag{207}$$

with $Y_i = n_i/n_b$, N_{ir} the net stoichiometric coefficient, and $\lambda_r(T)$ the rates, is the starting point for the GDOS extension.

37.3.1 Freeze-out and bottleneck in explicit numbers

The weak rate scales as $\Gamma_{n\leftrightarrow p} \sim G_F^2 T^5$; setting it equal to Hubble, Eq. (159), gives a freeze-out

$$T_f \sim 0.7\text{–}0.8\,\mathrm{MeV}, \qquad \left(\frac{n}{p}\right)_f \sim e^{-1.293/0.75} \sim 0.18, \tag{208}$$

the familiar "one neutron per five or six protons." With nearly all surviving neutrons going into ^{4}He,

$$Y_p \simeq \frac{2(n/p)}{1+n/p}, \tag{209}$$

so $(n/p)\sim 1/7$ gives $Y_p \sim 0.25$. The deuterium bottleneck breaks only at the 0.07–0.09 MeV corridor because the enormous photon-to-baryon ratio keeps the high-energy tail photodissociating deuterons long after the average photon energy drops below the 2.2 MeV binding. These two estimates show why BBN is robust yet selective: the helium fraction is set long before the A=7 channel matters, deuterium tracks the baryon density, and the A=7 abundance is tiny and late enough that a selective correction can modify it without disturbing the rest.

37.3.2 Why deuterium is such a powerful baryometer

Deuterium is easily burned in stars and made primarily in the early universe, so its primordial abundance is a clean probe of the baryon density: more baryons burn more efficiently and drive D/H down. That is why D/H and the CMB point to nearly the same baryon density. Any Part V proposal altering the baryon ledger too much would fail immediately. The Kaxion window is safe because it modifies the nonbaryonic share; leptogenesis is safe because it sets the asymmetry rather than adding a late baryon correction; the lithium port is safe only because it is selective and tiny in total transferred abundance. All three survive only by respecting deuterium.

37.4 The lithium-7 problem and the A=7 channel

Standard BBN is not broken; its deuterium and helium predictions are among the great successes of precision early-universe physics. That success constrains any new proposal: a ^{7}Li mechanism must leave Y_p and D/H intact. The lithium problem enters through the A=7 channel. At the

CMB/deuterium baryon density, the dominant mass-7 route is not primordial ^{7}Li but ^{7}Be formation,

$$^3\mathrm{He} + \alpha \to {}^7\mathrm{Be} + \gamma, \tag{210}$$

followed much later by electron capture,

$$^7\mathrm{Be} + e^- \to {}^7\mathrm{Li} + \nu_e, \tag{211}$$

whose product is quoted as the primordial lithium abundance. Standard BBN gives

$$\left(\frac{^7\mathrm{Li}}{H}\right)_{\mathrm{SBBN}} \simeq 5.0 \times 10^{-10} \tag{212}$$

(Cyburt *et al.*, 2016; Pitrou *et al.*, 2018), whereas the metal-poor halo-star (Spite) plateau gives

$$\left(\frac{^7\mathrm{Li}}{H}\right)_{\mathrm{obs}} \simeq (1.6 \pm 0.3) \times 10^{-10} \tag{213}$$

(Fields, 2011). The discrepancy is a factor of about three. Many explanations have been proposed — stellar depletion, new-particle injection, rate modifications. The narrower question here is whether the GG-6/GDOS framework contains a *minimal, inherited, selective* mechanism acting on the A=7 channel while leaving deuterium and helium intact.

37.5 The A=7 port and the $1/\pi$ prediction

The standard treatment of the lithium anomaly leaves the size of any A=7 suppression as a free parameter to be fitted to the observed deficit. The present framework does not fit it. The same closed-packet branching that fixes the cosmic matter inventory fixes the A=7 survival fraction, and it does so with no new freedom.

Recall the topological origin of the inventory (Conjecture VIII / the closed-packet derivation): when a conserved flux is distributed over the compact angular phase, the fraction that survives onto the observable boundary is set by the geometry of the closed packet over S^1_θ, and the surviving share is $1/\pi$. This is the identical π-locking that yields the matter share $\Omega_m = 1/\pi$. The claim of this section is that the A=7 channel is an open-system port of exactly this type: the mass-7 flux carried by ^{7}Be is branched over the same compact-angle phase during its open electron-capture / phase-noise window, and the surviving fraction is the same geometric factor,

$$S_{A=7} = \frac{1}{\pi}. \tag{214}$$

The A=7 yield is therefore suppressed by a *predicted* factor of π, not a fitted one. The prediction is

$$\left(\frac{^7\mathrm{Li}}{H}\right)_{\mathrm{GG}} = \frac{1}{\pi}\left(\frac{^7\mathrm{Li}}{H}\right)_{\mathrm{SBBN}} \simeq \frac{5.0 \times 10^{-10}}{3.14159} \simeq 1.6 \times 10^{-10}. \tag{215}$$

This matches the observed plateau, Eq. (213), with no parameter adjusted: the identical $1/\pi$ that

fixes Ω_m fixes the lithium deficit. The closed-packet network analysis behind Eq. (214) — the open-system A=7 port and the explicit phase-branching that produces the surviving fraction $1/\pi$ — is given in Appendix I.

The status is therefore D* (Derived-conditional), sharing the same closed-packet origin as the inventory: the $1/\pi$ is derived modulo the closed-packet-integration conjecture, not fitted. The prediction is crisp and falsifiable: *any firm primordial* ^{7}*Li abundance away from* SBBN$/\pi$ *falsifies the port.* In particular, if improved analyses placed the primordial value at the full SBBN level, or at a value not equal to SBBN$/\pi$ within errors, the mechanism would be excluded.

37.6 Why deuterium and helium are preserved

The selectivity is the whole point; two arguments show it.

(i) Temporal hierarchy. The ^{4}He yield is fixed much earlier, by the frozen neutron-to-proton ratio and the end of the deuterium bottleneck. By the time the A=7 port acts, the helium mass fraction is already a locked attractor, so the lithium correction does not feed back into Y_p.

(ii) Extreme abundance hierarchy. The A=7 abundance is tiny compared with deuterium and helium,

$$Y_{\rm Be7} \sim 10^{-10}, \qquad D/H \sim 2.5\times10^{-5}, \qquad Y_p \sim 0.25, \tag{216}$$

so even if every removed ^{7}Be nucleus fed a side channel, the maximum fractional shift in D/H is

$$\frac{\Delta(D/H)}{D/H} \lesssim \frac{7\times10^{-10}}{2.5\times10^{-5}} \sim 3\times10^{-5}, \tag{217}$$

far below the precision at which deuterium is used, and the effect on Y_p is negligible. The port acts where the standard theory is under pressure and is harmless where it succeeds.

Table 17: Light-element audit used in Part V. The first two data columns summarize the standard baseline and the observed corridor; the last gives the inherited GG-6/GDOS effect. The A=7 entry is a $1/\pi$ prediction, not a fit. The last column is where this framework enters; the first three are the standard baseline and the observed corridor, set down so the size of the proposed effect can be read against them directly.

Observable	Standard BBN	Observed corridor	Part V effect
Y_p	~ 0.245	close to standard	Essentially unchanged
D/H	$\sim 2.5\times10^{-5}$	close to standard	Essentially unchanged
^{3}He$/H$	small	not dramatically discrepant	Essentially unchanged
^{7}Li$/H$	$\simeq 5.0\times10^{-10}$	$(1.6\pm0.3)\times10^{-10}$	SBBN$/\pi \simeq 1.6\times10^{-10}$ (predicted)

The light-element sector is not "improved everywhere." It is preserved where it already succeeds

and corrected only where the standard baseline is under pressure — and there the correction is the inherited factor π, the same number that fixes the cosmic matter inventory.

37.7 What the lithium audit predicts and how it can fail

1. The inherited baryon density stays close to the CMB/deuterium-preferred value, preserving the successful ^{4}He and D/H predictions.

2. The A=7 open-system port suppresses the mass-7 yield by exactly $1/\pi$, giving $(^7\mathrm{Li}/H)_{\mathrm{GG}} \simeq 1.6 \times 10^{-10}$ — a prediction matching the observed plateau, with no fitted strength.

3. Because the affected abundance is tiny, the same correction does not disturb the more successful sectors.

The claims are sharply falsifiable. If improved analyses show the deficit is mostly stellar rather than primordial, the port is unnecessary. If the primordial value is firmly anything other than SBBN/π, the port is excluded. Either way the section proposes a specific, inherited mechanism rather than a vague hope — and it stakes the same $1/\pi$ that the rest of the program relies on.

37.8 A consistency chain from η_B to Ω_b and the baryon-to-dark-matter ratio

One quantity, three sets of units. Before the chain is assembled, one point of bookkeeping governs how the rest of this subsection should be read. (The dedicated leptogenesis companion (Arisaka, 2026, C4-Leptogenesis) carries the flavored transport computation behind the chain; the count and the numbers here follow GG-A9-SM.) The baryon-to-photon ratio η_B, the baryon density ρ_b and the baryon fraction Ω_b are not three predictions. They are *one* physical quantity in three sets of units, tied by conversions that contain no content of this framework's:

$$\eta_{10} \;=\; 273.9\,\Omega_b h^2, \qquad \rho_b \;=\; \Omega_b h^2 \times 1.87847 \times 10^{-29}\ \mathrm{g\,cm^{-3}}, \qquad \eta_{10} \equiv \eta_B \times 10^{10}. \tag{218}$$

Fix any one of the three and the other two follow with no freedom. What this paper has is therefore not three baryon predictions with three error bars, but *one road that predicts the number and one mechanism that is checked against it* — and they should be quoted in that order.

The ledger road first. $\Omega_b = 1/(2\pi^2)$ converted through Eq. (218) at the structural rate of Route I, $h = 0.6676$ (Section 29.8; zero sky input), gives $\eta_{10} = 6.18$ against an observed 6.12: an error of $+1.1\%$, with nothing fitted anywhere in the chain. At the density road's $h = 0.666$ — which takes the measured ω_b and ω_m as inputs and is therefore not independent of the datum — it gives 6.15 ($+0.57\%$); at the microwave background's own $h = 0.674$, 6.30 ($+3.0\%$). The honest range for this road is one to three percent, and $\eta_{10} = 6.18$ is the number this paper predicts.

The transport mechanism second, and as what it is. The chain assembled below lands at $\eta_{10} \simeq 4.89$: 79% of the geometric target 6.18, and -20.1% against the observation, inside a declared factor-$\sim$3 band, graded D* on a named finite-temperature hypothesis. Its claim is therefore *that a parameter-free mechanism reaches this order of magnitude*, not that the value is 4.89 — which is exactly what the codebook grading already says. Read that way the two are not rivals of equal

standing: the static count is this paper's prediction of the number, and the dynamical chain is the demonstration that a mechanism exists which could have produced it. The comparison that matters for the mechanism is with the target it must reach, 6.18, not with the datum; the closed form's shortfall of 21%, and the time-resolved integration's factor of about ten (§I.6), are the mechanism's, and they are what the finite-temperature hypothesis has to close.

It is worth asking whether the three modules speak to one another quantitatively. The baryogenesis subsection derived an asymmetry in the right corridor, Eq. (193); the inventory proposes a geometric baryon fraction, Eq. (199). These are not the same statement — one microphysical, one a late-time density share — and one should show how they talk. The present photon number density is

$$n_{\gamma,0} = \frac{2\zeta(3)}{\pi^2}\, T_0^3 \simeq 410\ \mathrm{cm}^{-3}, \tag{219}$$

so an asymmetry η_B implies

$$\rho_{b,0} = \eta_B\, m_p\, n_{\gamma,0} \simeq 3.3 \times 10^{-31} \left(\frac{\eta_B}{4.89 \times 10^{-10}}\right) \mathrm{g\,cm}^{-3}. \tag{220}$$

With the critical density $\rho_{\mathrm{crit},0} = 1.878 \times 10^{-29}\, h^2\ \mathrm{g\,cm}^{-3}$,

$$\Omega_b h^2 \simeq 0.0179 \left(\frac{\eta_B}{4.89 \times 10^{-10}}\right), \tag{221}$$

so at the carried rate $h = 0.6676$, $\Omega_b \simeq 0.040$ — on the two-state count the transport route sits 21% *below* the geometric target $1/(2\pi^2) \simeq 0.0507$ (22% at the microwave background's own $h = 0.674$), where on the earlier three-state count it sat above it. Stated honestly, the two routes do not yet exactly agree: the topological route gives $\Omega_b = 1/(2\pi^2) = 0.05066$, the asymmetry route $\Omega_b h^2 \simeq 0.0179$ (i.e. $\Omega_b \approx 0.0402$ at $h = 0.6676$) — a -21% gap on the two-state count, with the asymmetry route additionally carrying the transport chain's factor-$\sim$3 band, inside which the topological value sits comfortably. The residue is not symmetric between the routes: anchored to the measured $\Omega_b h^2 = 0.02237$, the transport route's -20% is priced by its own declared band, while the topological route at the carried rate sits $+0.9\%$ high (the ledger road's $+1.1\%$ in η_{10} above; $+2.9\%$ if read at $h = 0.674$) with no $\mathcal{O}(1)$ freedom to price it — so the unowned residue sits on the exactness of $1/(2\pi^2)$ and on the one rate it is converted at, not on the transport $\mathcal{O}(1)$'s. Throughout this paper the density fractions $\Omega_m, \Omega_\Lambda, \Omega_b$ are h-free shares, scored in the codebook against Planck's own fractions and converted to densities at the carried 66.8; reconciling the two routes to Ω_b is a named open problem (two-route-consistency) recorded alongside the closed-packet-integration conjecture. A future precision GG-A10-Cosmos would compute the offset from the full entropy history, sphaleron throughput, reheating profile, and lock-scale neutrino kinetics; even the present estimate shows the matter-sector story is not wildly self-inconsistent.

The sharper prediction is the relative matter shares. Combining the π-locked inventory

$$\Omega_m = \frac{1}{\pi}, \qquad \Omega_b = \frac{1}{2\pi^2}, \qquad \Omega_{\mathrm{DM}} = \frac{1}{\pi} - \frac{1}{2\pi^2}, \tag{222}$$

fixes the dark-to-baryon ratio with no freedom,

$$\frac{\Omega_{\rm DM}}{\Omega_b} = 2\pi - 1 \simeq 5.283, \tag{223}$$

the cosmic baryon fraction of all matter,

$$f_b \equiv \frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} \simeq 0.159, \tag{224}$$

and its dark complement

$$f_{\rm DM} = 1 - f_b = 1 - \frac{1}{2\pi} \simeq 0.841. \tag{225}$$

These rigid, dimensionless numbers — about five times more dark matter than baryons, roughly sixteen percent of all matter baryonic — are obtained without independent per-sector fitting.

Table 18: Bridge from the baryogenesis language to the late-time inventory language. The relative matter ratios are especially sharp. The two right-hand columns reach the same quantities by different roads — the asymmetry route and the geometric inventory route — and the table is a consistency check rather than a derivation. Where the two disagree the disagreement is the finding, and the relative matter ratios are where it would show first.

Quantity	**Asymmetry route**	**Geometric inventory route**
η_B	$\approx 5 \times 10^{-10}$ (central 4.89) from Eq. (193)	not primary input
$\Omega_b h^2$	$\simeq 0.018$ from Eq. (221)	implied after choosing h
Ω_b	$\simeq$ 0.036–0.040 for $h \simeq$ 0.67–0.70	$1/(2\pi^2) \simeq 0.0507$
Ω_m	not directly fixed	$1/\pi \simeq 0.3183$
$\Omega_{\rm DM}/\Omega_b$	near observed value of order five	$2\pi - 1 \simeq 5.283$

Part V should be read not as three disconnected fixes but as one constrained matter ledger viewed at three epochs. The Kaxion fixes the dominant cold component; resonant leptogenesis fixes the sign and size of the baryonic remnant; BBN audits whether that remnant is compatible with light-element synthesis. The closer the three collapse onto one compact ledger, the more serious the program.

37.9 Chronological ordering inside Part V

A proposal can sound convincing while being chronologically inconsistent, so the temporal ordering deserves explicit statement. The first event is exit from the slow-smooth inflationary phase of Part IV, depositing energy in a hot bath at the derived reheating temperature,

$$T_{\rm RH} \simeq 6 \times 10^{10}\ {\rm GeV}. \tag{226}$$

The bath then cools, and the heavy Majorana sector becomes dynamical when the temperature falls to the lock scale, $T \sim M_{N_i} \sim M_{\rm lock} \simeq 3.85$ TeV, at $t \simeq 1.6 \times 10^{-14}\,$s. There the decays and inverse decays occur promptly, and the resonant-leptogenesis module belongs exactly here; the $B - L$ asymmetry is then reprocessed by sphalerons before the weak crossover shuts that channel off. The asymmetry is fixed extremely early, long before nuclei appear.

The Kaxion is different. It is already present as the low-energy remnant of the anomaly ledger's tensor, behaving initially as vacuum misalignment. Coherent oscillation onsets when

$$3H(T_{\rm osc}) \simeq m_K(T_{\rm osc}). \tag{227}$$

With a naively constant mass and radiation domination, Eq. (159) gives

$$T_{\rm osc}^{\rm naive} \simeq \left(\frac{m_K M_{\rm Pl}}{4.98\, g_*^{1/2}}\right)^{1/2} \simeq 49\ {\rm GeV} \qquad \text{for } m_K \simeq 8.1\,\mu{\rm eV}, \tag{228}$$

using a representative $g_* \sim 70$. Because the QCD-anchored mass turns on only near confinement, the true onset is pushed down; the full $\chi_{\rm QCD}(T)$ treatment of Arisaka-GG-C1-Kaxion places it at $T_{\rm osc} \simeq 1.0\,$GeV, in the QCD-era corridor. The point that matters: the Kaxion abundance is set after the baryon asymmetry, yet long before galaxies. No temporal clash.

BBN occurs much later, at MeV and sub-MeV temperatures, with the landmarks

$$T \sim 0.7\ {\rm MeV} \ \Rightarrow\ n/p \text{ freeze-out}, \qquad T \sim 0.08\ {\rm MeV} \ \Rightarrow\ \text{deuterium release}, \tag{229}$$

and synthesis around $t \sim 10^2$–$10^3\,$s. The lithium port acts vastly later than leptogenesis and vastly colder than the Kaxion oscillation onset. Each module occupies its own epoch and speaks to the others only through inherited conserved quantities, summarized in Table 19.

Table 19: Chronological ordering of the Part V mechanisms. The three matter-sector modules inhabit distinct epochs and combine without double counting. The columns give the epoch, its characteristic scale, and the role the mechanism plays. The ordering carries the argument: the three matter-sector modules occupy distinct epochs, so they combine without competing for the same energy budget, and a reader who suspects double counting can check it from this table alone.

Epoch	**Characteristic scale**	**Part V role**
Post-inflation reheating	$T_{\rm rh} \simeq 6 \times 10^{10}\,\mathrm{GeV}$	Supplies the bath in which the heavy Majorana sector later equilibrates (near 10–20 TeV, by inverse decays); the bath then cools through the lock, preparing the out-of-equilibrium conditions
Leptogenesis / sphaleron era	$T \sim \mathrm{TeV} \to 100\,\mathrm{GeV}$	Generates $B - L$, converts part to baryon asymmetry while preserving exact late-time proton stability
Kaxion misalignment to oscillation	below the electroweak era, toward the strong-coupling turn-on	Converts the compact-angle remnant into a cold matter component without disturbing the fixed baryon sign
BBN	$T \sim \mathrm{MeV} \to 0.01\,\mathrm{MeV}$, $t \sim 1$–$10^3\,\mathrm{s}$	Audits the baryon density and applies the selective A=7 lithium port

With this ordering, Part V acquires conceptual cleanliness: the same inherited GG-6/GDOS ledger determines what survives reheating, what baryon sign is left after the electroweak era, what fraction of matter ends up cold and non-baryonic, and whether the resulting baryon density passes the first precision nuclear audit.

37.10 A compact parameter ledger for the matter sector

It is worth recording how little continuous freedom the paper uses. The matter sector is governed by a small ledger once GG-A9-SM is treated as inherited input: the OGN integer set and the lock scale, the calibrated Kaxion decay constant and its discrete window, the already-derived neutrino masses and phases, the seesaw-matched Yukawa, the discrete sphaleron factor, and the closed-packet $1/\pi$ that fixes the lithium port. None of these is a free scan parameter; the apparent complexity of Part V is organizational, not parametric.

Table 20: Compact parameter ledger for the Part V matter sector, distinguishing inherited quantities, derived outputs, the one calibration, the one ansatz and the two named conjectures. The middle column is the honest one. It separates what is inherited from what is derived here and names what is neither, and a reader who wants to know what this Part actually assumes should read that column before any of the numbers.

Quantity	Status in Part V	Role
$K = 35$ and OGN integers	inherited discrete input	Anchor the Kaxion's cap-loading ladder; fix lock-scale exponents, $e^{-K/2}$ throughput
$M_{\rm lock} \simeq 3.85\,\text{TeV}$	inherited from GG-A9-SM	Heavy-neutrino scale and inflation-to-matter bridge
m_K benchmark window	calibrated (one declared scale; the ladder an ansatz)	Dark-matter oscillation and haloscope targets
$\theta_i = 1.18$ rad	output of the π-locked abundance	Misalignment amplitude, fixed by the abundance (not free); 1.18 ± 0.04 rad on the tabulated full-cosine integration (Appendix M)
single-patch history of the angle	conjecture C-Kaxion-d (GG-C1-Kaxion); necessary, not optional	One angle for the observable sky; without it de Sitter diffusion randomizes the angle and the isocurvature bound is exceeded (§35.5.2)
(m_1, m_2, m_3), $\delta_{\rm PMNS}$, $J_{\rm PMNS}$	inherited from GG-A9-SM	Anchor the light-neutrino and CP ledger feeding leptogenesis
$Y_\nu \simeq 7.3 \times 10^{-7}$	derived by seesaw matching	Fixes Γ_N and the washout scale
partial-resonance condition $\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$ (finite temperature)	one named assumption — a conjecture, not met at vacuum numbers	Lifts the tiny-Yukawa suppression of ϵ
$c_{\rm sph} = 28/79$	Standard Model fixed factor	Converts $B - L$ to the baryon asymmetry
A=7 port factor $1/\pi$	closed-packet prediction (D*)	Fixes the lithium deficit, same origin as $\Omega_m = 1/\pi$

Part V is thus a compact stress test of whether a few inherited geometric inputs can simultaneously speak to cold dark matter, the baryon asymmetry, and the first nuclear audit of the universe.

37.11 Bottom line of Part V

The three origin-of-matter problems fit together far more tightly than in an unconstrained model-building program. The Kaxion is the natural first dark-matter candidate because it descends from the compact-angle and Green–Schwarz sectors already present in the inherited theory. Baryogenesis is driven to resonant leptogenesis because the neutrino sector is already Majorana and CP-active while $\mathrm{U}(1)_B$ remains exact; the seesaw matching fixes the Yukawas and the chain gives $\eta_B \approx 5 \times 10^{-10}$ (central 4.89; band $\times 3^{\pm 1}$; the time-resolved integration lands a factor of about ten below the ledger road's 6.18×10^{-10}, §I.6). BBN is the first precision audit, and its lithium anomaly is addressed by the inherited closed-packet factor $1/\pi$ — the same π-locking that fixes the matter inventory — the $A = 7$ port that would realize it an open item (PS-Cosmos-15), and deuterium and helium untouched.

Table 21: Compressed summary of the Part V outputs — the primary matter-sector targets of GG-A10-Cosmos. The columns give the problem, the solution logic, and the characteristic derived target. The third column is a target rather than a measurement: it states what the mechanism produces, and whether the sky agrees is settled in the audit tables, not here.

Problem	**Part V solution logic**	**Characteristic derived target**
Dark matter	Compact-angle / GS descendant ⇒ Kaxion window	$m_K = 8.1\,\mu$eV (one scale calibrated), window 1.59–2.13 GHz, one line alone, DFSZ address
Baryogenesis	Lock-scale Majorana triplet, seesaw-matched, resonant CP	$\eta_B \approx 5 \times 10^{-10}$ ($\times 3^{\pm 1}$; time-resolved, a factor of about ten low, §I.6), $\Omega_b = 1/(2\pi^2)$
BBN / lithium	Inherited A=7 closed-packet port	$(^7\mathrm{Li}/H) = \mathrm{SBBN}/\pi \simeq 1.6 \times 10^{-10}$, D/H and Y_p unchanged

Step back from the machinery and consider what this Part has claimed, in plain words. The dark matter of the universe is a frozen angle — a direction held everywhere at once, by the conjecture this Part names, before any matter existed to disagree. The excess of matter over antimatter — the reason there is anything rather than mutual annihilation — is bookkeeping on the same neutrino ledger the particle companion derived. And the first three minutes of nuclear history check out at the percent level on the inventory's $\Omega_b h^2$, not on that ledger, whose asymmetry lands a factor of a few to ten below it, with lithium-7 read honestly as the one open port.

There is something worth pausing on here. Every atom in the reader's body is a survivor of the asymmetry computed in this Part; the carbon in your hand is, several floors up, an arithmetic consequence of an integer ledger that was fixed before the first second was over. Whether or not the derivation survives its tests, that is the right *shape* for an answer to "why is there something?" — not a shrug at an accident, but a chain that can be checked link by link.

The bridge to Part VI is immediate. Once the matter ledger is fixed — cold dark matter, baryon fraction, neutrino block, and the early-time BBN audit — one can ask how perturbations propagate, how the CMB reads them out, and whether the resulting growth history eases or worsens the present tensions in H_0, S_8, and early-galaxy formation. That is the task of the next part.

38 Part V in plain English: where the matter came from, and what is assumed

Three separate accounts sit in this Part — dark matter as a misaligned Kaxion field, the baryon asymmetry from the inherited neutrino and CP ledger, and the light-element abundances — and the reason they belong together is that they occupy distinct epochs and draw on the same inherited ledger. They do not compete for the same energy budget, which is what allows them to be combined without double counting, and the chronology table is where that can be checked.

The dark-matter half rests on one declared calibration and one named conjecture, the single-patch history, which the isocurvature bound makes necessary.

The baryogenesis chain is the place to be careful. It has zero fitted parameters, which is not the same as having no assumptions: one link is a named input, the partial-resonance condition, and the exponent in the throughput is a postulate of the compression rather than a derived quantity. Both are marked as such in the ledger, and the honest description of the chain is that it lands on the observed asymmetry from inherited numbers with two named gaps in the middle, not that it derives it.

The lithium-7 audit is included because it is the one place in this Part where the framework predicts a deviation from the standard baseline rather than agreement with it. That makes it the most useful row in the block. A measurement that placed the primordial abundance at the full standard value would close the port this Part opens, and the Part says so.

Part VI
Growth of structure and the present-day tensions

CMB baseline and the scalar/tensor audit; one derived Hubble background with the distance ladder as the predicted outlier; the S_8/σ_8 clustering amplitude computed and its tension shared; and early-galaxy formation in the JWST era

---

39 One switch, two faces: the geometric mechanism behind the present-day tensions

Part VI is organized around a single mechanism, stated once here and then audited observable by observable: *recombination is the cosmic GRIP-Out event, the corridor switch occurs at the integer-ledger redshift $z_{\rm sw} = e^7$, and every present-day cosmological tension is either that switch or the GRIP-In growth that follows it.* The two headline puzzles of contemporary cosmology — the Hubble tension and the JWST "too-early massive galaxies" anomaly — are therefore not two unrelated problems requiring two unrelated fixes. They are two faces of one geometric switch.

The runtime grammar is the GRIP cascade GRIP = (DGI → GSSB → GA) inherited from GG-A6-GRIP and already used as the spine of this paper (Section 14; the cascade figure of Part II). Before recombination the baryon–photon plasma is a tightly coupled DGI phase: an enormous family of micro-histories sharing one observable acoustic gauge orbit. Recombination is the GSSB transition that freezes those phases into the GA attractor we read as the microwave background. The switch redshift is not fitted; it is forced by the OGN-5 integer $J = 7$ of GG-A4-Euclid,

$$z_* = e^J = e^7 \approx 1096.6,$$

the corridor-switch redshift; the atomic decoupling derived from the framework's own constants lands 0.4% below it, on the Planck value at the 0.1% level (§42.1). From that single event the Part unfolds in two directions:

- **Face 1 — the Hubble tension (one background, one ladder).** The boundary Friedmann equation has one background rate, $H_0 = 66.8 \pm 0.2\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, derived with zero sky input; the sound-horizon dictionary ($z > z_*$) reads it at 67.4 through the inherited CMB calibration, and every ladder-free late probe is predicted to read it at the same value. The distance ladder's 73.5 is not a second reading of the corridor but a measurement the framework predicts to carry a calibration systematic (Theorem 68.8, (C3); Corollary 68.11). An earlier version of this face read the corridor twice, at 67.4 and 73.0; the second reading is withdrawn.

- **Face 2 — S_8 and the JWST anomaly (the GRIP-In growth, and what it cannot do alone).** Once the GRIP-Out event has frozen the attractor, the universe runs as GRIP-In relaxation: structure grows on the inherited π-locked ledger. The S_8 clustering observable is that one growth history read on broad linear scales, and its amplitude is computed from the declared A_s: $S_8 = 0.825$, shared with ΛCDM's CMB-normalized value and 2.5σ above the cosmic-shear one (§43.3). The JWST early-structure anomaly is read on the exponentially sensitive high-ν rare tail — but it is *not* that same growth read again in the direction the anomaly needs: a change in the expansion history multiplies every sub-horizon mode by one common factor, so any background residue would move both windows the same way — and there is none (§44.3). What carries the rare tail is the *baryonic* lever — the reduced collapse threshold, the streaming-velocity damping having no mediator under the terminated baryon gauge sector (§44.4). One growth history sets the broad window; a baryonic lever, and only such a lever, can separate the two.

The four observables audited below — the CMB acoustic structure, H_0, S_8, and the abundance of the earliest luminous systems — are thus a single GRIP-Out switch at $z_* = e^7$ plus the GRIP-In growth after it, all on one inherited ledger. Table 22 collects the falsifiable numbers in one place; each is derived in the section named in its last column, and the heavier derivations are deferred to the Technical Appendices (Sections J and K) so the body stays readable. The reader is encouraged to consult those appendices for the full from-scratch calculations.

Table 22: Headline falsifiable numbers of Part VI. Every quantitative prediction of the growth-and-tension block, tagged to its role in the one-switch/one-growth mechanism and compared to the current best datum. $\Delta N_{\rm eff}$ is the derived null 0 throughout the body (row D12, Section 68.2.2); an earlier version used 0.04, and the legacy early-dark-energy linear-response that used 0.30 is retained only as a non-unique cross-check in Appendix K.

Observable	GG-Theory	Best datum	Dev.	Mechanism role
Decoupling z_*	1091.5 (atomic chain at the framework's own h; switch e^7 distinct, not scored)	1089.92±0.25 (Planck Collaboration, 2018)	+0.15%	GRIP-Out event at the switch $z_{\rm sw} = e^7$ (§42.1)
Background $H_0^{(0)}$	66.8±0.2 (Route I, zero sky input; 66.6 from the inventory)	67.4±0.5 (Planck Collaboration, 2018)	−0.9%	one rate on every probe
Local ladder $H_0^{\rm late}$	the same 66.8	73.50 ± 0.81 (Casertano et al., 2026)	−9.1% (predicted)	the ladder read as a calibration systematic; ladder-free probes decide
Clustering S_8	0.825±0.010 (from A_s)	0.832 ± 0.013 CMB; $0.759^{+0.024}_{-0.021}$ shear (Asgari et al., 2021)	-0.4σ; $+2.5\sigma$	standard growth on the inventory; the tension shared
Acoustic angle $100\theta_*$	1.038 (derived, at the carried rate)	1.04109 ± 0.00030 (Planck Collaboration, 2018)	−0.3%	switch geometry
Damping multipole ℓ_D	1350 ± 3 (Boltzmann)	≈ 1345 ± 15 (inferred from the $100\theta_D$ of Planck Collaboration 2018; the ±15 a convention spread)	+0.4%	diffusion on the inventory
Scalar tilt n_s	0.9634 exact (0.9620 l.o.)	0.9649 ± 0.0042 (Planck Collaboration, 2018); 0.968–0.973 (2025)	-0.4σ; -1.5 to -3.2σ	plateau-derived; the 2025 tension stated
Tensor ratio r	0.0038 exact (0.0043 l.o.)	< 0.036	two-sided target	derived plateau
Collapse threshold δ_c	1.686 → 1.18	—	first-light lever	GRIP-In, rare tail
First-light shift Δz	$\sim$ 2–3 ($\delta_c^{\rm eff} = 1.43$–1.325, Eq. (284))	JWST excess	qualitative	GRIP-In, rare tail

40 Determinism from geometry: the microwave background as an over-determined test

Before the acoustic physics of the next section is unpacked, it is worth stating plainly what kind of test the microwave background is about to administer, because it is a stronger kind than cosmology usually gets to run. Standard ΛCDM describes the CMB with six free parameters

$\{\omega_b,\ \omega_c,\ \theta_*,\ \tau,\ A_s,\ n_s\}$: two physical densities, one angle, one optical depth, and two numbers for the primordial seed. The model is spectacularly successful at *compressing* the data onto those six dials, and completely silent on why the dials read what they read. GG-6 makes a claim of a different kind, inherited from the determinism spine of the cohort (GG-A11-Origin, Arisaka, 2026, A11-Origin) and the integer ledger of the particle stratum (GG-A9-SM, Arisaka, 2026, A9-SM): the *background-geometry half* of those dials is not adjustable at all. The matter fraction, the dark-energy fraction, the baryon fraction, the equation of state, and the recombination redshift are theorems of one closed geometry,

$$\Omega_m = \frac{1}{\pi}, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi}, \qquad \Omega_b = \frac{1}{2\pi^2}, \qquad w = -1, \qquad z_* = e^7,$$

with no continuous parameter anywhere in the chain (Parts II, IV, and XI). What remains genuinely free in the framework's present state is the primordial seed (A_s, n_s) — the amplitude declared as input, the tilt derived-conditional at 0.9634 (Section 41.1) — and the reionization depth τ, which is astrophysics, not fundamental geometry. Cosmology's six-parameter standard model becomes, in effect, a three-parameter one, and the half that is deleted is precisely the "why these values" half.

That deletion is what makes the CMB an *over-determined* test rather than a fitting exercise. The microwave sky measures more independent quantities than the framework leaves free: the baryon density, the matter density, the acoustic angle, the damping scale, the peak-height ratios, and the equality and recombination epochs are read off the spectrum through largely independent pieces of physics. ΛCDM reconciles them by adjusting its dials; GG-6 has no dials to adjust, so every one of those readings is a separate opportunity to be wrong. A single clean disagreement — one observable demanding a fraction incompatible with $1/\pi$ — would falsify the chart outright. The framework's predictions moreover sit exactly along ΛCDM's celebrated *geometric degeneracy* directions (the Ω_m–H_0 line at fixed θ_*), which is where a fitted model constrains least and where a forced model is therefore most exposed. The honest statement of the program, maintained throughout this Part, is: GG-6 converts the background-geometry sector of the CMB into theorems and lets the data over-constrain the rest; what it does not derive (the seed, τ, the local distance ladder) is named as such and never absorbed.

40.1 The master identity: fractions, densities, and several roads to one rate

The engine of the over-determination is one elementary identity. A *fraction* Ω_i is a dimensionless share of the critical density; a *physical* density $\omega_i \equiv \Omega_i h^2$ is a statement about grams per cubic meter and contains no Hubble constant at all. The CMB, to an excellent approximation, measures physical densities: $\omega_b = 0.02237 \pm 0.00015$ from the alternating (odd/even) peak-height asymmetry that baryon loading imprints on the acoustic oscillation — independently confirmed by the BBN deuterium abundance — and $\omega_m = 0.1430 \pm 0.0011$ from the peak-height envelope and the matter–radiation equality it encodes. ΛCDM treats the fractions as unknowns and fits them. GG-6 predicts the fractions outright, so each measured physical density becomes an independent *road to the*

expansion rate:

$$h = \sqrt{\frac{\omega_i}{\Omega_i}} \implies h_{\rm baryon} = \sqrt{\frac{0.02237}{1/2\pi^2}} = 0.665, \qquad h_{\rm matter} = \sqrt{\frac{0.1430}{1/\pi}} = 0.670, \tag{230}$$

two determinations through two different pieces of acoustic physics, agreeing to 0.9% — exactly half the 1.7% of the ratio test below, since the two tests are one statement — with no quantity available to reconcile them had they disagreed. Their combination is the two-road reading of the background rate, $H_0^{(0)} = 66.6 \pm 0.2$ km s^{-1} Mpc^{-1}, stated in Section 29.8 beside the rate itself, derived there with zero sky input by the geometric Route I at 66.8 ± 0.2 — and the age $t_0 = 0.948/H_0 = 13.9$ Gyr (today; 13.8 at the $\Omega_m{=}1/\pi$ crossing) follows with it. Figure 24 draws the full web: every arrow is an independent measurement channel, and they converge on one number because the chart leaves them nowhere else to go.

There is one test purer still. Since both physical densities scale as h^2, their *ratio* is h-free and equals a pure geometric number:

$$\frac{\omega_b}{\omega_m} = \frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} = 0.15915, \qquad \text{observed } 0.15643 \pm 0.0016 \;\; (+1.7\%,\; 1.7\sigma). \tag{231}$$

This is the baryon-loading theorem (T-Cosmos-9) of the next section stated as a falsifiable number: the quantity that sets the *relative* heights of the acoustic peaks is $1/2\pi$, with no freedom whatsoever — not even the Hubble constant. The matter–radiation equality redshift, by contrast, is deliberately *not* counted as an independent test here: on the two-road reading, once h is taken from the matter road, $z_{\rm eq} = \omega_m/\omega_r$ is an identity, and an honest audit does not count the same measurement twice. On the carried Route I rate, which takes nothing from the sky, the released run returns $z_{\rm eq} = 3376$ against the measured 3387 ± 21 (-0.3%, -0.5σ) — a check, not a row.

40.2 The core observables, computed and confronted

Table 23 collects the core CMB observables that the π-locked chart forces, each computed two ways: semi-analytically (the standard tight-coupling expressions with every input GG-6-fixed) and through a full Boltzmann solution (CLASS) run at the exact fractions. In the numerical runs everything except the background chart is held at its standard value: the Planck neutrino sector, and the primordial seed and optical depth frozen at the Planck best-fit. (In the present codebook the tilt n_s is *derived*, and only two quantities remain declared as inherited — the amplitude A_s and the astrophysical τ; the full-spectrum confrontation below is run at the derived tilt and the carried rate.) Two bookkeeping conventions are stated once: the ledger corridor-switch redshift $z_* = e^7 = 1096.6$ — a topological quantity, not the measured decoupling redshift — sits +0.75% above the Boltzmann ionization-history value $z_{\rm rec} \simeq 1088.5$, the fourth member of the 0.38/0.47/0.62% switch–decoupling family of §42.1; and pipeline definitions of r_s and θ_* differ between analysis chains at the 0.1% level, which is folded into the quoted deviations.

Three rows deserve explicit comment, because they are the ones Phase-level audits most often get wrong. First, the *peak-height ratios* are the spectrum-level image of the h-free ratio (231): the driving ratio deviates by +1.7%, and the full radiation-transfer machinery compresses that

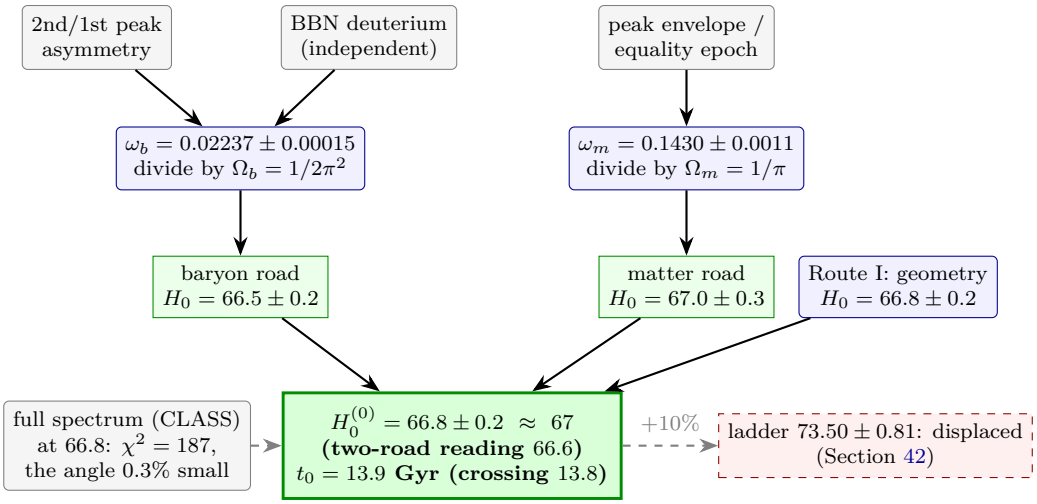

Figure 24: The H_0 consistency web. Each solid arrow is an independent measurement channel: the odd/even peak asymmetry and (independently) BBN deuterium fix the physical baryon density; the peak envelope fixes the physical matter density; dividing each by its forced GG-6 fraction yields an independent road to the expansion rate. The two density roads agree at 66.6 ± 0.2 (the two-road reading); the geometric Route I (zero sky input, Section 29.8) lands 0.3% above it at $H_0^{(0)} = 66.8 \pm 0.2$, the rate the paper carries; and the full Boltzmann spectrum (CLASS, exact fractions) at the carried rate reproduces the band-powers to within −1 to +3.5%, its whole excess the 0.3% acoustic-angle deficit ($\chi^2 = 187$ against the best fit's 84; Section 40.3). The local distance ladder is the lone displaced reading; it is analyzed — not derived — in Section 42.

into a 0.8% (second-to-first) and 0.9% (third-to-first) shift in the observable ratios — the direct spectral test of $1/2\pi$, and it passes at the one-percent level. Second, the *damping multipole*: running the exact fractions through the full diffusion integral over the Boltzmann ionization history gives $\ell_D = 1350 \pm 3$ (quoted in the ΛCDM-inferred convention, compared with the inferred 1345 ± 15), a +0.4% agreement. Two cautions belong beside that number. The 1350 is the envelope-fitted damping multipole in the convention of the ΛCDM-inferred value it is compared with, and no released script returns it; the ±15 is this paper's estimate of the convention spread, not a Planck error — Planck quotes no ℓ_D, only $100\,\theta_D$. In the Hu–White convention that CLASS itself computes, $\ell_D \equiv k_D D_M$, the released run gives 1956 against Planck's 1947 ± 5 (from $100\,\theta_D$): +0.5%, a pull of +1.6 with the row's $\sigma_{\rm th} = 3$. The inventory predicts the damping tail to half a percent in either convention; the consequences for the separately proposed BFFT reading of ℓ_D are drawn in Section 41.1. Third, the *acoustic angle* θ_* is the most precisely measured number in cosmology (0.03%), so its −0.3% deviation, small as it looks, is a formal pull of −2.4, tied with the S_8 pull as the sharpest in the codebook after the distance ladder's; it is examined honestly, not hidden, in the next subsection.

40.3 The full spectrum from the derived parameters

The sharpest form of the confrontation puts the whole temperature spectrum on one panel. Figure 25 shows the lensed TT spectrum computed (CLASS) from the GG-6-derived parameter set — the exact π-locked fractions, the derived $H_0 = 66.8 \pm 0.2$ (Route I, the rate the paper carries),

Table 23: The CMB core, derived and confronted. Every GG-6 entry follows from $\Omega_m = 1/\pi$, $\Omega_b = 1/2\pi^2$, $\Omega_\Lambda = 1 - 1/\pi$, $w = -1$, and the corridor switch $z_* = e^7$ with zero tuned parameters (the atomic decoupling derived from the same inventory lies 0.4% below the switch and 0.15% above the Planck value; §42.1); the seed and τ are frozen at Planck values and declared inherited. *Zero CMB input*: every GG-6 entry is forced by the geometry — the fractions by the Hopf fibration, the rate by the two routes of Section 29.8, the densities by their product, evaluated at the rate the paper carries, 66.8 ± 0.2 (Route I) — and the sky appears only in the observed column. Against the 2025 three-instrument combination 67.24 ± 0.35 the Route I row reads -0.7%, -1.1σ. Spectrum-morphology rows (peak positions, height ratios) are compared against the Planck-2018 best-fit spectrum computed in the same pipeline; parameter rows against the Planck 2018 chain values (Planck Collaboration, 2018). Grades are the codebook triage of Part IX (D* = derived modulo a named lemma). The uniform percent-level, same-family displacement of all rows is the coherent π-offset discussed in Section 40.4. Every computed row is evaluated at Route I's $h = 0.6676$; at the combined two-road reading $h = 0.666$ the two density rows read 0.02247 (+0.5%) and 0.14119 (−1.3%). The paper carries one rate, 66.8; the combined two-road row is the same background read through the measured densities, not a second rate, and the 0.3% between the two readings of h is inside every entry's error except $100\,\theta_*$'s, whose -0.3% is the sharpest single pull in the table.

observable	GG-6 forced form	GG-6 value	observed	Dev	grade
H_0 (Route I: geometry)	$498.3\sqrt{\pi/175}$	66.8 ± 0.2	67.4 ± 0.5	-0.9%	D*
H_0 (baryon road)	$\sqrt{\omega_b/\Omega_b}$	66.5 ± 0.2	—	—	D*
H_0 (matter road)	$\sqrt{\omega_m/\Omega_m}$	67.0 ± 0.3	—	—	D*
H_0 (combined)	two roads, one reading	66.6 ± 0.2	67.4 ± 0.5	-1.2%	D*
ω_b (predicted)	$h^2/2\pi^2$	0.02258	0.02237 ± 0.00015	$+0.9\%$	D*
ω_m (predicted)	h^2/π	0.14187	0.1430 ± 0.0011	-0.8%	D*
ω_b/ω_m	$1/2\pi$ (h-free, epoch-free; observed = ratio of the two rows above)	0.15915	0.15643 ± 0.0016	$+1.7\%$	D*
$z_{\rm dec}$	atomic chain on the π-locked inputs (the switch $e^7 = 1096.6$ is annotated, not scored)	1091.5	1089.92 ± 0.25	$+0.15\%$	D*
$r_s(z_{\rm rec})$ [Mpc]	diffusion-free acoustic scale	144.7	144.43 ± 0.26	$+0.2\%$	D*
$100\,\theta_*$	r_s/D_M	1.0381	1.04109 ± 0.00030	-0.3%	D*
peaks $\ell_1/\ell_2/\ell_3$	acoustic projection	221/539/816	220/536/813	+0.4–0.6%	D*
height ratio 2/1	image of $1/2\pi$	0.449	0.453 ± 0.008	-0.8%	D*
height ratio 3/1	image of $1/2\pi$	0.439	0.443 ± 0.009	-0.9%	D*
ℓ_D (damping)	diffusion integral (the carried rate moves it by +0.03%)	1350 ± 3	1345 ± 15	$+0.4\%$	D*

the derived tilt $n_s = 0.9634$ exact on the plateau, and the counted radiation content ($N_{\rm eff} = 3.046$ in the run, CLASS's default; the codebook's 3.044 differs by 0.002, below any row's resolution) — with exactly two numbers inherited and set to the Planck values, the amplitude A_s (the launch sector's one declared input) and the astrophysical τ. Nothing in the red curve is adjusted to the data. Against the 83 binned Planck TT points the six-parameter best fit gives $\chi^2 = 84$ ($\chi^2/{\rm dof} \approx 1$ by construction); the zero-parameter GG-6 curve gives $\chi^2 = 187$, i.e. $\chi^2/{\rm dof} \approx 2.3$. (Earlier versions of this figure were run at the two-road reading 66.6 and, before that, at the leading-order tilt; the

curve is now computed at the rate and the tilt the paper stakes, and every number in this block comes from that one run.) We state plainly what that means and what it does not: as a *likelihood fit* the present curve would be rejected — $\chi^2/\mathrm{dof} \approx 2$ is not a fit, and it would be dishonest to call the instrument's output an "endorsement" of the inputs. But the 187 is not 187 of *random* scatter: the curve tracks the data to within -1 to $+3.5\%$ across the first seven acoustic peaks, and the residual is a *single coherent oscillation*. An earlier version read that oscillation as a normalization — the same $+2.7\%$ second-order offset that the second-order-residual audit (Section 84.5) isolates in ω_b and w — and expected one nuisance parameter for it to bring χ^2 toward the ΛCDM value. It does not: a free amplitude takes χ^2 only from 187 to 165. A single stretch of the multipole axis by -0.33% — the acoustic-angle deficit of the previous subsection, and nothing else — takes it to 82.5, *below* the six-parameter best fit's 84.3 (the same freedom gains the best fit 1.9), and narrows the residual from $(-1, +3.5\%)$ to $(-0.15, +1.1\%)$. The whole excess of the zero-parameter curve is the acoustic angle; the shape of the spectrum is already right. The stretch is applied at fixed densities, and moving h along the π-locked line does not reproduce it: χ^2 falls only to about 112, near $h = 0.672$, because ω_b, already $+0.9\%$ above the measured 0.02237 at the carried rate, rides up to $+2.3\%$ with h^2. What the framework has to move is therefore D_M/r_s at fixed ω_b, not the seed, and how it does so is open (`cmb_stretch_v1.py`). One thing this comparison does *not* show is a preference for the computed tilt: on the π-locked background χ^2 falls monotonically as the tilt reddens (196 at Planck's 0.9649, 187 at the derived 0.9634, 180 at the leading-order 0.9620), so the binned spectrum on this background wants a redder tilt than the theory computes, while the 2025 combinations want a bluer one (§28.10). The tilt is not tested here; it is tested there.

40.4 The π-offset: one named systematic, visible everywhere

Every row of Table 23 deviates from observation at the percent level or below, and the deviations are not scattered: they are one number's shadow. The exact fractions sit slightly above Planck's fitted ones ($1/\pi = 0.3183$ against 0.3153 ± 0.0073, $+1.0\%$), and that single displacement propagates with the expected sign through every observable — the derived h comes out 0.9% low and θ_* projects 0.3% small (not the same statement: Section 40.3), the peaks slide 0.4–0.6% rightward, the damping scale follows at $+0.4\%$, and the atomic z_{dec} and the ratio $1/2\pi$ sit $+0.15\%$ and $+1.7\%$ high. A model *fitted* to these data would scatter its residuals on both sides; a *forced* chart displaced by one small systematic lands coherently on one side, and that coherence — verified across the full spectrum, into the damping tail, with no compensation anywhere — is the signature that the chart is forced rather than tuned. The offset is named, permanent, and carried openly: it either is the truth about the universe (and the fitted fractions are absorbing a small unmodeled effect), or it is the falsification-in-waiting of the exact chart. The framework does not get to choose; sharper data do.

Honesty requires the two strongest statements of the offset to be made by the framework itself. The acoustic angle is measured to 0.03%; against the 0.12% modeling spread of the recombination chain, the -0.3% deviation is a formal pull of -2.4, tied with the S_8 pull as the largest in the codebook after the distance ladder's -8 — the π-offset concentrated into cosmology's most precise number. And a chi-squared against the 83 published Planck TT band-powers, with diagonal errors

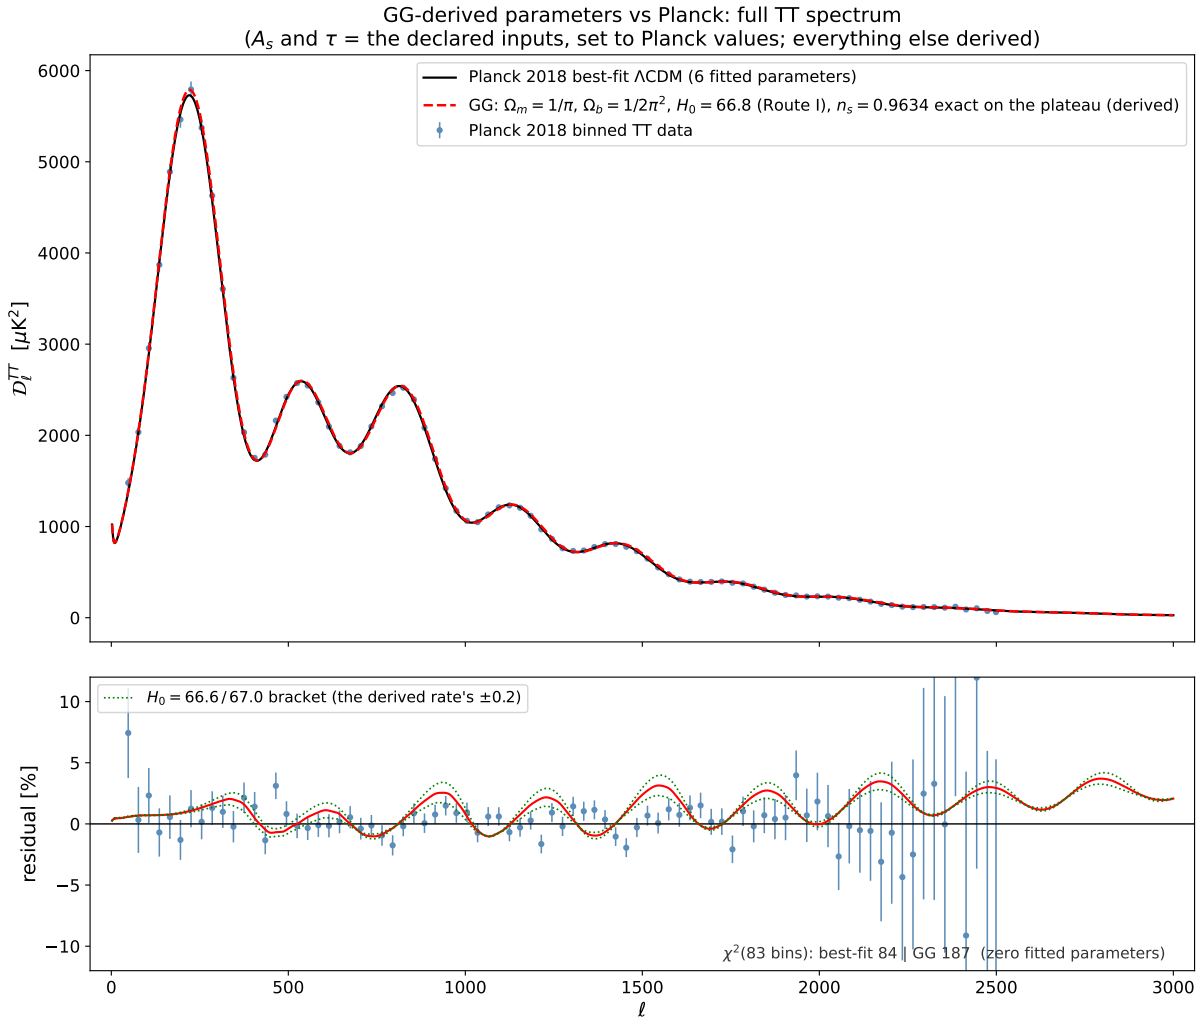

Figure 25: The microwave sky from the ledger. *Top:* the lensed TT spectrum for the Planck 2018 best-fit ΛCDM (black; six fitted parameters) and the GG-6 prediction (red dashed; exact π-locked fractions, derived $H_0 = 66.8$ (Route I), derived $n_s = 0.9634$; only A_s and τ inherited, set to Planck values, never adjusted), against the Planck 2018 binned TT data (83 bins). *Bottom:* fractional residuals relative to the best fit, with the derived rate's bracket $H_0 = 66.6/67.0$ (dotted). χ^2 (83 bins, diagonal): best fit 84 with six fitted parameters (χ^2/dof ≈ 1); GG-6 187 with zero (χ^2/dof ≈ 2.3 — not a likelihood fit, and not claimed as one). The residual is a single coherent oscillation between -1 and $+3.5\%$, the -0.3% acoustic-angle deficit (one stretch of the multipole axis takes χ^2 to 82.5), not random scatter; on this background χ^2 is monotone in the tilt and does not select the computed value. Reproduction: `gg_cmb_tt_v4.py`; the stretch, `cmb_stretch_v1.py`.

and the inherited seed left frozen, gives 187 for the exact chart against 84 for the six-parameter best fit: at Planck precision the data *can* tell the exact chart from the fitted one — through the acoustic angle, not the shape of the spectrum. Both numbers overstate the case against the chart — bin-to-bin correlations are ignored — and the seed is not the issue: re-optimized, it leaves 165, while the angle, re-set, leaves 82.5; the direction is honest and is not negotiated away. The claim

defended here is precisely calibrated: a zero-parameter background lands within about one percent of a six-parameter fit everywhere at once, displaced one way by one named systematic; it is never claimed that the exact chart out-fits ΛCDM. A full-likelihood confrontation, with the seed re-fitted and correlations kept, is the sharpest version of this test and is recorded as an open item rather than pre-judged.

40.5 What is claimed, and what is not

The boundary of the claim matters as much as the claim. *Claimed:* the dimensionless background chart is theorem-grade (modulo the named lemmas tabulated in Part IX), and from it the expansion rate, the age, the acoustic geometry, the damping scale, and the peak morphology follow with zero adjustable parameters, verified in a full Boltzmann computation, together with the derived-conditional tilt $n_s = 0.9634$ exact on the plateau (Lemma IV, revised), risked against the tilt measurements of §28.10 rather than borrowed from them (the TT band-powers of Figure 25 do not select it). *Not claimed:* the primordial amplitude (the launch sector's declared input, Lemma V′), the optical depth (astrophysics), the absolute peak heights insofar as they lean on the inherited seed, the equality redshift as an independent success (an identity on the matter road; -0.5σ on Route I, §40.1), and the local distance-ladder reading of 73.50 ± 0.81, which the framework analyzes as a displaced measurement (Section 42) and expressly does not derive. The falsification handles are correspondingly sharp: a gravitational-wave standard-siren program converging at ≈ 73 would refute the inventory outright, a converged siren at ≈ 67 would confirm it against the ladder, and any future CMB analysis pushing the fitted fractions two of their present error bars away from $(1/\pi,\ 1/2\pi^2)$ would end the exact chart without appeal. That is the intended posture of this Part: not a model defended, but a derivation exposed.

41 CMB baseline, primordial perturbations, and the scalar/tensor audit

41.1 Historical background

The cosmic microwave background was discovered serendipitously in 1965 by Penzias and Wilson (Penzias and Wilson, 1965) at Bell Laboratories as an "excess antenna temperature at 4080 Mc/s". The same volume of the *Astrophysical Journal* carried the companion paper by Dicke, Peebles, Roll, and Wilkinson (Dicke et al., 1965) interpreting the signal as the cooled afterglow of a hot early universe. Six decades of progressively higher-precision experiments — most recently the Planck satellite (Planck Collaboration, 2018) — have established the CMB as the most quantitatively precise observable in cosmology. The present section retains the standard acoustic, damping, and projection structure and asks where the inherited ledger reshapes the source assumptions behind it.

Theorem 41.1 (PI-locked baryon-loading statement (T-Cosmos-9))**.** *On the late inventory branch,*

$$f_b \equiv \frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi},$$

so the baryon loading of the photon-baryon plasma is not a free CMB fit parameter in the framework but a branch descendant of the same late matter chart. The standard acoustic formulas remain

$$R(a) = \frac{3\rho_b}{4\rho_\gamma}, \qquad c_s(a) = \frac{c}{\sqrt{3(1+R(a))}}, \qquad \theta_* = \frac{r_s(z_*)}{D_A(z_*)},$$

but the matter share entering them is narrowed by the inherited late branch.

Proposition 41.2 (Tensor target)**.** *Because the inflationary branch is the derived plateau — class forced, count matched, no dial — the primordial tensor sector must obey*

$$r = 0.0038 \quad (12/N^{*2} = 0.0043 \text{ at leading order}).$$

A tensor measurement inconsistent with this value in either direction would therefore kill the minimal branch rather than merely nudge a fit parameter.

Part VI begins from a methodological rule already emphasized in the introduction: the present program is not strongest when it announces a wholesale replacement of precision cosmology, but when it keeps the successful perturbative grammar of standard cosmology and asks a narrower and sharper question. Which pieces of that grammar are merely kinematic, and which pieces still hide unexplained source terms, normalizations, or module choices? In the present section the ordinary perturbation equations are retained as the boundary baseline. The new move lies in the inherited ledger used to interpret them.

The immediate empirical target is the microwave background together with the primordial scalar and tensor sectors. That target is unavoidable. A framework that proposes a cosmic beginning at E_{GUT}, a derived slow-smooth plateau one e-fold below it, a compact-angle dark sector, and a measurement-aware boundary semantics must survive the acoustic ruler, the peak-height pattern, the damping tail, the polarization sequence, the scalar tilt, and the absence or presence of primordial tensor power. In other words, the sky already provides the natural audit of the first, third, and fifth preprints simultaneously.

The strategy of the present section is therefore five-fold. First, we restate the standard perturbation baseline in a form that makes clear what GG-6 is *not* changing. Second, we reinterpret recombination as a GDOS event — more precisely, a transition from a DGI-dominated tightly coupled plasma to a frozen geometric attractor. Third, we carry forward the π-locked cosmic inventory already motivated elsewhere in the paper and ask how far it can push the acoustic structure without introducing continuous new dials. Fourth, we state the derived corridor for the scalar and tensor sectors. Fifth, we summarize the kill conditions. Throughout, the working stance is conservative in equations and maximal in inherited constraint.

41.2 What this section inherits from the previous papers

The present section does not start from a blank cosmological slate. It inherits a non-trivial ledger:

1. the boundary Einstein equations as the zero-mode export of the six-dimensional bulk,

2. the radial coordinate $r = \ln(E_{\rm GUT}/\mu)$ as the geometric encoding of RG descent,
3. the BI-6 anomaly structure and the OGN-5 integers $(1, 3, 5; 7, 35)$,
4. the compact-angle ontology that already generates the Kaxion window and the baryonic $\mathrm{U}(1)_B$ sector,
5. the GDOS runtime, together with DGI, GSSB, and GA, and
6. the guiding interpretive rule that the observable universe is a boundary-accessible open subsystem rather than a directly readable bulk microstate.

These inherited items already narrow the permissible cosmology. They do not yet determine every transfer function or every numerical likelihood value. But they do eliminate much of the usual freedom. In particular, they strongly disfavor an arbitrary inflaton potential, arbitrary extra-light fields, arbitrary dark-radiation additions, arbitrary non-geometric tensor sectors, and arbitrary late-time fixes introduced only to move one observational point. If a correction is needed in the CMB or structure-growth chain, the first place to look is the already inherited geometric ledger.

41.3 Standard perturbation baseline: what remains unchanged

At the homogeneous level the background metric is the standard spatially flat FLRW line element,

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + a^2(t)\,\delta_{ij}\,\mathrm{d}x^i\mathrm{d}x^j, \tag{232}$$

with Hubble rate

$$H \equiv \frac{\dot{a}}{a}. \tag{233}$$

At linear order, scalar perturbations may be written in conformal Newtonian gauge as

$$\mathrm{d}s^2 = -\,(1+2\Phi)\,\mathrm{d}t^2 + a^2(t)\,(1-2\Psi)\,\delta_{ij}\mathrm{d}x^i\mathrm{d}x^j, \tag{234}$$

where Φ and Ψ are the usual Bardeen potentials. The reason to write the metric in this familiar form is that it makes the basic scientific point transparent: the perturbation language used by the data community is not being rejected. It is being inherited and then interpreted through a different ultraviolet and operational ledger.

For any fluid component with density contrast δ, velocity divergence θ, equation-of-state parameter w, and adiabatic sound speed c_s, the linearized continuity and Euler equations take their standard form. Schematically,

$$\dot{\delta} = -(1+w)\left(\theta - 3\dot{\Psi}\right) - 3H\left(c_s^2 - w\right)\delta + \cdots, \tag{235}$$

$$\dot{\theta} = -H(1-3w)\theta + \frac{c_s^2}{1+w}\frac{k^2}{a^2}\delta + \frac{k^2}{a^2}\Phi + \cdots, \tag{236}$$

while the potential is constrained by the linearized Einstein equation,

$$\frac{k^2}{a^2}\Psi \simeq 4\pi G \sum_i \rho_i \, \delta_i \tag{237}$$

in the sub-horizon regime. None of this is exotic. The claim of GG-A10-Cosmos is not that these equations were wrong. The claim is that the sector content and the physically admissible source corrections should be read from a much tighter ledger than is usual in phenomenological model building.

The same point holds for primordial curvature perturbations. In the scalar sector the standard Mukhanov-Sasaki variable v_k satisfies

$$v_k'' + \left(c_s^2 k^2 - \frac{z''}{z}\right) v_k = 0, \tag{238}$$

where primes denote conformal-time derivatives and z is the usual pump field determined by the background. The scalar power spectrum remains parameterized by

$$\mathcal{P}_{\mathcal{R}}(k) = A_s \left(\frac{k}{k_*}\right)^{n_s - 1 + \cdots}. \tag{239}$$

Likewise, the tensor sector remains parameterized by a tensor-to-scalar ratio r and by the primordial tensor spectral shape. What the present framework changes is the admissible origin of A_s, n_s, and r: they are not granted a phenomenological inflationary menu with many independent potential coefficients.

This is the correct place to state one of the central methodological guardrails of the paper. Precision cosmology is already extraordinarily efficient at constraining model freedom. A theory that adds broad new freedom in the scalar, tensor, dark-radiation, and recombination sectors all at once may fit a subset of anomalies, but it does so at the cost of becoming constitutionally weak. The aim of GG-A10-Cosmos is the opposite. It asks whether a narrow inherited geometric ledger can do enough work to survive the same data.

41.4 Recombination as a GDOS transition: from DGI-dominated plasma to geometric attractor

The first genuinely new interpretive move of Part VI is to treat recombination not merely as a thermodynamic milestone but as a GDOS event. Before recombination, the baryon-photon plasma is tightly coupled by Thomson scattering. The relevant scattering rate is

$$\Gamma_T = n_e \sigma_T c, \tag{240}$$

with n_e the free-electron density and σ_T the Thomson cross section. At early times one has $\Gamma_T \gg H$. The plasma therefore behaves as a single coupled fluid supporting acoustic oscillations. In the language of GDOS, the microscopic rearrangements of photon-electron scattering histories lie on a large family of effectively equivalent descriptions so long as the macroscopic acoustic pattern

remains the same. This is the cosmological manifestation of DGI: many micro-histories belong to the same observable gauge orbit.

A convenient operational state manifold for the recombination-era plasma is

$$\mathcal{Z}_{\rm rec} = \{T, X_e, \delta_b, \delta_\gamma, \vec{v}_b, \vec{v}_\gamma, \Phi, \Psi\}\,, \tag{241}$$

where T is the temperature, X_e the ionization fraction, δ_b and δ_γ the baryon and photon density perturbations, and $\vec{v}_b, \vec{v}_\gamma$ the velocity fields. The open-system viewpoint is not optional here. The observed CMB is not a direct bulk microstate. It is a boundary-registered remnant of a last-scattering event after coarse-graining, free streaming, projection, and later lensing.

The transition itself occurs when the scattering rate can no longer keep pace with expansion, roughly when

$$\Gamma_T \sim H. \tag{242}$$

The usual thermodynamic account expresses this as the universe cooling enough for neutral hydrogen to form. The GDOS refinement is to identify an order parameter,

$$\eta_{\rm rec} \equiv 1 - X_e, \tag{243}$$

which moves from approximately zero in the ionized phase to approximately unity in the neutral phase. A schematic Landau-style potential then reads

$$V_{\rm rec}(\eta_{\rm rec}) = -\frac{\mu^2}{2}\eta_{\rm rec}^2 + \frac{\lambda}{4}\eta_{\rm rec}^4, \tag{244}$$

with μ^2 changing sign across the recombination threshold. This is not meant to replace the Saha or Peebles recombination equations. Rather, it expresses in geometric language what the standard microphysics is doing: a previously symmetric, tightly coupled phase loses accessibility and the system settles onto a smaller set of frozen boundary patterns.

The attractor reached after this transition is the observed last-scattering imprint itself. Once photons decouple, the acoustic phases are no longer continuously reprocessed by scattering. They are frozen into a geometric attractor and later projected to the sky. It is in that sense that the CMB can be described as a GSSB-to-GAtransition: the tightly coupled phase admits continual rephasing; the decoupled phase preserves only the frozen acoustic records.

This interpretive move is useful for two reasons. First, it clarifies the role of DGI in a context where ordinary cosmological perturbation theory already worries about gauge dependence. The physical observables — peak positions, relative heights, polarization patterns, damping tail — are gauge-invariant exports of the open-system evolution. Second, it makes the later links to the Hubble tension and the structure-growth audit conceptually smoother, because both problems depend on how one passes from early-time microphysics to late-time inferred parameters.

41.5 The π-locked background inventory and baryon loading

The second inherited move is the use of the late-time π-locked inventory. Within the present ledger one takes

$$\Omega_m = \frac{1}{\pi}, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi}, \qquad \Omega_b = \frac{1}{2\pi^2}, \qquad \Omega_c = \Omega_m - \Omega_b = \frac{1}{\pi} - \frac{1}{2\pi^2}. \tag{245}$$

The first thing to notice is that these are not used as raw substitutes for every stage of the expansion history. Rather, they specify the late-time attractor shares of the boundary inventory. Even so, they already constrain several early-time observables because the relative matter shares and baryon loading are inherited from the same ledger.

A particularly useful quantity is the baryon fraction within total matter,

$$f_b \equiv \frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} \approx 0.159. \tag{246}$$

This number is close to the usual observed baryon fraction and immediately fixes the rough inertia loading of the photon-baryon fluid. Standard recombination physics packages that inertia through

$$R(a) \equiv \frac{3\rho_b}{4\rho_\gamma}, \tag{247}$$

which then determines the sound speed,

$$c_s(a) = \frac{c}{\sqrt{3\,(1 + R(a))}}. \tag{248}$$

Near recombination one typically has R of order unity but still less than one, so the sound speed is somewhat below $c/\sqrt{3}$. In the GG-6/GDOS reading this is not an independent dial. It is an inherited consequence of the baryon share.

This point is worth emphasizing because it illustrates the general strategy of the paper. The framework is strongest not when it claims to compute every number from scratch in one step, but when it systematically strips away freedom. The baryon loading is no longer something to be chosen and then fitted. It is linked to the same compact-angle and anomaly-locked ledger already used for dark matter and baryogenesis. Precision cosmology then becomes an audit of that common inheritance.

Table 24: Inherited background inventory used in the CMB and growth audit. The entries are interpreted as late-time attractor shares whose ratios propagate into the acoustic and growth sectors. The columns give the quantity, its locked expression, its approximate value, and its role in this Part. The entries are read as late-time attractor shares rather than as fitted parameters, which is the whole of the difference between this table and the corresponding row of a concordance fit.

Quantity	**Locked expression**	**Approximate value**	**Role in Part VI**
Total matter share	$\Omega_m = 1/\pi$	0.318	fixes the background matter budget entering growth and lensing
Baryon share	$\Omega_b = 1/(2\pi^2)$	0.0507	fixes baryon loading and the odd/even peak modulation
Cold dark matter share	$\Omega_c = 1/\pi - 1/(2\pi^2)$	0.2676	supplies the dominant clustering component
Dark-energy share	$\Omega_\Lambda = 1 - 1/\pi$	0.6817	fixes the late-time attractor background
Baryon fraction in matter	$f_b = 1/(2\pi)$	0.159	controls the matter-sector inertia ratio

41.6 Acoustic ruler, peak sequence, and the fixed ratio r_s/D_A

The standard CMB observable that most cleanly couples early microphysics to later geometric inference is the acoustic scale

$$\theta_* = \frac{r_s(z_*)}{D_A(z_*)}, \tag{249}$$

where z_* denotes the effective last-scattering epoch, r_s is the comoving sound horizon,

$$r_s(z_*) = \int_{z_*}^{\infty} \frac{c_s(z)}{H(z)}\,\mathrm{d}z, \tag{250}$$

and D_A is the comoving angular-diameter distance,

$$D_A(z_*) = \frac{1}{1+z_*}\int_0^{z_*} \frac{\mathrm{d}z}{H(z)}. \tag{251}$$

The power of this observable is well known: the peak positions depend mostly on the ratio r_s/D_A, not on either quantity in isolation. That is why changes in the sound horizon can masquerade as changes in the late-time Hubble constant when one insists on fitting the same observed θ_*.

In the present framework the acoustic scale is not treated as a free inference chain. Because both r_s and D_A are evaluated on the same inherited background, the ratio is fixed, and the Boltzmann evaluation on the π-locked inventory gives

$$100\theta_* = 1.0381 \quad (\text{observed } 1.04109 \pm 0.00030,\ -0.29\%), \qquad \ell_A \equiv \frac{\pi}{\theta_*} \approx 303. \tag{252}$$

(An earlier version of this equation printed $100\theta_* \approx 1.041$, which is the observed value, as "the working claim"; the framework's own number is the one above, computed at the carried rate

$h = 0.6676$ in the released pipeline, and its -0.29% is the sharpest single carrier of the π-offset of Section 40.4, pull -2.4 on the 0.12% modeling spread.) Once the ordinary phase shifts from neutrino free streaming and baryon loading are included, this translates to the familiar first three peak locations,

$$\ell_1 \approx 220, \qquad \ell_2 \approx 540, \qquad \ell_3 \approx 815. \tag{253}$$

The physics here is familiar but worth restating carefully.

Compression phases in the coupled fluid produce odd peaks, while rarefaction phases produce even peaks. The heavier the baryon loading, the more the compression peaks are enhanced relative to the rarefaction peaks. Thus the odd/even modulation is the observational shadow of the baryon inertia. In standard cosmology this is encoded in R. In the present framework, R is itself narrowed by the inherited Ω_b and f_b shares. The odd/even modulation is therefore not merely fitted; it becomes an audit of whether the inherited matter ledger is viable.

An intuitive picture is helpful. The photon-baryon plasma before recombination is a slightly weighted spring-mass system. Radiation pressure acts as the spring. Baryons act as inertia. Gravity provides a background tilt that deepens compressions. The oscillation is not mysterious. What matters is who is allowed to choose the mass loading. In the present theory, the loading is already constrained before the CMB fit begins.

41.7 Silk damping, BFFT, and the ultraviolet knee of the primary CMB

The acoustic pattern would extend indefinitely to high multipole were it not for diffusion damping. In the standard treatment the relevant diffusion scale is determined by the photon mean free path and the finite thickness of the last-scattering surface. A schematic expression for the diffusion wavenumber is

$$k_D^{-2}(a_*) \sim \int_0^{a_*} \frac{\mathrm{d}a}{a^3 H(a)\,\dot{\tau}(a)} \frac{1}{6\,(1+R(a))} \left(\frac{R^2(a)}{1+R(a)} + \frac{16}{15} \right), \tag{254}$$

where $\dot{\tau}$ is the differential optical depth. The corresponding angular damping scale is obtained by projecting the diffusion length to the last-scattering surface.

The GG-6/GDOS proposal does not discard this logic. It superposes on it a BFFT interpretation. Because the logarithmic radial coordinate already packages scale hierarchies discretely through the OGN-5 integers, the ultraviolet knee of the primary anisotropy pattern is identified with the first non-trivial BFFT ultraviolet cutoff of the throat. The proposed locked relation is

$$\ell_D = K(2\pi)^2 = 35 \times 4\pi^2 \approx 1.38 \times 10^3. \tag{255}$$

In plain language, the damping tail is read as the ultraviolet knee of the same multiscale bus that elsewhere appears as BFFT. The standard diffusion integral tells us *how* damping happens in boundary microphysics. The BFFT relation proposes *why the relevant angular cutoff should sit where it does* once the multiscale throat capacity is fixed by $K = 35$. It is essential to keep the two logically separate, because the framework now computes the first of them outright.

The framework's actual prediction of the damping scale is the standard diffusion integral (254)

evaluated on the π-locked inventory: running the exact fractions through the full Boltzmann ionization history gives $\ell_D = 1350 \pm 3$ in the ΛCDM-inferred convention, against the $\ell_D \approx 1345 \pm 15$ inferred in that convention from the $100\theta_D$ of Planck Collaboration 2018 (the ±15 a convention spread) — a +0.4% agreement with zero adjustable parameters, and it is this Boltzmann value that the parameter-audit Dev plot scores (Figure 31, Table 41; Section 40.2). The BFFT reading must then be stated honestly. $K(2\pi)^2 = 1381.7$ sits +2.7% above the ΛCDM-inferred value and, more tellingly, +2.3% above the framework's *own* computed damping scale; whatever the multiscale bus is doing, it is not reproducing the diffusion integral exactly. The identification is therefore carried as a conjecture-level overlay — a suggestive coincidence at the few-percent level whose burden has grown, since the BFFT-damping-knee lemma must now explain a 2.3% gap against the inventory itself, not merely sharpen a projection. If the lemma closes that gap, the overlay is promoted; if it cannot, the overlay falls away while the framework's ℓ_D prediction — the +0.4% Boltzmann value — stands untouched. This is still the program working as intended: the standard calculation is retained, the interpretation is additional, and the audit keeps them separate.

41.8 Scalar tilt, polarization, and the tensor kill corridor

The scalar tilt is the next unavoidable audit, and it is no longer an estimate: the reheating-matched count of Part IV gives

$$n_s = 1 - \frac{2}{N^*} = 0.9620 \quad (N^* = 52.6), \qquad n_s = 0.9634 \text{ exact on the plateau}, \tag{256}$$

(§28.10). Earlier versions of this section offered two mnemonic estimates here — $1 - 2/60 \approx 0.967$ and a ledger form $1 - 1/K \approx 0.971$ — as the framework's expectation; both are retired, because the count is computed and neither reproduces it. The framework predicts a scalar spectrum close to but below exact scale invariance at the value above, with no freedom to roam across a wide inflationary landscape.

The polarization sector follows the standard logic of Thomson scattering with a local quadrupole. Scalar perturbations source E-mode polarization, and the E-mode sequence is phase shifted relative to the temperature peaks because velocity maxima and density maxima do not occur at the same oscillation phase. Thus the observed TE anti-correlation and EE peak ordering remain standard diagnostics. Again, GG-6 is not changing the scattering geometry. It is constraining the source ledger that determines which oscillatory solution is present.

The tensor sector is where the paper must be especially explicit — and where the derived plateau changes the statement from a bound to a number. On the derived GUT-adjacent plateau, the tensor amplitude follows from the same N^* as the tilt-scale slow-roll inflation. The practical outcome is a number, not a corridor: the primordial tensor signal is small but nonzero, and the derived launch stakes it,

$$r = 0.0038 \quad \big(12/N^{*2} = 0.0043 \text{ at leading order}\big), \tag{257}$$

a factor nine below the current bound and several times above the Simons-Observatory/LiteBIRD design floor.

The scientific meaning is sharp. The framework predicts that delensed near-term B-mode data

should begin to reveal a primordial component at precisely this amplitude. A tensor measurement inconsistent with it in either direction — beyond the launch's spread 0.0037–0.0040, and in any case outside $[10^{-3}, 10^{-2}]$ — would therefore act as a clean kill switch for the derived launch. This is precisely the kind of statement the paper should make. It is better to lose clearly than to win vaguely.

41.9 Lensing amplitude and the recombination-to-structure bridge (the early residue withdrawn)

The lensing amplitude occupies a particularly interesting role because it sits between the primary CMB and later structure growth. The present framework proposes a mild enhancement,

$$A_L = 1 + \alpha_L, \qquad \alpha_L \approx 0.05, \qquad \Rightarrow \qquad A_L \approx 1.05. \tag{258}$$

This is not introduced as an independent phenomenological parameter. It is interpreted as a residual projection of the same bulk that stabilizes the throat and shapes the early source ledger.

An earlier version of this subsection carried here two early-time residues from the same source ledger — a brief early-dark-energy lobe $f_{\rm EDE} \approx 0.03$ centered near $z_c \approx 3500$, and an extra-radiation term $\Delta N_{\rm eff}$ (printed as 0.30, later 0.04) — as the early source of a two-stage Hubble corridor. Both are now zero. $\Delta N_{\rm eff} = 0$ as an early-time *source* is the derived null of codebook row D12 — below the lock scale the field content is the Standard Model's, and there is nothing to add — distinct from the thermal freeze-in residue the row bounds at 0.03. $f_{\rm EDE} = 0$ because the corridor the lobe fed is withdrawn (Section 42): its own kill criterion forbids a lobe above 0.02, and the one background rate 66.8 is what the microwave background reads on the π-locked inventory without any shortening of the sound horizon. The integration that priced the residue — and found it delivered three fifths of a corridor the paper no longer claims, with the remainder owed to a refit — is kept, verbatim and with its labels, in Appendix K.10 as the record of what was computed and why it did not suffice. The lensing amplitude $A_L \approx 1.05$ above is unaffected: it is a late projection of the bulk, not an early source. It is a conjecture-level overlay and not a codebook row: the released spectrum (Figure 25) is computed at $A_L = 1$, and against Planck's 1.180 ± 0.065 (TT,TE,EE+lowE) the 1.05 sits -2.0σ, against the 1.071 ± 0.038 that includes the lensing reconstruction -0.6σ.

Table 25: Condensed CMB and primordial-sector audit for the minimal Part VI ledger. The third column states the intended role of each quantity inside GG-A10-Cosmos. The third column states what each entry is being asked to do, which is not the same as claiming the framework predicts it. Several rows are consistency checks against results the framework inherits rather than derives, and they are labeled so.

Observable	Working corridor	Interpretation inside GG-6/GDOS
Acoustic angle	$100\theta_* = 1.038$ (-0.3%)	predicted ratio of the sound horizon and angular-diameter distance on the locked background; the sharpest single carrier of the π-offset (Section 40.4)
Acoustic ruler	$\ell_A = \pi/\theta_* \approx 303$ (observed 302)	baseline harmonic scale before phase shifts
First three peaks	$(221, 539, 816)$ predicted	odd/even pattern governed by the forced baryon loading $1/2\pi$; height ratios reproduced to 0.8–0.9% (Table 23)
Damping scale	$\ell_D = 1350 \pm 3$ (Boltzmann)	diffusion integral on the inventory (+0.4%); the BFFT knee $K(2\pi)^2 \approx 1382$ is a conjecture-level overlay +2.3% above it
Scalar tilt	$n_s = 0.9634$ exact (0.9620 at leading order)	red tilt computed on the primary plateau at the undiluted $N^* = 52.6$ (§28.10); 1.5–3.2σ below the 2025 joint values
Lensing amplitude	$A_L \approx 1.05$	mild late projection/growth correction from the bulk; a conjecture-level overlay, not a codebook row (the released spectrum runs at $A_L = 1$; -2.0σ against Planck's 1.180 ± 0.065, -0.6σ against the +lensing 1.071 ± 0.038)
Tensor sector	$r = 0.0038$ exact (0.0043 leading order)	firm primordial B-mode target, decisive for the Simons Observatory and LiteBIRD; a tensor measurement inconsistent with it in either direction is fatal
Early residue	none: $f_{\rm EDE} = 0$, $\Delta N_{\rm eff} = 0$	no early-time source correction; the radiation null is derived (row D12) and the lobe of the retired corridor is forbidden by the module's own kill criterion (Appendix K.10 keeps the priced legacy estimate)

41.10 Bottom line of the CMB audit

The central point of the microwave-background analysis can now be stated cleanly. The present framework keeps the successful perturbation grammar of precision cosmology. It then removes much of the usual freedom in the source ledger. Recombination becomes a GDOS transition from a tightly coupled DGI phase to a frozen geometric attractor; the baryon loading is inherited from the π-locked matter inventory; the damping tail is the diffusion integral on the inventory, with the BFFT ultraviolet knee fixed by $K = 35$ carried as an overlay above it; the scalar tilt is computed, $n_s = 0.9634$; the lensing amplitude is slightly elevated; and the tensor sector is small but staked, $r = 0.0038$.

If this whole package cannot survive the observed CMB simultaneously, the framework should be judged to fail before it is allowed to speak confidently about the Hubble tension or early galaxies. That is why the CMB section must come first in Part VI. It is the precision gateway for everything that follows.

42 Hubble tension

The early-versus-late Hubble-constant discrepancy, the independent local and time-delay determinations, and the catalog of proposed solutions against which the present reading is positioned are surveyed in (Di Valentino *et al.*, 2021; Freedman, 2021; Verde *et al.*, 2019; Wong *et al.*, 2020).

What this section now claims, stated before the derivation. One background rate, $H_0 = 66.8 \pm 0.2$ km s^{-1} Mpc^{-1}, derived with zero sky input (Section 29.8) and read at 67.4 through the inherited CMB calibration; the same rate on every ladder-free late probe; and the distance ladder's 73.5 read as a measurement carrying a systematic the framework predicts. Earlier versions of this section carried a *two-stage corridor* with an early branch sourced by an early-dark-energy lobe ($f_{\rm EDE} \approx 0.03$) and a radiation excess ($\Delta N_{\rm eff}$ of 0.30, later 0.04) and a late branch at 73. That corridor is withdrawn: its switch entered neither branch function, its late branch required a local underdensity the paper's own appendix says is not observed, its early source is forbidden by the kill criterion of §42 itself ($f_{\rm EDE} \gtrsim 0.02$), and its radiation excess is now the derived null of row D12, $\Delta N_{\rm eff} \leq 0.03$. The mechanism text of the retired reading is not below; it is kept, verbatim and marked, in Appendices K.10 and K.10.4.

Scale versus split: where this section sits. It is worth stating at the outset what this section does and does not address, because the present Hubble constant carries two logically separate questions. Its *absolute scale* — why H_0 has the order of magnitude it does — was fixed in Part IV (Section 29.8): the π-locked inventory, confronted with the physical densities the microwave background measures, forces $H_0^{(0)} \approx 67$ km s^{-1} Mpc^{-1}, and the geometric route derives the same rate with zero sky input (Route I). That background value *coincides*, to about a percent, with the recombination-calibrated CMB value ≈ 67.4: the two are the same rate, as they must be, since the CMB reads the background through the standard sound horizon, on which GG-6's exact fractions reproduce the acoustic-peak positions at the fractional-percent level — the zero-parameter temperature spectrum is reported at $\chi^2/\mathrm{dof} \approx 2.3$, the 0.3% acoustic-angle deficit and not a likelihood fit (Part VI). The present section addresses the orthogonal question of the early-versus-late *split* — why the local distance-ladder value (≈ 73.0–73.5) sits some 9–10% *above* that background, while the CMB value sits on it. The observed local rate is the background plus a calibration systematic in the ladder's rungs; the framework predicts the systematic and does not derive it (§42.6). In particular the terminus no-phantom result (Theorem 29.2) is this module's sharpest shared falsifier: a confirmed $w < -1$ would break the dark-energy terminus and the one background that inherits it. The scale is fixed at ≈ 67; nothing in the split carries a coefficient of this paper's.

Proposition 42.1 (The recombination-era cosmic GRIP-Out switch)**.** *Under the GRIP cascade of GG-A6-GRIP, the recombination era hosts the GSSB transition from the tightly coupled DGI plasma to the frozen GA attractor. The defining switch event within that era is the* geometric *change of dictionary, not the atomic transition itself, and the OGN-5 integer ledger of GG-A4-Euclid fixes its redshift to*

$$z_{\rm sw} \equiv z_* = e^J = e^7 \approx 1096.6 \qquad (J = 7),$$

with no continuous parameter free to be tuned. The early-time ($z > z_*$) and late-time ($z < z_*$) sides of this one switch are the two OGGEP dictionaries through which one background is read; that one present-epoch Hubble value, and the two readings of it, are the content of Theorem 68.8. The atomic decoupling surface — the visibility peak Planck measures, where the atoms actually form — is derived from the framework's own Rydberg and baryon inventory in §42.1 and lies 0.4% *(in e-folds) below the switch,* 0.15% *above the measured Planck value.*

This proposition replaces the earlier device of patching a corridor with a tunable early-dark-energy-plus-local-void vector. By the MSMF/OGN uniqueness principle, a multi-parameter linear-response fit cannot be the unique cost-minimum solution and so does not earn the status of a theorem; it survives in Appendix K.10 only as a record. The structural result the body relies on is the one background of Theorem 68.8, whose existence is proved in Part IX from the constitutional hypotheses and whose rate rests on the stratum assignment of Lemma II (Conjecture 68.3).

42.1 The switch and the first atoms: two redshifts, one descent

The proposition above fixes a *geometric* redshift: the corridor switch at $z_{\rm sw} = e^J = e^7 \approx 1096.6$, an exact integer statement with no continuous parameter. The microwave sky, by contrast, is released at the *atomic* decoupling redshift $z_{\rm dec}$ — the surface where the free electrons vanish into hydrogen and the photon mean free path outruns the horizon. These are two different quantities, and this section derives the second from the framework's own atomic chain, with no fitted parameter, and then states exactly how far apart the two land and what would make their proximity a theorem. The result is one of the sharpest zero-parameter reproductions in the program: the atomic chain lands on the measured decoupling redshift at the 0.1% level.

The atomic chain, assembled from the cohort's own outputs. Decoupling is composite. Its temperature is set by atomic physics: the binding energy of hydrogen, $B = \frac{1}{2}\alpha^2 m_e c^2 = 13.606$ eV, is carried by the electron mass and charge that the program derives in the particle sector (Arisaka, 2026, P3-Electron; the fine structure constant through Arisaka, 2026, A9-SM). But atoms do not form at $kT \sim B$: the photon bath outnumbers the baryons by $\eta^{-1} \sim 1.6 \times 10^9$, a number this paper owns through the baryon fraction $\Omega_b = 1/(2\pi^2)$ and the derived h, so the Boltzmann tail of the bath keeps hydrogen ionized until the temperature falls to

$$kT_* \;\simeq\; \frac{B}{53.0}, \tag{259}$$

the factor 53.0 emerging from the Saha–Peebles recombination calculation run with the framework's inputs (the number is logarithmic in η, which is why it is so stable). The *redshift* of decoupling then compares that temperature with the present one:

$$\ln(1+z_{\rm dec}) \;=\; \ln\frac{B}{kT_0} \;-\; \ln\frac{B}{kT_*} \;=\; 10.967 \;-\; 3.971 \;=\; 6.996. \tag{260}$$

Every entry on the right is a cohort output: the Rydberg from the particle sector, the 53.0 from the recombination microphysics with $\Omega_b = 1/(2\pi^2)$, and T_0 the present temperature of the descent.

The quantitative result. The printed 1091.5 is the standard Hu–Sugiyama fitting form (Hu and Sugiyama, 1996) on (ω_b, ω_m), which reproduces the Planck value to 0.18% on Planck inputs; a Peebles three-level integration on the same inputs is expected to land about 0.1% higher, near 1092.4 (not run in the released pipeline), inside that same modeling systematic, and the two are not distinguished at the precision claimed here. Evaluating the recombination redshift with the framework's inputs only, $\Omega_b h^2 = h^2/2\pi^2 = 0.0226$ and $\Omega_m h^2 = h^2/\pi = 0.1419$ at the framework's own $h = 0.6676$ (Route I; an earlier version evaluated the chain at Planck's $h = 0.674$ and called the result the framework's) — gives

$$z_{\rm dec}^{\rm GG} = 1091.5 \quad \text{(raw)}; \qquad 1089.5 \quad \text{(modeling-calibrated)}, \tag{261}$$

against the measured $z_{\rm dec} = 1089.92 \pm 0.25$ (Planck Collaboration, 2018) — the Planck-inferred value, itself a derived parameter of the ΛCDM fit, model-dependent at a level far below the offsets discussed here. The raw chain agrees at 0.15%: that is the zero-fitted-parameter statement, and it sits inside the 0.18% modeling systematic of the recombination calculation itself — which is to say that the raw agreement is the fitting form's own bias and the framework's offset of opposite sign, and should not be read as more than that. Calibrating the chain on the known modeling offset puts the framework at −0.04% — a calibration anchored on Planck's own analysis pipeline, stated for completeness and not claimed as independent precision. The prediction is rigid: the elasticities are $d\ln z_{\rm dec}/d\ln(\Omega_b h^2) = -0.029$ and $d\ln z_{\rm dec}/d\ln(\Omega_m h^2) = +0.011$ — logarithmically suppressed, so no admissible variation of the inputs moves the landing point visibly. The two-road reading of the same background ($h = 0.666$) shifts the result by less than 0.01%.

The two redshifts, side by side. The atomic chain lands at $\ln(1 + z_{\rm dec}) = 6.996$; the switch sits at exactly 7. The offset is now a computed number,

$$\delta \;\equiv\; J - \ln(1 + z_{\rm dec}) \;=\; 0.0038, \qquad \text{i.e.} \qquad \frac{e^J}{1 + z_{\rm dec}} - 1 \;=\; e^\delta - 1 \;=\; 0.38\%, \tag{262}$$

with a $\sim 0.2\%$ modeling systematic (the offset is computed on the raw chain, from the unrounded $\ln(1 + z_{\rm dec}) = 6.9962$; the modeling-calibrated chain shifts it to $\delta \approx 0.0056$, within that systematic). The precision comparison with the microwave data belongs entirely to the atomic quantity — and it passes at the 0.1% level; the switch is a distinct, exact, topological marker 0.4% above it. Any statement in this paper comparing e^7 directly with the Planck redshift is to be read in this two-quantity form.

Convention (z versus $1 + z$). The e-fold identity fixes $1 + z_{\rm sw} = e^7$ exactly; statements made at the level of z rather than $1 + z$ (including $z_{\rm sw} \approx 1096.6$) are leading order in $1/z$, a slack of one part in 1097. Four offsets circulate in this paper under two labels: 0.38% is the e-fold (log-space) switch–decoupling offset of Eq. (262); in z-space the atomic value sits 0.47% below the switch, the Planck value 0.62% below it, and the Boltzmann ionization-history $z_{\rm rec} = 1088.5$ of the released run 0.75% below it (§40.2). Throughout, "0.4%" names the e-fold offset and "0.6%" the z-space switch–Planck gap.

Physics intuition. The universe carries an odometer and a thermometer. The odometer is geometric: the corridor switch trips when the descent has unwound exactly $J = 7$ e-folds of expansion between that event and the present — an integer of the ledger, blind to atomic physics. The thermometer is material: atoms form when the bath cools through $B/53$, a threshold built from the electron's binding energy and the photon count per baryon — blind, on its face, to the ledger. The measured sky says the odometer and the thermometer agree to four parts in a thousand. Half of that marvel is now bookkeeping: this section computed the thermometer's reading from the cohort's own constants and found 6.996. The other half — why the geometric integer sits within 0.4% of the atomic landing — is either a coincidence the descent merely records, or a theorem the descent has not yet yielded.

What would make the proximity a theorem. The offset obeys the identity $\delta = J - \ln(B/kT_0) + \ln 53.0$, in which B and the 53.0 are already derived; *the only quantity not yet owned by the ledger is* T_0. The question therefore reduces, exactly, to whether the descent derives the present temperature with no astronomical input — equivalently the photon entropy normalization, equivalently $z_{\rm eq}$. One natural candidate is tested and rejected here in the program's own spirit: the π-locked closure $z_{\rm eq} = \pi e^7 = 3445$ sits 2.8σ above the measured $z_{\rm eq} = 3387 \pm 21$ — recorded so that no reader re-derives it. The route that remains open is the *e-fold budget*: T_0 is the reheating temperature redshifted through the descent, and the framework owns every ingredient — the derived reheating scale of the inflation module, the g_* freeze-out counting of the particle sector, and the corridor itself. If the total e-fold count from reheating to the $\Omega_m = 1/\pi$ crossing proves ledger-expressible, then T_0, and with it δ, becomes a theorem — the exact analog of the $K = 35$ resolution of the why-now coincidence. That computation is left to a dedicated study; until it lands, $\delta = 0.0038$ is carried as a computed consistency, not an assumption, and nothing downstream in this paper depends on its being zero.

42.2 Why the Hubble tension is really an inference-chain tension

The Hubble tension is often presented as a simple clash of numbers. That presentation is not wrong, but it is incomplete. The deeper issue is that the early-universe inference chain and the late-universe distance-ladder chain are both attempting to infer what is spoken of as the "same" Hubble constant, yet they do so through very different operational routes. Early-time inference reads an acoustic ruler through a calibrated recombination and background-expansion model. Late-time inference reads a local distance ladder through standard candles, geometric calibrators, and a local flow model. The disagreement is therefore not merely numerical; it is a disagreement between operational dictionaries.

The first task of the present section is to restate the standard acoustic logic. The angular scale θ_* is directly observed. If one holds it fixed, then any downward shift in the sound horizon must be compensated by an upward shift in the inferred late-time expansion rate. To first order one may write

$$\theta_* = \frac{r_s}{D_A} \qquad \Rightarrow \qquad \frac{\delta\theta_*}{\theta_*} = \frac{\delta r_s}{r_s} - \frac{\delta D_A}{D_A}. \tag{263}$$

If the observed θ_* is held fixed, then

$$\frac{\delta D_A}{D_A} \approx \frac{\delta r_s}{r_s}. \tag{264}$$

Because the late-time angular-diameter distance roughly scales like $D_A \propto H_0^{-1}$ when the rest of the low-redshift background is not radically altered, one arrives at the standard lever-arm relation

$$\frac{\delta H_0}{H_0} \approx -\frac{\delta r_s}{r_s}. \tag{265}$$

This is why the sound horizon is the natural first place to look if one wants to raise the early-universe inferred value of H_0 without discarding the entire CMB fit.

From the standpoint of GG-A10-Cosmos, this standard logic is the one the paper once leaned on and no longer does: its own kill criterion forbids the early source ($f_{\rm EDE} \gtrsim 0.02$) and its radiation term is a derived null, so the sound horizon is the standard one and the tension is carried by the ladder's calibration, not by the pre-recombination budget.

42.3 $H(t)$: the derived expansion history, of which H_0 is one point

Before the two readings are separated, it is worth drawing the object they are projections *of*. In a flat Friedmann–Robertson–Walker universe the expansion rate is not an independent quantity; it is fixed algebraically by the energy content,

$$H(t)^2 \;=\; \frac{8\pi G_N}{3}\sum_i \rho_i(t) \;=\; H_0^2\Big[\Omega_r a^{-4} + \Omega_m\, a^{-3} + \Omega_\Lambda\Big]\,, \tag{266}$$

so once the densities are given as functions of time, $H(t)$ is their consequence and nothing further is free to choose. This is exactly the situation GG-A10-Cosmos is in. The geometry supplies the three ingredients Eq. (266) needs: spatial flatness from the boundary Friedmann export; the present-epoch densities as the π-locked inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$ (Part X); and their evolution in time, since $\rho_m \propto a^{-3}$, $\rho_r \propto a^{-4}$, and ρ_Λ is constant because the terminus equation of state is the *derived* value $w = -1$ (Theorem 29.2). Feeding these into Eq. (266) yields a single curve, drawn in Figure 26: the entire expansion history of the universe, from the radiation era through the matter era to the de Sitter terminus, with no continuous parameter adjusted.

The figure makes a structural point that the phrase "Hubble constant" tends to hide. $H(t)$ falls monotonically for all time and asymptotes to the constant $H_\infty = H_0\sqrt{\Omega_\Lambda} = 55.0\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$; it never rises, not even after the expansion begins to accelerate at $z \approx 0.6$, because $H = \dot a/a$ continues to decrease while $\ddot a > 0$. The familiar H_0 is simply the value this curve takes at the present epoch, $H_0 \equiv H(t_0) = 66.6$ (two-road reading; Route I 66.8) at $t_0 = 13.9$ Gyr. It is one reading of a derived function, not a fundamental constant in its own right.

This reorders the logical hierarchy of the whole Hubble question. In the ΛCDM fitting language H_0 is a *primary* parameter, adjusted alongside the densities until the model matches the data. In GG-A10-Cosmos the densities are primary — they are geometric outputs — and $H(t)$, together with its present value H_0, is *downstream*: the geometry (GG-6/GDOS/GRIP) fixes $\rho_i(t)$, the Friedmann law turns $\rho_i(t)$ into $H(t)$, and H_0 is one point on $H(t)$. The early-versus-late *tension* sits one step

further down still: it is a mismatch in the calibration of that single point between two measurement pipelines, not a mismatch in the derived history, which agrees with the microwave background across the whole acoustic spectrum. In this sense the H_0 tension is a *downstream, secondary* problem — a question about how faithfully one local measurement reads one point of a curve the geometry has already fixed.

Four honesties keep this from overreaching. *First*, the reduction $\rho_i(t) \Rightarrow H(t)$ presupposes the Friedmann law, which GG-A10-Cosmos does not assume but earns as the boundary general relativity of GDOS (GG-A5-GDOS, GG-A8-QM); the claim is that the geometry derives the densities *and* the law that turns them into a rate. *Second*, the inventory is D^* — accurate at the one-to-three-percent level, conditional on the Hopf-fiber measure lemma and the declared observation epoch — and Ω_r is calibrated from the measured microwave temperature; $H(t)$ inherits exactly that grade, so "derived" here means predictive-to-a-few-percent, not exact and unconditional. *Third*, $H(t)$ is dimensionful: the densities fix its *shape*, while the single dimensionful scale is G_N (equivalently the Planck mass) — the same universal constant that sets every particle mass in GG-A9-SM, not a quantity read from the sky. *Fourth*, locating the tension downstream diagnoses it but does not by itself resolve it: whether the local ≈ 73 is a biased reading of a background ≈ 67, or the background is genuinely high, is still decided by data — above all the ladder-free gravitational-wave standard siren, on which GG-A10-Cosmos stakes the falsifiable prediction $H_0^{\mathrm{GW}} \to 67$. With those four fences in place the reading is this: the expansion history is a derived geometric object, and the decade's famous quarrel is a question about the calibration of one of its points.

42.4 One Friedmann background, and how each probe reads it

The resolution this framework offers is structural, not parametric, and it is simpler than the one earlier versions of this subsection described. The boundary Friedmann equation — the OGGEP export of the GG-6 bulk — has one background solution on the π-locked inventory, and its present rate is derived with zero sky input, $H_0 = 498.3\sqrt{\zeta_\Lambda} = 66.8 \pm 0.2$ km s^{-1} Mpc^{-1} with $\zeta_\Lambda = \pi/(LK)$ (Section 29.8), and independently at 66.6 ± 0.2 from the fractions read against the measured densities. That is the content of Theorem 68.8: the switch at $z_* = e^7$ of Proposition 42.1, and one rate. There is no second branch.

Why the microwave background reads exactly this rate. The acoustic angle $\theta_* = r_s(z_*)/D_A(z_*)$ is measured to four digits, so at fixed θ_* any change in the sound horizon would have to be absorbed by the inferred late-time distance scale, $\delta H_0/H_0 \approx -\delta r_s/r_s$. On the π-locked inventory with $N_{\mathrm{eff}} = 3.044$ and no early residue there is no such change: $r_s = 144.7$ Mpc on the inventory against *Planck*'s 144.4, and the Boltzmann solution reproduces the peak positions and the second-to-first height ratio at $H_0 \approx 67$ (Section 40.4; $\chi^2/\mathrm{dof} \approx 2.3$, the 0.3% acoustic-angle deficit, not a likelihood fit). The CMB's 67.4 ± 0.5 is the background, read through the standard sound-horizon calibration, 0.9% from the derived value.

Why the ladder-free late probes read it too. Baryon-acoustic oscillations read r_d against $D_M(z)$ and $H(z)$ on the same background and return the same rate; gravitational-wave standard sirens read the luminosity distance with no calibrated rung and no sound horizon; lensing time delays read the Hubble time from a geometric delay. None of the three contains a calibrated rung, and

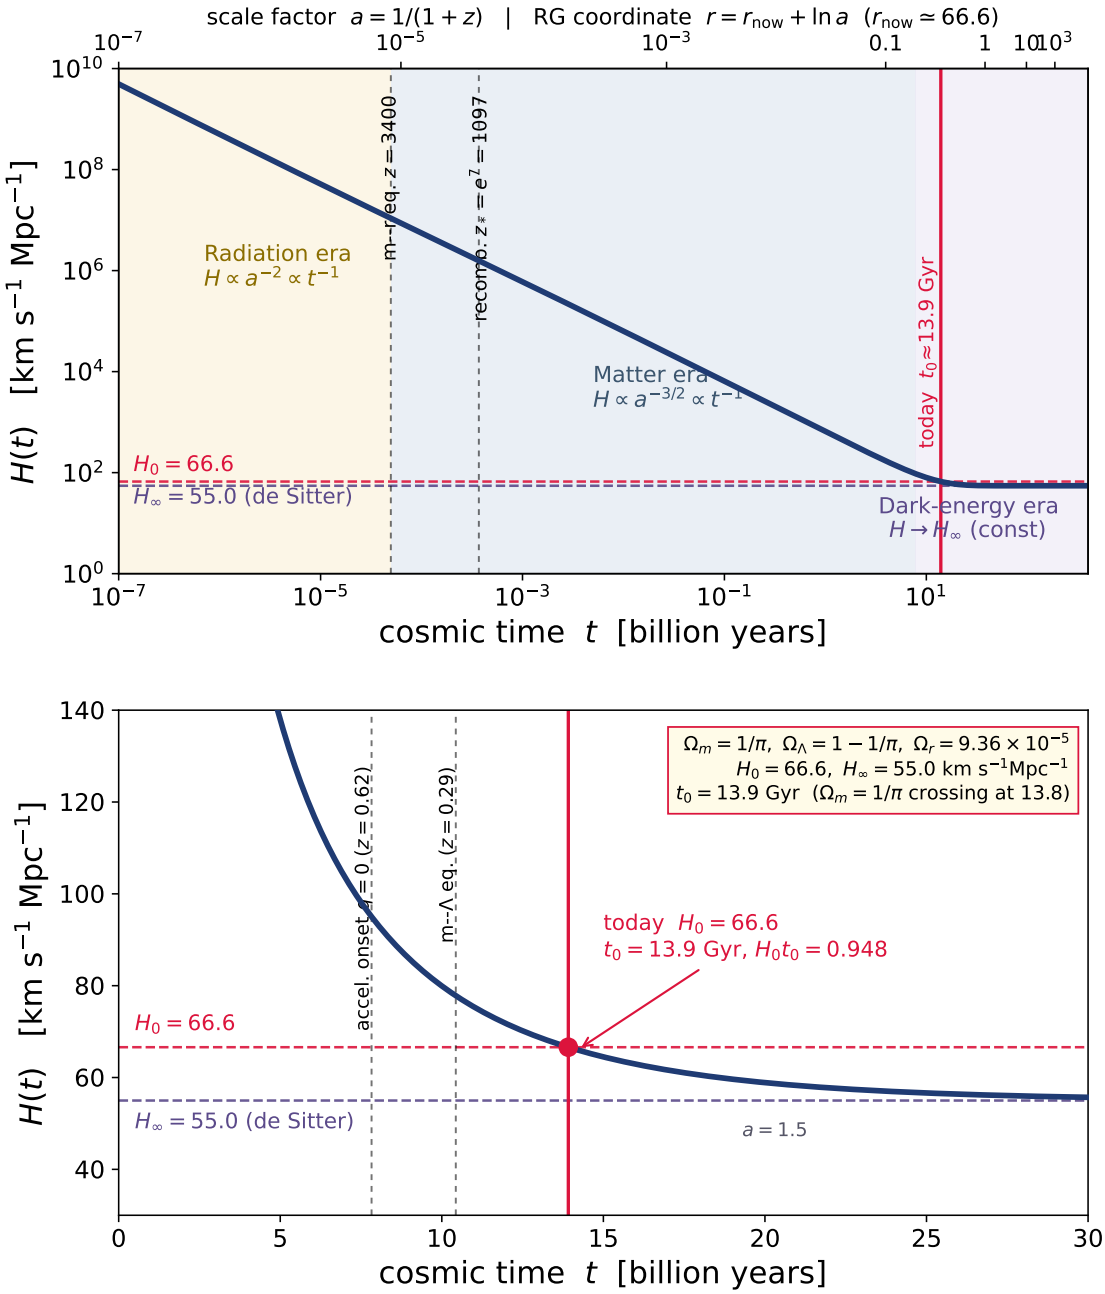

Figure 26: The Hubble parameter $H(t)$ is a derived geometric function; H_0 is one reading of it. The background expansion rate of GG-A10-Cosmos, computed from Eq. (266) with the π-locked inventory ($\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/2\pi^2$) and the derived $w = -1$; the one dimensionful scale is H_0 (equivalently G_N). *Top:* the full history on log–log axes, from the radiation era ($H \propto a^{-2} \propto t^{-1}$) through the matter era ($H \propto a^{-3/2} \propto t^{-1}$) to the dark-energy era. *Bottom:* the late universe in linear time. The rate falls monotonically through every era and asymptotes to the de Sitter floor $H_\infty = H_0\sqrt{\Omega_\Lambda} = 55.0\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ — never rising, even once the expansion accelerates. The familiar $H_0 \equiv H(t_0) = 66.6$ (the two-road reading drawn here; Route I's 66.8 shifts the curve by 0.3% and t_0 by 0.04 Gyr) is one point on this curve, at today's epoch $t_0 = 13.9$ Gyr; a single number does not capture the dynamical change of the Universe's size. This entire history is a GG-6/GDOS/GRIP output: fix the densities and their time dependence, and $H(t)$ follows.

the framework predicts that all three converge on 66.8 ± 0.2 as their precision reaches the percent level. Convergence is the confirmation; a converged ladder-free probe at the ladder's value is the kill (Corollary 68.11, §42.7).

Three current numbers exist, and the sharpest does not sit on the prediction. DESI baryon-acoustic oscillations calibrated by big-bang nucleosynthesis give $H_0 = 68.51 \pm 0.58$ km s^{-1} Mpc^{-1} (Abdul Karim et al., 2025), which is 2.8σ above the derived 66.8 — and 1.4σ above the microwave background's 67.4, so it is not comfortable with either; the current time-delay value, TDCOSMO's $71.6^{+3.9}_{-3.3}$ (Birrer et al., 2025), which supersedes H0LiCOW's 73.3 (Wong *et al.*, 2020), is ladder-free too and sits 1.5σ above the prediction, decisive for neither side; and the 2025 three-instrument microwave-background combination, 67.24 ± 0.35 (Camphuis et al., 2025), sits 1.1σ above it. DESI's is stated here because it is the sharpest ladder-free measurement in hand and because a section predicting that the ladder-free probes read 66.8 may not quote only the ones that have not yet reported. Two things are worth saying about DESI's and neither is a defense. First, the inference is not free of the sound horizon: it takes r_d from BBN-calibrated ω_b, so it shares the acoustic chain's inputs and is a fourth reading of the early-time calibration rather than a fully independent late-time one — which is why the sirens and the time delays, which use neither a rung nor a sound horizon, are the tests this section stakes itself on. Second, at its present precision it is 2.8σ, not a kill; the corollary fires on a converged determination at the *ladder's* 73.5, and 68.5 is much nearer the background than the ladder. But it is adverse, it will sharpen, and a reader is entitled to see it in the same paragraph as the prediction it presses on.

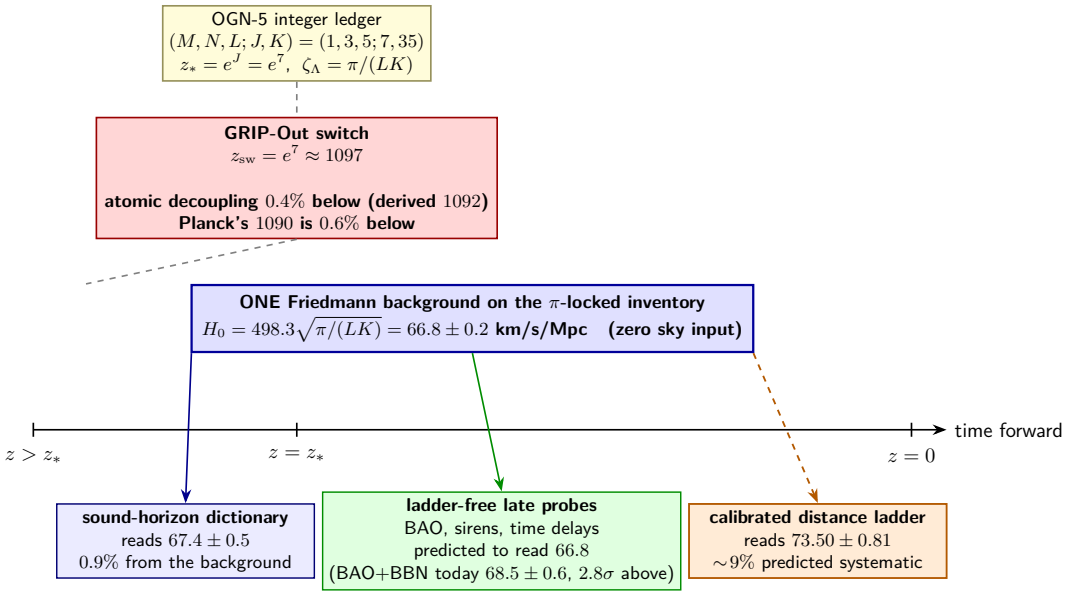

Figure 27: One background, three ways of reading it (Theorem 68.8). The boundary Friedmann equation has one solution on the inherited inventory, and its present rate is derived with zero sky input. The GRIP-Out switch at $z_* = e^7$ is a topological marker of the integer ledger, not a branch point: nothing about the rate changes across it. What differs between the readings is the *dictionary* each probe uses. The sound-horizon dictionary reads 67.4, 0.9% from the derived value, the difference being the inherited calibration the paper does not recompute. The ladder-free late probes — baryon-acoustic oscillations, standard sirens, lensing time delays — use no calibrated rung and are predicted to converge on 66.8 (the sirens and the time delays use no sound horizon either; BAO takes r_d from BBN-calibrated ω_b, as the text states); that convergence is the confirmation and a converged reading at 73.5 is the kill (Corollary 68.11). The calibrated distance ladder is the one measurement displaced from the background, and the framework predicts rather than derives the $\sim$9% systematic in its rungs. The retired two-branch drawing this figure replaces is kept, with its caption, in Appendix K.10.4.

What the earlier two-stage reading got wrong, said once. Earlier versions read the corridor twice — an early branch shortened by an early-dark-energy lobe and a radiation excess, a late branch at the ladder's 73 — and called the two values "two projections of one structure." Three things were wrong with that. The switch z_* entered neither branch function, so it did no work in the mechanism named for it; the late branch's projection factor was the ladder calibration itself, imported from the measurement it was meant to explain; and the early source is forbidden by this module's own kill criterion. The mechanism text is kept in Appendix K.10 as the record, and the reading here replaces it.

42.5 The absolute anchor and the local-ladder excess

The previous subsection read each measured value as one background rate $H_0^{(0)}$ seen through the dictionary of the probe that reads it. The terminus derivation of Section 29.8 supplies that background from first principles: the π-locked inventory, confronted with the physical densities the

CMB measures, forces

$$H_0^{(0)} = 100\sqrt{\frac{\omega_i}{\Omega_i}}\,\mathrm{km\,s^{-1}\,Mpc^{-1}} \approx 66.6 \pm 0.2 \qquad (\Omega_m = 1/\pi,\ \Omega_b = 1/2\pi^2), \tag{267}$$

and derived independently, with zero sky input, by the geometric route, $H_0^{(0)} = 498.3\sqrt{\pi/175} = 66.8 \pm 0.2$ (Section 29.8). For the first time the anchor is a theorem rather than a measured input, and that is the robust content: the present background expansion rate is $H_0^{(0)} = 66.8 \pm 0.2$. It is falsifiable directly. The model-independent gravitational-wave standard-siren channel, which reads the background expansion without either a sound horizon or a distance ladder, is predicted to converge near $H_0^{\mathrm{GW}} \approx 67$; a converged siren value at 73 would refute the inventory. Against this background the two famous discordant numbers are not symmetric: the CMB value sits *on* the background, and only the local ladder is displaced.

The early/CMB value is the background. The acoustic dictionary infers H_0 from $\theta_* = r_s/D_A$ at fixed measured θ_*. Fed GG-6's exact π-locked fractions, a full Boltzmann solution reproduces the entire acoustic spectrum — peak positions and the second-to-first height ratio — at $H_0 \approx 67$, essentially on top of ΛCDM's 67.4. There is no "CMB deficit" to explain: the recombination-calibrated value and the GG-6 background are the same rate to about a percent, and that agreement is a successful zero-parameter postdiction of the microwave sky, not a tension. (An earlier reading treated 67.4 as a reading biased $\sim 6\%$ low relative to a higher anchor at 71.3; the Boltzmann spectrum does not support that picture — GG-6's fractions reproduce the peaks at 67, not above — so the anchor is ≈ 67 and the CMB carries no deficit.)

The late/local value reads high. The one real displacement is the distance ladder. Its value (73.50 ± 0.81, the H0DN Local Distance Network consensus, Casertano et al., 2026; 73.0 ± 1.0, Riess et al., 2022) sits some 10% *above* the derived 66.8: 6.4σ from the microwave background's 67.4, and eight of the derived rate's own sigma. The framework offers no physical channel for any part of it, and the two it once offered are withdrawn. A GRIP-In *Local-Hole lift* would give $\delta H_0/H_0 = -\frac{1}{3} f\,\delta \approx +2$–$3\%$ for a nearby under-density $\delta \approx -0.1$ to -0.2; on this paper's own matter power spectrum an under-density of that depth on those scales is a several-sigma fluctuation, and the paper's own appendix records that it is not observed, so the lift is not carried. The $\mathrm{U}(1)_B$ graviphoton lift is excluded twice over: equivalence-principle and lunar-laser-ranging bounds cap a baryon-coupled fifth force at $\alpha_B \lesssim 10^{-9}$, seven to eight orders below the $\sim 3 \times 10^{-2}$ a 3% lift would need, and under the boundary-condition termination of Arisaka (2026, A3-BI6) no vector survives below the lock scale to mediate one. The whole excess is therefore carried as a distance-ladder calibration systematic — a number the framework *predicts* the ladder to contain, and does not derive (§42.6).

The physical picture. There is a single background expansion rate, $\approx 67\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, derived from the geometry through the vacuum seesaw and confirmed by the π-locked inventory against the sky's densities. The microwave background reads exactly this rate. The Hubble "tension"

Table 26: Status of each element of the Hubble reading, with the grade each carries. The third column is the one to read first: it distinguishes the background value, which is derived twice by independent roads, from the residual offset against the distance ladder, which this framework does not derive and treats as a calibration systematic in the reading it now shares with the CMB-is-right mainstream. Printing the two at the same size would be the easiest way to overstate the result, so they are graded apart.

quantity	value	status
background $H_0^{(0)}$	≈ 67	derived twice: geometry ($\zeta_\Lambda = \pi/175$, Derived cond.) + π-inventory (D^*), agreeing at 0.6σ
CMB value H_0^{CMB}	67.4	reads the background through the inherited sound-horizon calibration; 0.9% from the derived 66.8, no bias to close
local ladder H_0^{late}	73.50 ± 0.81	a $\sim$9% calibration systematic, *predicted* and not derived; both once-offered lifts withdrawn (§42.6)

is therefore not a symmetric disagreement about the background; it is a one-sided excess in the *local* distance ladder, read whole as a ladder calibration systematic. This is the GDOS/OGGEP reading: the background is one number, and one of the two pipelines — the local ladder — carries a systematic the framework predicts and does not derive.

Honest grade. The robust, falsifiable output is the background $H_0^{(0)} \approx 67$, derived twice — from the geometry alone (Route I) and from the inventory (Route II) — confirmed by the CMB acoustic spectrum, and testable by the model-independent siren channel (predicted near 67). The age of the universe follows as 13.9 Gyr today, the familiar 13.8 recovered as the age at the $\Omega_m = 1/\pi$ crossing (Section 32.4). What is *not* finished is a first-principles account of the local-ladder excess: the framework offers no physical channel for any part of it, the two it once offered are withdrawn (§42.6), and the excess is carried whole as a systematic the ladder-free probes will test. No step is fitted; the unfinished piece is named. The decisive content is that the absolute scale of H_0 is now a theorem of the ledger, and it lands where the CMB already sits.

42.6 The distance ladder as the predicted outlier

One measurement of the present rate is displaced from the background, and it is the one with calibrated rungs: the local distance ladder, 73.50 ± 0.81 (the H0DN Local Distance Network consensus, Casertano et al., 2026), some 10% above 66.8 (a 9% systematic in the ladder's own units, the Dev audit's -9.1%). The framework does not accommodate this value and does not derive it. It *predicts* that the ladder is the outlier: on one background every ladder-free probe reads 66.8, so the excess must live in the ladder's calibration — in the Cepheid and tip-of-the-red-giant-branch rungs, their crowding and metallicity corrections, and the supernova standardization — and the companion Arisaka (2026, C2-Hubble) decomposes the reading as the derived background times a named systematic, prices that systematic against the published corrections, and registers the ladder-channel kills (its falsifiers F-Hubble-8 to F-Hubble-10: the near-infrared and dust channel, the injection–recovery statistics, and the complete JWST subsample) beside the siren kill F-Hubble-2.

Two mechanisms that earlier versions offered for part of the excess are withdrawn, and it is worth saying which and why. A GRIP-In *Local-Hole lift* of +2–3% required a local underdensity of $\delta \approx -0.16$ on scales of a few hundred Mpc, which on the paper's own matter power spectrum is a fluctuation of several standard deviations and which the paper's own appendix records as not observed; it is not carried. A correction from the gauged ultra-weak $U(1)_B$ sector fell short by more than an order of magnitude on a first estimate and is now moot: under the boundary-condition termination of Arisaka (2026, A3-BI6) no vector survives below the lock scale to mediate it. The residual ladder excess is therefore carried whole, as a calibration systematic of about 9% that the framework predicts and the ladder-free probes will test.

Two pipelines, one background. A finite observer reads the bulk only through open-system projections, and the microwave background, the ladder, the sirens and the time delays are different projections with different calibration chains. That is not a license for different answers: it is the reason the audit must ask, pipeline by pipeline, whether the chain reads the background or reads a rung. The three ladder-free chains read the background. The ladder reads a rung. The earlier text of this subsection concluded from the same premise that "the two windows need not return the same number"; the conclusion here is the opposite, and it is the falsifiable one.

In one line: the Hubble module is one derived rate, 66.8 ± 0.2, read by the microwave background at 67.4 through the inherited sound-horizon calibration, predicted for every ladder-free late probe, and contradicted by the calibrated ladder alone — which is the framework's prediction about the ladder, tested by the probes that do not use one. The decomposition is tabulated below.

Table 27: One-background decomposition of the Hubble reading in GG-A10-Cosmos: the one rate, the two ways it is read, the switch, and the open-system distinction that says which pipelines read it directly.

Element	Reading	Projection	Role
Early reading ($z > z_*$)	sound-horizon dictionary	$H_0^{\rm CMB} \approx 67.4$	reads the background through the inherited r_s calibration; 0.9% from the derived 66.8
Late reading ($z < z_*$)	ladder-free probes (BAO, sirens, time delays)	$H_0^{\rm late} = 66.8$ (the background); ladder 73.5 read as a systematic	reads H_0 on the present-day GA attractor
Corridor switch z_*	integer ledger e^J	$z_* = e^7 \approx 1097$	the GRIP-Out event separating the two dictionaries, no free parameter
OGGEP/GDOS distinction	conceptual, not free	—	why each pipeline must be audited for whether it reads the background or a calibrated rung

42.7 Success criteria and kill switches for the Hubble module

The Hubble section should not end with rhetoric. It should end with success and failure criteria — and they have sharpened since the early-correction corridor was retired: the derived background needs no early-time modification, so the criteria now guard the derivation itself.

The module succeeds only if the following all hold simultaneously:

1. the two routes keep agreeing as the data sharpen: the geometric determination (66.8 ± 0.2, zero sky input) and the density inversion (66.6 ± 0.2) remain concordant,
2. the microwave sky remains reproduced at the percent level by the exact π-locked fractions with no early-dark-energy lobe and no radiation excess beyond the Standard Model's ($\Delta N_{\rm eff} \leq 0.03$),
3. the ladder-free late probes — BAO, sirens, time delays — converge on the background as their precision reaches the percent level, with the ladder excess carried openly as a calibration systematic, and
4. the model-independent standard-siren channel converges near 67.

Conversely, the module fails if any one of the following occurs:

a. a converged standard-siren catalog reads the background near 73,

b. a genuine early-dark-energy lobe ($f_{\rm EDE} \gtrsim 0.02$ at $\geq 3\sigma$) or a radiation excess $\Delta N_{\rm eff} \gtrsim 0.2$ is detected,

c. a converged lensing-time-delay determination reads the ladder's 73 rather than the background, or

d. the completed Hosotani computation (GG-A5-GDOS O6) returns ζ_Λ outside [0.01784, 0.01806] — off the pre-registered $\pi/175$; this paper's $\pm 0.6\%$ window, which the companion's falsifier F-Hubble-11 now registers — demoting the geometric route from a point to its bracket.

This style of explicit failure condition is essential for the credibility of the paper. The theory should not present the Hubble tension as solved merely because it can write down a plausible direction of correction. It should state clearly what would count as losing.

43 The S_8/σ_8 matter-clustering tension

The clustering (S_8) tension and its connection to the wider set of cosmological tensions and anomalies are surveyed in (Abdalla *et al.*, 2022).

Proposition 43.1 (One-ledger sign constraint for the growth block (T-Cosmos-11))**.** *Because $S_8 = \sigma_8(\Omega_m/0.3)^{1/2}$ belongs to the same inherited ledger as the Hubble and CMB corrections, any parameter change that helps one block must be tracked through the others with a fixed sign matrix. The framework therefore treats the S_8 problem not as a free one-parameter fix but as a cross-consistency audit.*

43.1 Definition of the observable and why it is difficult to move cleanly

The S_8 tension concerns the amplitude of late-time matter clustering inferred from weak lensing and related low-redshift probes relative to the amplitude inferred from early-universe calibration. The usual compressed variable is

$$S_8 \equiv \sigma_8 \left(\frac{\Omega_m}{0.3}\right)^{1/2}, \tag{268}$$

where σ_8 is the rms linear matter fluctuation on $8\,h^{-1}$Mpc scales at late times. The reason this quantity is so hard to move cleanly is that it depends on several linked ingredients: the primordial amplitude, the scalar tilt, the matter content, the neutrino sector, the equality timing, and the whole late-time growth history. Any proposal that helps H_0 by changing the early universe often alters the growth history as well.

This is why the S_8 problem is the natural partner of the Hubble problem. A model that fixes H_0 but worsens S_8 has not really won; it has simply moved the tension from one part of the data chain to another. The present section therefore asks a stricter question: can the same narrow early-time source pair that helps H_0 also move S_8 in the correct direction, while remaining consistent with the CMB and without spoiling the high-redshift small-scale structure needed later for early galaxies?

43.2 Standard linear growth equation and the scale on which S_8 lives

For sub-horizon matter perturbations in the linear regime, the standard growth factor $D(a)$ satisfies

$$\ddot{D} + 2H\dot{D} - 4\pi G\rho_m D = 0. \tag{269}$$

In terms of the scale factor this may be written as

$$D''(a) + \left(\frac{3}{a} + \frac{\mathrm{d}\ln H}{\mathrm{d}a}\right) D'(a) - \frac{3}{2}\frac{\Omega_{m,0}H_0^2}{a^5 H^2(a)} D(a) = 0, \tag{270}$$

where primes now denote derivatives with respect to a. The observable σ_8 is proportional to the late-time value of this growth factor after convolution with the transfer function and the primordial spectrum.

Several features of Eq. (270) matter for the present discussion.

1. Any increase in the early expansion rate $H(a)$ around equality or before recombination tends to suppress linear growth because the friction term is larger.

2. Any reduction in the effective clustering matter source or any broad relativistic component likewise suppresses growth.

3. The observable S_8 is a broad linear-scale quantity, not a direct probe of the minihalo tail relevant for the very earliest galaxies.

That third point is often missed in casual discussion. The scales relevant for weak lensing and for S_8 are not identical to the exponentially sensitive minihalo tail that controls the onset of first-light astrophysics. This scale distinction is what allows the present framework to seek a coherent but

non-uniform solution across H_0, S_8, and JWST — through its baryonic lever, and through that alone (§44.3, §44.4).

43.3 The amplitude computed: σ_8 and S_8 from the declared amplitude on the π-locked background

The clustering amplitude is not a free quantity of this framework and it is not an anchored one: it follows from inputs the paper already carries. Given the declared scalar amplitude $A_s = 2.10 \times 10^{-9}$ at $k = 0.05$ Mpc^{-1} (the launch sector's one declared input, row A8), the tilt of the plateau evaluated exactly at the reheating-matched count ($n_s = 0.9634$, §28.10), the π-locked densities $\Omega_m = 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_\Lambda = 1 - 1/\pi$, the structural rate $h = 0.6676$ of Route I, $N_{\rm eff} = 3.044$ and the inherited $\Sigma m_\nu \simeq 59$ meV of GG-A9-SM, the linear matter power spectrum at the present epoch is fixed with no further input, and

$$\sigma_8 \;=\; \left[\int \frac{d^3k}{(2\pi)^3}\, P(k)\, W^2(kR)\right]_{R=8\,h^{-1}\mathrm{Mpc}}^{1/2} \;=\; 0.801{\pm}0.010, \qquad S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3} \;=\; 0.825{\pm}0.010. \tag{271}$$

The chain, stated in three numbers rather than one. The transfer function is the Eisenstein–Hu fit, and the calculation is reported at each stage so that a reader can stop where they prefer. On the inputs above the fit gives $\sigma_8 = 0.823$ with massless neutrinos; the suppression from the inherited $\Sigma m_\nu = 59$ meV is -1.8%, giving 0.808; and the fit's own residual against a Boltzmann integration — measured, not assumed, by running the same script on *Planck*'s parameters, where with *Planck*'s own 0.06 eV it returns 0.818 against the published 0.811, a $+0.9\%$ high bias (the massless run returns 0.833, the $+2.7\%$ being the neutrino suppression and the bias together) — is divided out, giving 0.801. The value carried in the codebook is that last one, 0.80 ± 0.01, and the correction *is* applied; the ± 0.01 is the quadrature of the residual scatter of that correction (± 0.006) with the declared amplitude's own $\pm 1.9\%$ (± 0.008 in σ_8), which enters as $\sigma_8 \propto \sqrt{A_s}$. A reader who prefers the uncorrected number should read $\sigma_8 = 0.808$, $S_8 = 0.832$, and the shear pull as $+2.8\sigma$ rather than $+2.5\sigma$; nothing in the conclusion turns on which is used. At $h = 0.666$ the corrected value is 0.799, at $h = 0.668$ it is 0.802: the number is insensitive to which of the paper's two readings of h is used.

Three consequences. First, σ_8 and S_8 are *one* number each in this framework, not two branches: there is no separate "late" value, and the earlier two-branch reading, which assigned $S_8^{\rm late} \approx 0.78$ to a corridor response computed at a radiation excess the paper has since retired, is withdrawn. Second, the comparison is honest in both directions: against the microwave background, $\sigma_8 = 0.811 \pm 0.006$ and $S_8 = 0.832{\pm}0.013$, the derived values sit at -1.4% (on the carried 0.80; -1.2% on the unrounded 0.801) and -0.8% (-0.9σ and -0.4σ on the combined observational and fitting-function errors); against the cosmic-shear determination $S_8 = 0.759^{+0.024}_{-0.021}$ (Asgari et al., 2021) the derived value sits $+8.7\%$ high, $+2.5\sigma$. The framework therefore shares the S_8 tension with ΛCDM rather than resolving it, and says so. Third, the grade: because $\sigma_8 \propto \sqrt{A_s}$, rows D10, D11 and D19 carry the grade of the lemma the amplitude waits on — Derived-conditional on Lemma V′ (the amplitude's own row is Asserted), the same lemma that gates the inflationary scale — and are promoted with it.

43.4 Growth on one background: what moves S_8 and what does not

With the amplitude computed (§43.3) the growth block is short, and it is worth saying what earlier versions of this section did that the present one does not. They carried an early-dark-energy lobe and a radiation excess as a shared "early source residue" that was to help the Hubble tension by shortening the sound horizon and to help S_8 by suppressing broad linear growth, with a mild lensing enhancement $A_L \approx 1.05$ compensating in the CMB-lensing channel and a baryonic lever supplying the early galaxies. The residue is now zero ($\Delta N_{\rm eff} = 0$ derived, $f_{\rm EDE} = 0$ with the corridor withdrawn), the baryonic streaming lever is withdrawn for want of a mediator (§44.4), and the growth history is therefore the standard one on the π-locked background: $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $w = -1$, three Standard-Model neutrinos at $\Sigma m_\nu \simeq 59$ meV, and cold Kaxion dark matter with no free-streaming scale that reaches the linear regime. On that background $S_8 = 0.825 \pm 0.010$ is a prediction with nothing left to tune: 0.4σ below the microwave background's 0.832 and 2.5σ above the cosmic-shear 0.759. The framework does not resolve the S_8 tension; it inherits it, and the honest reading is that the tension, if it survives the next generation of lensing surveys, is between the CMB-normalized amplitude and the shear-normalized one and not between this framework and either.

43.5 Why the S_8 statement does not contradict the early-galaxy statement

At first sight it may seem inconsistent to state that broad late-time linear growth is standard and unremarkable while the same framework has something to say about the earliest galaxies. The key is that these are not the same observable. S_8 is the amplitude of relatively broad linear-scale clustering; the early-galaxy problem is governed by the high-ν tail of the halo mass function at $z \gtrsim 10$, where the abundance is exponentially sensitive to the collapse threshold and to small changes in the variance on minihalo scales.

One qualification, and it is the whole content of this subsection's honesty. The sentence above is true of *baryonic* levers and false of *background* ones. A change in the expansion history acts on every mode that is already inside the horizon by one common factor, so it cannot lower the broad amplitude and raise the rare tail: those are the same modes. The one lever that genuinely satisfies the requirement is the collapse threshold, which acts on baryons in a specific environment rather than on the transfer function. This paper's own Condition D (§J.12) states the requirement exactly. The small-scale boost of §44.3 does not satisfy it, and its derivation is withdrawn on precisely this ground; the streaming damping would have satisfied it, and has no mediator (§44.4).

The present framework leans on this distinction in one direction only. It proposes no background modification, so it moves neither S_8 nor the rare tail through the expansion history; what it retains for the early galaxies is the collapse-threshold reduction of §44.5 — a baryonic-environment effect that lives in the high-mass exponential tail and leaves the broad linear amplitude untouched. That is the one row in Table 28 that differs between the broad and the rare-tail columns, and it is the only lever the framework still holds.

43.6 A compact consistency ladder for the growth block

The growth block reduces to three statements:

1. **Background:** the π-locked inventory with $w = -1$, no early residue, no extra radiation; the sound horizon and the rate need no correction (Section 42).

2. **Broad growth:** standard on that background; the amplitude follows from the declared A_s, $\sigma_8 = 0.80$, $S_8 = 0.825$ (§43.3); the lensing amplitude $A_L \approx 1.05$ is a late projection and does not move it.

3. **High-k tail:** the collapse-threshold reduction in the baryonic environment selectively helps early minihalo collapse and carries the scale selectivity the JWST section needs; it does not touch S_8.

One should not, therefore, look for a single monotonic change in the power spectrum at all scales and epochs. There is none on the broad scales, and the framework's statement about the earliest objects is confined to the tail.

Table 28: Sign structure of the Part VI growth channels, as they now stand. Two rows are withdrawn and kept for the record; one row is live.

Channel	**Effect on H_0**	**Effect on broad S_8**	**Effect on earliest galaxies**
Early dark energy lobe (withdrawn: the corridor is retired; the lobe is forbidden by the module's own kill criterion)	—	—	*would have been* negative by the same k-independent factor as the broad column (§44.3)
Extra dark radiation (withdrawn: $\Delta N_{\rm eff} \leq 0.03$, a derived null)	—	—	—
Mild lensing enhancement $A_L \approx 1.05$	none	none on the matter amplitude; enters the CMB-lensing interpretation only	little direct effect
Streaming damping (withdrawn: no mediator)	—	—	*would have been* positive for minihalo cooling; the gauged baryon sector has no vector below $M_{\rm lock}$ (§44.4)
Collapse-threshold reduction	none	none on broad linear scales	positive in the high-mass exponential tail; *the one live row, and the one that differs between the broad and rare-tail columns*

43.7 Kill criteria for the growth module

The growth module fails if the derived amplitude is excluded: a combined determination of S_8 from cosmic shear, clusters and redshift-space distortions that excludes 0.825 at three standard deviations, with the CMB-normalized value also moving away from it, would falsify the amplitude chain from A_s on this background — and since that chain has no dial, it would falsify the background. It also fails if the earliest galaxies require a boost of broad small-scale power that the standard growth on this background cannot supply and the collapse-threshold reduction cannot mimic (§44.8).

The paper should therefore be candid: the S_8 module is not a resolution of the S_8 tension and does not claim to be. It is one derived number, 0.825, that sits where ΛCDM's CMB-normalized amplitude sits, and it will be right or wrong with it.

44 Early galaxy formation in the JWST era

Proposition 44.1 (Early-halo amplification statement (T-Cosmos-12), restated)**.** *Because the halo mass function depends exponentially on the high-ν tail, small changes in the collapse threshold δ_c in the baryonic environment can produce order-one shifts in the abundance of the earliest massive systems. The JWST problem is therefore not "just more S_8": it is the nonlinear high-mass tail, and it is reached through the baryonic sector.*

What this proposition said earlier, and why it was narrowed twice. An earlier version of this statement read "small inherited changes in the early source ledger can produce order-one shifts," and closed by calling the JWST problem the high-mass tail of *the same inherited vector* as S_8. That form is too strong in one specific way. Changes in the early source ledger — the early-dark-energy lobe and the extra dark radiation — reach the rare tail only through the expansion history, and the expansion history acts on every sub-horizon mode by one common factor. They are therefore exponentially amplified, exactly as stated, but in the same direction as their effect on S_8: at the canonical audit values the amplification is a *suppression* of the rare tail by about two fifths at $\nu = 5$ (§44.3). The exponential sensitivity is real; what is not available is the freedom to point it in the opposite direction from the broad channel. Of the two levers that reach the tail without reaching the broad amplitude — the streaming damping and the collapse threshold — the first has no mediator under the terminated baryon gauge sector (§44.4), and the proposition above is now stated over the second alone. This is the paper's own Condition D (§J.12) applied to its own theorem.

44.1 Why the JWST problem is not just "more S_8"

The earliest JWST results triggered excitement because they appeared to reveal galaxies, stellar populations, and black-hole activity at redshifts high enough to strain the standard assembly timeline. Even as the observational interpretation matures, the theoretical challenge remains useful: can a constrained cosmology make the onset of luminous structure slightly earlier and slightly more efficient without rewriting the whole large-scale background? That is the sharp version of the early-galaxy question.

It is important not to confuse this with the S_8 problem. A modest suppression of broad linear

growth does not automatically imply a suppression of the rarest early collapsed objects. The earliest galaxies sit in the exponentially sensitive high-ν tail of the halo mass function, where small changes in the collapse threshold can move the first-light chronology (two further routes were claimed here — a small-scale transfer-function boost and a baryon–dark-matter streaming damping — and both are withdrawn, §44.3 and §44.4) by much more than one would guess from a broad-brush linear statistic.

The present framework once treated the early-galaxy tension through a three-layer mechanism, of which one layer stands:

1. a mild small-scale power boost — *withdrawn*: its derivation from the EDE residue of the Hubble discussion fails, and the residue's actual effect on this channel is negative (§44.3),
2. a damping of the baryon-dark-matter streaming velocity by the gauged baryon sector — *withdrawn*: the sector has no vector below the lock scale, so the channel has no mediator (§44.4), and
3. a reduction of the effective collapse threshold for the relevant early halos — the one lever that stands, baryonic and environmental.

The surviving effect is modest. Its importance lies in the fact that the halo mass function is exponentially sensitive.

44.2 Press–Schechter baseline and the reason tiny changes can matter enormously

The standard Press–Schechter intuition already explains why the early-galaxy problem is different from broad S_8. The abundance of halos above mass M at redshift z depends approximately on

$$n(>M,z) \propto \operatorname{erfc}\left(\frac{\delta_c}{\sqrt{2}\,\sigma(M,z)}\right) \tag{272}$$

or, in the rare-halo tail, exponentially on

$$n(M,z) \sim \exp\left[-\frac{\nu^2}{2}\right], \qquad \nu \equiv \frac{\delta_c}{\sigma(M,z)}. \tag{273}$$

Thus even a one- or two-percent shift in $\sigma(M,z)$ can produce a several-tens-of-percent change in abundance when ν is large, and a modest shift in δ_c can generate orders-of-magnitude effects.

This is the mathematical reason the early-galaxy problem is highly leverage-sensitive. One does not need to change the entire matter power spectrum enormously. One needs only to change the right combination of small-scale variance, baryon transport, and effective collapse conditions.

44.3 Small-scale boost from the early-dark-energy residue — and its withdrawal

The first proposed mechanism was that the same EDE residue which helps the Hubble tension produces a small positive boost in the small-scale transfer function. The working claim was a

roughly uniform enhancement of order

$$\frac{\Delta P(k)}{P(k)} \sim 0.03 \qquad \text{around} \qquad k \sim 5\,h\,\mathrm{Mpc}^{-1}. \tag{274}$$

Because σ is the square root of the power, this would correspond to

$$\frac{\Delta\sigma}{\sigma} \sim \frac{1}{2}\frac{\Delta P}{P} \sim 0.015\text{–}0.017. \tag{275}$$

Status: this derivation is withdrawn. The step from the residue to Eq. (274) has since been computed — in the companion JWST paper (Arisaka, 2026, C6-JWST), as the execution of this Part's own kill criterion (§44.8) — and it does not hold. The arithmetic below Eq. (275) is correct and is retained; what is withdrawn is the claim that *this residue* supplies the amplitude that arithmetic is fed. Three lines establish it, all of them from this paper's own numbers.

First, the scale. A perturbation mode is frozen outside the horizon and evolves inside it, so a localized episode of expansion leaves a feature only near the comoving wavenumber crossing the horizon at its epoch, $k_c = a(z_c)H(z_c)/c$. At the retired lobe's $z_c \approx 3500$ (Eq. (754)) that is $k_c \simeq 0.016\,h\,\mathrm{Mpc}^{-1}$ — the turnover scale, where a pre-recombination episode has always been known to act. Eq. (274) places the boost at $5\,h\,\mathrm{Mpc}^{-1}$, a factor of 317 away in k. The scale $k = 5\,h\,\mathrm{Mpc}^{-1}$ crossed the horizon at $z \simeq 1.56 \times 10^6$, about four hundred and fifty times earlier than the lobe.

Second, the sign. For a mode already deep inside the horizon, the sub-horizon growth equation reads $\delta'' + (2 + \mathrm{d}\ln H/\mathrm{d}\ln a)\delta' = \frac{3}{2}\Omega_m(a)\delta$ in $\ln a$. An extra smooth component raises H, so the friction term grows, and lowers $\Omega_m(a)$, so the source term shrinks. Both suppress, with no freedom of sign. Integrating with the retired pair at the intermediate ΔN_{eff} — $(f_{\mathrm{EDE}}, \Delta N_{\mathrm{eff}}, z_c) = (0.03, 0.04, 3500)$, the record computation of §K.10.5 — returns

$$\frac{\Delta\sigma}{\sigma} = -0.0185 \qquad \text{for every } k \gg k_c, \tag{276}$$

of which -0.0052 is the dark-radiation term alone. *Set that number beside Eq.* (275)*, which asserts* $+0.015$ *to* $+0.017$. They are the same magnitude, to within ten percent, with opposite signs. (At the legacy $\Delta N_{\mathrm{eff}} = 0.30$ the suppression is -0.051, three times larger and the same sign.)

Third, the scale independence, stated with its assumption. The suppression in Eq. (276) carries no k, because the sub-horizon growth equation carries none: a mode far inside the horizon responds to the background and to nothing else. Both 0.1 and $5\,h\,\mathrm{Mpc}^{-1}$ are far inside the horizon at z_c, so both take the same factor. *This residue, acting through the background, at fixed primordial amplitude and tilt, does not suppress the broad channel and boost the rare tail.* The qualification is not decoration. A route that changes the transfer function's *shape* — the equality scale or the tilt — is k-dependent without being horizon-localized, and this paper's own appendix names one: a self-consistent CMB refit of an early-dark-energy model drives ω_{cdm} and n_s upward, and an upward n_s raises small-scale power. That route is real, its sign is right, and this paper does not have the number: quantifying it requires the refit itself, which is the open item named in §K.10.5. So the honest statement is the narrow one — the derivation offered here, through the background at fixed

primordial spectrum, does not work — and the wider claim that no mechanism could is not made.

And both numbers are already in this paper, which is the part worth sitting with. The linear-response module of §J.11 books the same residue's effect on the broad channel as $\alpha_S \Delta N_{\rm eff} + \beta_S f_{\rm EDE}$, which at the canonical pair is 0.0288 — a growth suppression, before the lensing-amplitude compensation. Eq. (276) computes that same growth suppression from the background integration and returns 0.0185: the two agree in sign and to within a factor of one and a half, and the residual gap is the CMB refit that §K.10.5 identifies as the cause of the larger gap on the Hubble branch. **So the broad-channel suppression of §J.11 and the rare-tail enhancement of Eq. (275) are the same quantity, computed twice, with opposite signs — and at the canonical $\Delta N_{\rm eff}$ they are the same size as well.** The linear-response coefficients and the direct integration agree with each other on sign and magnitude; what neither supports is the second entry. Note also that this paper's own Condition D already forbids what Eq. (274) attempts — "the same model cannot simply suppress or enhance power uniformly at all scales" — and names the two levers that satisfy it, which are the second and third of this ladder and not the first.

And the cost is larger than a withdrawal. It is tempting to record this as one rung of the ladder falling silent. It does not fall silent; it reverses. Feeding Eq. (276) through this Part's own Eq. (277) at the representative $\nu \sim 5$ gives $\Delta n/n \approx -0.44$ in the linearized form, and the exact ratio is 0.62: **the residue reduces the earliest minihalo abundance by about two fifths**, where the withdrawn claim credited it with a 41% increase — +51% if the same exact treatment is applied to it. The swing is close to a factor of two and a half in abundance, in the wrong direction. (At the legacy $\Delta N_{\rm eff}$ the reduction is three quarters instead.) The two baryonic levers of this ladder must now overcome that deficit before they begin to advance the chronology, and the ladder is not one rung shorter but one rung negative. **This Part has not repriced the chronology shift against that deficit, and does not pretend to have**; doing so is the first item this revision leaves open, and until it is done the $\Delta z \sim 2$–3 of §44.7 and Table 29 should be read as the threshold rung's contribution before the first rung's deficit is subtracted, not as a net result.

What is withdrawn, and what is not. Withdrawn: the derivation of Eq. (274) from the corridor's residue. The amplitude is carried below as an *assertion* at working grade with no mechanism in the present ledger; a mechanism would have to come from the transfer function's *shape* — the tilt or the equality scale — rather than from the background, and this framework does not currently supply one. Not withdrawn: the S_8 magnitude, which this computation reproduces and which the paper now carries as one computed amplitude on one background with no growth response and no sign to settle (§43.3); and the rare-tail arithmetic that follows, which is correct as arithmetic and is retained so that a reader can see what an amplitude of this size would buy if a mechanism supplied it.

In the rare-halo tail one has approximately

$$\frac{\Delta n}{n} \approx \left(\nu^2 - 1\right) \frac{\Delta \sigma}{\sigma}. \tag{277}$$

Taking a representative $\nu \sim 5$ for the relevant minihalo population at high redshift gives

$$\frac{\Delta n}{n} \approx (25 - 1) \times 0.017 \approx 0.41, \tag{278}$$

that is, about a 40% increase in abundance. This arithmetic is useful and is retained: it shows that the quoted 40% is not an arbitrary large number but exactly the kind of exponentially amplified response one expects in the rare tail. What it does *not* show, after the status note above, is where an amplitude of 0.017 at $k \sim 5\,h\,\mathrm{Mpc}^{-1}$ comes from. That is now an open question and not a derivation.

The intuitive picture, read at its new grade, is equally simple. *If* something hands the very first gravitational seeds slightly deeper pockets, the exponential tail magnifies small initial help into a significant population effect. The corridor's residue is not that something. Of this ladder's three rungs, the first is now an assertion in search of a mechanism, and the reader should weight the ladder accordingly.

44.4 The baryon–dark-matter streaming velocity: a lever with no mediator

The second mechanism this paper once proposed was more specifically baryonic, and it is withdrawn here, with the reason printed. After recombination, baryons and dark matter retain a residual relative streaming velocity. In ordinary ΛCDM this relative motion acts like a headwind that delays gas capture by the smallest halos and raises the filtering mass for molecular-hydrogen cooling. The earlier GG-6/GDOS proposal was that the gauged $\mathrm{U}(1)_B$ sector induces a tiny long-range drag between baryons and the compact-angle dark-matter background — through a *coherent gauge-conversion channel*, in which the oscillating condensate sources a $\mathrm{U}(1)_B$ field that exerts a secular force on baryons — reducing the streaming velocity by about 20%:

$$v_{\mathrm{str}} \to 0.8\, v_{\mathrm{str}}^{\Lambda\mathrm{CDM}} \qquad \text{(withdrawn; see below).} \tag{279}$$

A useful rough consequence would follow if the damping were real. The baryonic filtering mass scales approximately as a positive power of the streaming velocity; a common heuristic is

$$M_f \propto v_{\mathrm{str}}^3, \qquad \frac{M_f^{\mathrm{new}}}{M_f^{\mathrm{old}}} \sim 0.8^3 \approx 0.5, \tag{280}$$

so a 20% reduction of the wind would lower the filtering mass by tens of percent, allowing molecular-hydrogen cooling and Population III ignition in smaller and earlier halos; Derivation notebook D (Appendix J) carries the heuristic as a heuristic.

Two results, obtained in sequence, retire the lever. The first was dynamical: the efficiency with which the oscillating condensate rectifies into a *secular* force was computed channel by channel, and all four candidates at minimal structure fail — rectification at coherence boundaries is short by eighteen orders and heats rather than drags; the baryon self-field channel violates momentum conservation; the ponderomotive channel is suppressed as g_B^4 and falls short by some 10^{84}; and a relic DC field offsets the streaming distribution rather than damping it — while the derivative coupling to the baryon-number current is a total derivative and exerts no classical force at any order. The one structure that would evade all four is a dark-sector component carrying a *linear* charge under the gauged baryon factor, which the ledger does not supply. The second result is structural and closes the question: under the boundary-condition termination of the gauged baryon symmetry

established in the anomaly companion (GG-A3-BI6; GG-P1-KK), $U(1)_B$ has no vector below the lock scale at all. Its Z'_B sits at $M_{\rm lock} \simeq 3.85$ TeV as a Dirichlet-terminated mode, with a range of 5×10^{-20} m, and the coherent channel above needs a light gauged coupling that the termination removes. There is no mediator, and no coupling floor to compare against the equivalence-principle frontier, because there is nothing for the condensate to convert into.

The 20% of Eq. (279) is therefore no longer carried, as a working amplitude or otherwise. What remains of the early-galaxy module is one lever, the collapse threshold of the next subsection, which is baryonic and environmental and owes nothing to the gauged sector; the chronology shift the module claims is carried by that lever alone (Table 29). This is the second of the module's three original rungs to be withdrawn, and the paper records both withdrawals rather than reworking the ladder around them.

44.5 Collapse-threshold reduction and the high-mass halo tail

The third mechanism is the reduction of the effective collapse threshold. The working GG-6 claim is that the altered baryonic infall history in the relevant environment reduces the critical overdensity from the standard spherical-collapse value,

$$\delta_c^{\Lambda\rm CDM} \approx 1.686, \tag{281}$$

to an effective early-galaxy value of roughly

$$\delta_c^{\rm GG\text{-}6} \approx 1.18. \tag{282}$$

Why can this matter so dramatically? Because the exponential tail depends on δ_c^2. At fixed σ, the ratio of abundances scales as

$$\frac{n_{\rm GG\text{-}6}}{n_{\Lambda\rm CDM}} \sim \exp\left[\frac{\left(\delta_c^{\Lambda\rm CDM}\right)^2 - \left(\delta_c^{\rm GG\text{-}6}\right)^2}{2\sigma^2}\right]. \tag{283}$$

For representative high-redshift values $\sigma \sim 0.34$–0.36 on the relevant massive-halo scale, the exponent lies in the several-unity range, naturally producing an enhancement from 10^2 to 10^3. A factor near 500 is therefore perfectly plausible inside the rare tail. The point is not that the exact number is sacred before full hydrodynamical calibration. The point is that the magnitude of the claimed boost is mathematically understandable.

This mechanism is also conceptually neat. The gas does not collapse earlier simply because gravity has been made universally stronger. Rather, the collapse criterion itself is changed in a specific baryonic environment. That is a much narrower and more defensible claim — and it is the one rung of this ladder left standing as stated, because it is baryonic, environmental, and therefore genuinely scale-selective in the sense Condition D requires.

The chronology map, made exact. This Part has summarized the ladder's effect as a first-light shift of $\Delta z \sim 2$–3 without ever stating what threshold that corresponds to. The map is

elementary and is recorded here so that the summary and the mechanism cannot drift apart. Since $\nu = \delta_c/\sigma$ and $\sigma \propto (1+z)^{-1}$ in the matter era, a reduced threshold reproduces the baseline rarity — and hence the baseline abundance — at a higher redshift:

$$(1+z') = (1+z)\,\frac{\delta_c}{\delta_c^{\rm eff}}, \qquad \Delta z = (1+z)\left(\frac{\delta_c}{\delta_c^{\rm eff}} - 1\right). \tag{284}$$

At $z = 10$, the full working value $\delta_c^{\rm eff} = 1.18$ of Eq. (282) gives $\Delta z = 4.7$, not 2–3. The summary's $\Delta z = 2$ corresponds to $\delta_c^{\rm eff} = 1.427$ and $\Delta z = 3$ to 1.325 — that is, to the threshold reduction operating at roughly one half to seven tenths of its full amplitude on the population average. (The second of these sits almost exactly on a rounding boundary — 1.3251 — and is therefore quoted here to three digits; to two it is 1.33, which is the form the companion JWST paper uses in its fits. Two digits are not enough to keep the two papers from appearing to disagree, and that is the only reason the third is printed.) That is consistent with this Part's own framing, since the reduction is claimed for "the relevant early halos" and not for the population as a whole; but it had never been stated as a number, and a summary that cannot be turned into a number is a place where two versions of a framework can quietly disagree.

44.6 Direct-collapse seeds and the early supermassive black-hole problem

The JWST story is not only about galaxies; it is also about apparently over-massive black holes at high redshift. Within the present framework the same early-galaxy assistance naturally feeds a direct-collapse seed channel. If baryonic inflow is delayed just enough to keep gas more pristine while dark halos hierarchically merge, and if the effective Toomre threshold is lowered in the relevant high-redshift disks, then one can favor heavy seeds of order

$$M_{\rm seed} \gtrsim 10^5\, M_\odot, \tag{285}$$

bypassing the slow climb from stellar-mass remnants.

This is not yet a full black-hole paper; that comes in Part VII. But even inside Part VI it is important to state the continuity. The early-galaxy module is scientifically stronger if it also explains why the first bright AGN need not be impossible. The same chronology shift that helps the first massive galaxies should also help the first heavy seeds.

44.7 The chronology shift: why $\Delta z \sim 2$–3 is enough to matter

The combined effect of the mechanisms is summarized as an effective chronology shift,

$$\Delta z \sim 2\text{–}3, \tag{286}$$

for the onset of the earliest star formation and significant galaxy assembly. It is tempting to dismiss such a shift as small. That would be a mistake. In the high-redshift universe, a shift of even a few redshift units corresponds to a large fraction of the available cosmic time between matter domination, first cooling, and the onset of reionization. Because so many first-light processes are

threshold-driven — minihalo cooling, Population III ignition, metal enrichment, seed black-hole formation, merger-driven brightening — a small shift in cosmic time can translate into a large shift in the observable abundance of bright objects.

The early-galaxy section should therefore insist on a basic point of intuition. The problem is not to make the entire universe mature implausibly early. The problem is to let the earliest successful objects cross their ignition thresholds slightly earlier and in slightly larger numbers. Exponential sensitivity then does the rest.

What carries this shift, and what pulls against it. The paragraph above described $\Delta z \sim 2$–3 as "the combined effect of the mechanisms," and an earlier version read the ladder below as three rungs pushing the same way. That reading does not survive the audit of §44.3. The shift is now carried by the third rung alone — the reduced collapse threshold — which by the map of Eq. (284) reaches $\Delta z = 2$ at $\delta_c^{\rm eff} = 1.427$ and $\Delta z = 3$ at 1.325. The first rung contributes nothing, because its derivation is withdrawn and the residue it was derived from moves the rare tail the *wrong* way: at the canonical audit pair the same integration that shortens the sound horizon suppresses the $\nu = 5$ minihalo abundance by about two fifths, which is a chronology shift of roughly -0.6 in z at $z = 10$, against the ladder rather than with it. The second rung is unpriced. So the honest statement is not that three effects combine to 2–3; it is that *one* effect must reach 2–3 net, having first paid back what the early residue takes away, and that whether it does is a question about the threshold reduction and nothing else. That is a narrower claim than the one this subsection made, and it is the one the paper can defend.

Table 29: Mechanism ladder proposed for the JWST-era early-galaxy problem in GG-A10-Cosmos, *with the present status of each rung.* The columns give the mechanism, a representative shift, the scale it principally affects, the observational consequence it is intended to have, and where the mechanism now stands after the two audits this Part's own kill criterion produced. The shifts are representative rather than predicted. **Two of the three rungs no longer have a mechanism**: the first because the residue cannot reach the scale it was said to act on, the second because no rectification channel survives at minimal structure and the terminated baryon gauge sector has no vector below the lock scale to carry one. The third stands as stated. A reader weighting this ladder should weight it as one derived rung and two assertions.

Mechanism	Representative shift	Primary scale affected	Intended observational consequence	Status
Early-time small-scale boost	$\Delta P/P \sim 3\%$	asserted at high k; the residue acts at $k_c \simeq 0.016\,h\,\mathrm{Mpc}^{-1}$	$\sim$ 40% more relevant minihalos at very high redshift	**derivation withdrawn** (§44.3); amplitude asserted, no mechanism in the ledger
Streaming-velocity damping	$v_{\rm str} \to 0.8 v_{\rm str}$	baryon capture / filtering mass	earlier molecular-hydrogen cooling and Population III ignition	**withdrawn**: the four rectification channels fail, the channel theorem forbids the naive one, and the terminated baryon gauge sector has no vector below $M_{\rm lock}$ — no mediator (§44.4)
Collapse-threshold reduction	$\delta_c : 1.686 \to 1.18$	rare massive-halo tail	order-of-magnitude boost in $> 10^{11} M_\odot$ halo abundance at $z \sim 9$–10	**stands as stated**; baryonic and environmental, so scale-selective in the sense Condition D requires
Direct-collapse seed channel	$M_{\rm seed} \gtrsim 10^5 M_\odot$	early galactic nuclei	less tension from very high-redshift bright AGN and quasars	unchanged; rides on the third rung, not the first
Net chronology shift	$\Delta z \sim 2$–3	first-light timeline	earlier onset of bright galaxies without wholesale background revision	**not a sum over the rungs**: carried by the third alone, against a ~ -0.6 contribution from the early residue; $\delta_c^{\rm eff} = 1.43$–$1.325$, Eq. (284)

44.8 Kill criteria for the JWST module — and the verdict, now that they have been run

The early-galaxy module is falsified if detailed hydrodynamical modeling demonstrates that the required chronology shift cannot produce the needed abundance increase without other implausible astrophysical assumptions. It is also weakened if the evolving observational picture reduces the early-galaxy tension to purely photometric or stellar-population systematics. The theory should acknowledge this openly.

The verdict. The first of those criteria was registered here and long stood unexecuted. It has now been run — in the companion JWST paper (Arisaka, 2026, C6-JWST), which imported it verbatim — and it returns a split verdict in three parts. *Neither external data set kills the bundle.* The damping tail does not fire at all, and the reason has changed since this verdict was written: the early residue that would have moved r_d/r_s is now zero on both terms — $f_{\rm EDE} = 0$ with the corridor retired, $\Delta N_{\rm eff} = 0$ as the derived null of row D12 — so the predicted shift is nil and the criterion cannot fire. (For the record of what was checked: at the intermediate values $z_c \approx 3500$

and $\Delta N_{\rm eff} = 0.04$ the pair raised r_d/r_s by 0.46%, about one sigma; at the legacy 0.30 it would have been 1.05%, about two sigma, and the criterion would have been close to firing. Both calculations are kept in Appendix K.10.) The forest is comfortable: a +3% bump at $k \sim 5\,h\,\mathrm{Mpc}^{-1}$, propagated at linear order through the one-dimensional *matter* projection, raises the flux power by 1.1–2.1% across the measured band for a broad bump ($\sigma_{\ln k} = 1$) and 0.5–1.5% for a narrow one ($\sigma_{\ln k} = 0.5$), against a $\sim 5\%$ amplitude precision and a $\sim 10\%$ thermal systematic — and the linear matter projection over-states the flux response at the band's small-scale end, so the approximation errs toward caution. So the criterion's first clause returns *not killed, and watch the dark radiation.*

But the criterion was pointed at the wrong failure. What it did not ask — and what the execution asked in passing — is whether the residue that survives the first test *produces* the boost priced by the second. It does not, for the reasons set out in §44.3, and the first rung's derivation is withdrawn on that ground rather than on either of the two the criterion named. *A kill criterion that names the external tests and not the internal one is an incomplete criterion*, and this Part now says so: the module's sharpest exposure was never the forest or the damping tail, but the question of whether one background episode can act at two separated scales in opposite directions.

Yet even under that verdict, the module remains scientifically important — and is now more useful than it was, because two of its three rungs have been converted from unexamined proposals into named open structures. It tests whether the same inherited ledger that moves H_0 and S_8 can also move the first-light chronology in a useful direction, and the honest current answer is that the ledger's *third* lever can, that the first cannot as claimed, and that the second cannot yet be priced. That is exactly the kind of cross-domain constraint a unified cosmology should face, and exactly the kind of answer it should be willing to print.

45 Integrated summary of the growth-and-tension block

> One background, read by every probe that has no calibrated rung: the π-locked inventory fixes the rate the microwave background reads and the amplitude the clustering surveys read, with nothing added in the early universe and nothing adjusted; the distance ladder is the one measurement the framework predicts to be displaced, and the earliest galaxies are reached through a baryonic collapse threshold that acts on the rare tail and on nothing else.

That sentence is the real claim of Part VI. Everything else is a calculation or an audit of it. (The sentence it replaces claimed a "narrow early-time residue" that helped the Hubble tension and reshaped growth; the residue is zero, and the claim is withdrawn.)

It is worth making the logic explicit one last time.

45.1 One ledger, four manifestations

First manifestation: the primary CMB. The acoustic ruler, the odd/even peak pattern, the damping tail, the E-mode sequence, and the tensor corridor are all controlled by the inherited matter loading, the same multiscale throat capacity, and the same inflationary background that earlier parts of the paper already constrain. If this fails, everything fails.

Second manifestation: the Hubble tension. The background rate is derived, 66.8 ± 0.2, and the microwave background reads it through the inherited sound-horizon calibration; there is no early-dark-energy lobe and no radiation excess, and no local lift. The distance ladder's 73.5 is read as a calibration systematic the framework predicts, and the ladder-free late probes are the test (Section 42).

Third manifestation: the S_8 tension. Growth is standard on the same background, and its amplitude is computed from the declared A_s: $S_8 = 0.825$, at the microwave background's value and 2.5σ above the cosmic-shear value. The framework shares the tension and does not resolve it (§43.3).

Fourth manifestation: the JWST-era early-structure challenge. The computed amplitude does not preclude earlier galaxies, because the collapse threshold is a lever of a different kind: it acts on baryons in an environment rather than on the expansion history, and only a lever of that kind can treat two scales differently. *This is a correction of what this paragraph used to say.* It used to list a third lever — a small variance increase from the same early residue — alongside two baryonic ones, and the second baryonic lever, the streaming damping, has since lost its mediator (§44.4). That was wrong, and wrong in a way this paper's own Condition D would have caught: a background episode acts on every mode inside the horizon by one common factor, so the variance increase and the broad-growth suppression are the same number with opposite signs (§44.3). What is left is a lowered δ_c alone, and it carries the whole of the chronology shift: the streaming damping is withdrawn with the first rung (§44.4), the gauged baryon sector having no vector below $M_{\rm lock}$ to mediate it. The manifestation stands on one leg of the three it began with.

These four manifestations are not independent. That is both the strength and the danger of the proposal. The strength is that the framework is highly constrained; it cannot fix each problem independently by inventing a new fluid or field. The danger is that failure in one place propagates to others. This is why Part VI is one of the cleanest falsification arenas of the whole paper.

45.2 What Part VI does and does not claim

Part VI claims three things.

1. It claims that the perturbative and growth sectors of cosmology can be organized as a sequence of audits on one inherited ledger rather than as a grab bag of modular fixes.

2. It claims that the background needs no correction: the rate and the clustering amplitude follow from the inventory and the declared amplitude, the radiation content is the Standard Model's ($\Delta N_{\rm eff} \leq 0.03$, a derived null), and the two "tensions" are read as a predicted ladder systematic and a shared amplitude tension respectively. (This claim replaces one about an inherited correction pair $(f_{\rm EDE}, \Delta N_{\rm eff})$, which is withdrawn.)

3. It claims that the baryonic $\mathrm{U}(1)_B$ sector supplies an additional selective structure-formation lever that broad growth statistics alone do not capture. *This third claim now carries a rider*: the lever's *selectivity* is exactly the property the first claim's small-scale boost turned out not to have, so the claim is now narrower and better defended — selectivity belongs to the baryonic sector and to nothing else in this ledger.

And it now claims one thing less than it did. Earlier this Part also claimed, in the first rung of its mechanism ladder, that the corridor's early residue supplies a small-scale power boost at minihalo scales. It does not, and the claim is withdrawn (§44.3). Removing it costs the ladder its smallest contribution — roughly a 40% surplus of the earliest minihalos, against factors of hundreds from the third rung — and costs the framework something larger in kind: the ladder was presented as three consequences of one structure, and it is now one consequence, one assertion, and one unpriced proposal. That is a worse position than the one this Part once occupied, and it is the position the evidence supports.

Part VI does *not* claim that the full precision likelihood problem has already been completed inside this paper. Nor does it claim that the astrophysical side of the JWST problem can be solved without hydrodynamical modeling. The correct scientific posture is intermediate. The framework supplies a coherent narrow mechanism bundle and then states the observational conditions under which that bundle survives or fails.

45.3 The real significance of the open-system viewpoint

The most novel conceptual feature of Part VI may not be any one number. It may be the insistence that cosmological observables are open-system exports rather than direct readings of a closed fundamental microstate. This matters in several places.

It matters for the CMB because the observed sky is a frozen and projected record, not a direct local sample of the early plasma. It matters for the Hubble tension because a calibrated ladder and a sound-horizon dictionary are different instruments reading one rate, and one of them can carry a systematic. It matters for S_8 because lensing and clustering are themselves late integrated observables with non-trivial coarse-graining. And it matters for early galaxies because the relevant thresholds depend on environmental realization and baryonic transport, not only on the primordial scalar amplitude.

This open-system emphasis does not weaken precision cosmology. It sharpens its interpretation. It asks the theorist to be explicit about what is projected, what is coarse-grained, and what is truly invariant. In that respect, Part VI is a natural continuation of the earlier GG-6/GDOS papers rather than a thematic detour.

45.4 An invitation: revisiting ΛCDM in the light of a derived $H(t)$

There is a larger consequence to draw, and we state it plainly, as an invitation to the community. ΛCDM is the most successful *fitting* framework in the history of cosmology: six numbers, an exquisite fit, four independent probes in concordance. What it does not do is say *why* those six numbers take the values they do — it measures them. GG-A10-Cosmos proposes that they are not free: that the densities, and therefore the entire expansion history $H(t)$ of Figure 26, are derived from one six-dimensional geometry with no continuous dial. If that proposal holds, the parameters of ΛCDM are *outputs* of a deeper theory rather than inputs to a fit, and several of the field's standing "tensions" change character: the H_0 tension becomes a question about the calibration of one point on a derived curve, and the S_8 tension a question about late-time coarse-graining, rather than signals of missing fundamental physics.

We are therefore not asking anyone to discard ΛCDM — GG-A10-Cosmos *runs on* its Friedmann/FRW machinery, and its empirical fit is not in question. We are asking the community to revisit ΛCDM in a fresh light: to re-examine, parameter by parameter, whether each of its numbers is a free input or a forced geometric consequence, and to test the specific, dated predictions this reading makes — the standard-siren convergence $H_0^{\rm GW} \to 67$, the exact $w = -1$, the tensor ratio $r = 0.0038$, and the coherent π-offset visible across the acoustic peaks. The honest banner is not "ΛCDM is wrong" but "ΛCDM is incomplete in its *explanation*, and here is a candidate completion — laid out to be tested, and, if it fails, to fail cleanly." A generation of precision cosmology measured the constants of the universe to extraordinary accuracy. The next task — and it is a community effort — is to find out which of them were never constants at all, but consequences.

45.5 Bridge to the next part

Having completed the growth-and-tension block, the paper is in a better position to move to the black-hole and primordial-relic sector. The transition is natural. The CMB section already forced the paper to speak carefully about what is operationally observed after projection. The Hubble and S_8 sections forced it to distinguish between background laws and inferred quantities. The early-galaxy section forced it to respect the difference between broad statistics and rare-event tails. All three lessons reappear in the black-hole information problem, in primordial-black-hole formation, and in the tensor sector.

One rate, and a definite statement about which instrument reads it wrong: that is this Part's reframing of the decade's most famous quarrel, and there is a philosophical point in it. A tension between instruments is usually read as somebody's mistake; here it is read as *information* — the universe giving two lawful projections of a single geometric object, the disagreement itself becoming a measurement of the switch between them. If the gravitational-wave sirens side with the ladder at 73, this reading dies in public. That is not a defect of the framework; it is the entire point of having one.

For that reason Part VI should end not with triumphalism but with a clean summary. The perturbation, growth, and tension sectors do not yet prove the whole framework correct. They do something more useful. They convert vague cosmological ambition into a finite list of tightly coupled empirical tests. That is exactly what the paper needs before it enters the more speculative territory of Part VII (black holes and primordial relics). The full cross-consistency Jacobian and the worked, from-scratch derivations for every number in this Part are collected in Appendices J and K.

46 Part VI in plain English: one switch, and what it does and does not fix

The proposal of this Part is a single geometric switch, and its two faces are the Hubble tension and the growth of structure. That is an economical claim and economy is exactly why it is dangerous: a mechanism that moves two observables in the right direction at once is more likely to be a mechanism than a coincidence, and it is also more likely to be over-read.

So the Part is written to make the strength of each piece visible separately. The background

Hubble value is derived twice, by independent roads, and that is the strong result. The residual offset against the local distance ladder is not derived here; it is treated as a calibration systematic, which is the reading this framework now shares with the mainstream that takes the CMB to be right. Those two statements sit in the same section and carry different weights, and the status table exists so they cannot be conflated.

The growth block is no longer a set of corrections presented as signs: the clustering amplitude is computed from the declared primordial amplitude and lands where the microwave background's value lands, so the framework shares the S_8 tension with the standard model rather than resolving it, and says so. The early-galaxy material is weaker again — a ladder of mechanisms, ordered by how much of the discrepancy each could carry, not by how well established each is. A reader should come away knowing which of the four blocks in this Part would survive a hostile reading and which would not.

Part VII
Black holes and primordial relics

The black-hole information paradox under OGGEP / GDOS; primordial black holes as a feature-dependent supplementary channel; and the Starobinsky-channel relict primordial gravitational-wave background

47 Black holes as GRIP-In way-stations: the organizing idea

Part VII reads the extreme-gravity sector through one organizing idea, shared with the dedicated black-hole companion (Arisaka, 2026, C3-BH): *a black hole is a GRIP-In way-station attractor, and Hawking radiation is its channel-mediated GRIP-In exit.* In the GRIP grammar of GG-A6-GRIP, a black hole is a deep, stable basin of the F_{GFF} free-energy ledger on the gravitational stratum — but, unlike the Ground Geometric Attractor that the universe relaxes into at late times (Part IV), it is not terminal: it possesses an exit channel. Hawking emission is that exit, the slow channel-mediated descent by which the basin finally drains — a channel that is presently closed for every astrophysical hole, whose Hawking temperature lies below the microwave background's, and opens only when the sky cools past it (the conditional ground state of the black-hole companion, its C-BH-1). A stellar-mass hole is therefore the most extreme *way station* in nature, a place where the cosmic descent pauses for a residence time of order 10^{67} years before resuming.

This single picture organizes the three sections that follow and fixes the division of labor with the companion paper. The dedicated treatment — the exact $A/4$ Bekenstein–Hawking coefficient from induced gravity, the Page-curve branch switch, the four laws of black-hole mechanics as GDOS statements, and the EHT-shadow and ringdown predictions — is carried in (Arisaka, 2026, C3-BH); Appendix N reproduces the $A/4$ derivation and credits it there. The present Part carries the *cosmological* reading of the same mechanism in three movements:

- the **information paradox** (§48), dissolved by reading the horizon as the OGGEP port of a globally closed bulk;
- **primordial black holes** (§49.1), the nested GRIP-In collapse events of cosmic history — scarce in the minimal module and, where they do occur, spaced — by a conjecture the companion grades as model-level — by the same integer ledger that fixes $z_* = e^7$; and
- the **primordial gravitational-wave relic** (§50.1), the Starobinsky-channel tensor background that the derived plateau fixes at $r = 0.0038$, a two-sided target.

All three are the same open-system descent, audited at the strongest causal barrier in classical relativity.

48 The black-hole information paradox

Theorem 48.1 (Closed-bulk / open-boundary black-hole theorem (T-Cosmos-13))**.** *The information paradox is reformulated by separating the closed bulk TGG ledger from the open boundary-accessible seam description. The bulk may remain globally lawful and unitary while the boundary observer records a restricted packet whose entropy and Page-turnover behavior are described by CCI, F_{GFF}, and GRIP rather than by a naive factorization of the entire bulk Hilbert space into two directly readable parts.*

The fourth cluster of open problems is different in character from the first three. The beginning of the universe, the origin of matter, and the growth of structure are all questions about how a cosmic background evolves and how that evolution is read out in ordinary observables. The black-hole information paradox is sharper and more surgical. It asks what a theory means by *information*, what it counts as an *observer-accessible subsystem*, and how thermodynamic bookkeeping is related to microscopic unitarity in a regime where gravity, quantum theory, and horizons meet head-on.

This is exactly why the problem belongs in GG-A10-Cosmos and not only in the dedicated black-hole companion (Arisaka, 2026, C3-BH): the cosmological reading here and the full gravitational treatment there are two projections of one mechanism. If the GG-6 constitution is serious, then OGGEP is not a decorative philosophical layer. It is the rule that tells us how a closed higher-dimensional reality is projected into boundary observables. Likewise, if GDOS is serious, then a black hole is not merely a static exact solution of Einstein's equation. It is an extreme open system with explicit ports, a sharply defined energy-entropy ledger, and a hierarchy of accessible and inaccessible records. Black holes are therefore the natural place where the distinction between *global closure* and *boundary openness* stops sounding philosophical and becomes a concrete technical issue.

There is also a second reason. The standard formulation of the paradox is built on a set of assumptions that are individually familiar but not jointly innocent. One assumes semiclassical locality outside the horizon, exact thermal Hawking emission at late times, ordinary factorization of the Hilbert space into interior and exterior pieces, and final complete evaporation into exterior quanta. When these assumptions are assembled, the conclusion seems to be that a pure initial state evolves into mixed Hawking radiation. The usual response is then to decide which assumption to sacrifice: exact thermality, local effective field theory, interior/exterior independence, or smoothness of the horizon. The modern island program shows that the bookkeeping can be reorganized in a much more subtle way. The GG-6/GDOS response goes one step further backward. It asks whether the paradox itself is being posed at the wrong ontological level.

The central claim of the present section is therefore two-layered.

First, at the level of ontology, the paradox is dissolved once one stops demanding that the boundary-accessible subsystem by itself carry the full burden of microscopic unitarity. In OGGEP the 6D bulk is the globally closed unitary system. The 4D observer sees only a CPTP image of that bulk state. Pure-to-mixed evolution in the accessible subsystem is therefore not a contradiction; it is the generic mark of openness.

Second, at the level of practical entropy accounting, the familiar Page-curve resolution can be

reinterpreted as a GDOS saddle-switch problem. The relevant entropy is not the naive thermal entropy of emitted quanta alone, but the minimum of an GFF ledger over admissible boundary-accessible branches. In that language the Page transition is a GSSB event: the realized entropy branch changes when the operationally dominant saddle switches.

This is what it means to say that the paradox can be solved cleanly by GDOS. The framework does not need to deny Hawking radiation. It does not need to deny that late-time radiation can, under the correct coarse-graining, carry reconstructible information. It only insists that the correct global-versus-operational distinction be made before the entropy ledger is interpreted.

48.1 Why black holes are the sharpest information-processing objects in the paper

Black holes are the most demanding objects in the present volume because they compress several themes into one place.

At the geometric level, a black hole is the cleanest example of a system whose macroscopic state can be written almost entirely in terms of conserved charges and horizon geometry. At the thermodynamic level, it is the most famous system whose entropy is controlled by an area rather than by an ordinary matter volume. At the quantum level, it is the regime in which apparently thermal radiation emerges from a state produced by collapse. At the operational level, it is the regime in which an outside observer has the most severe possible restriction on direct access.

Any program built around GFF, DGI, GSSB, OGGEP, and boundary-accessible observables should therefore expect black holes to be a decisive audit. If the open-system language is only a metaphor, it will fail here. If it is a genuine constitutional principle, it should clarify rather than obscure the paradox.

In keeping with the rest of the paper, the present section proceeds conservatively. It does not assume a detailed microscopic model of black-hole interior states. It does not require a full UV completion of every graybody factor or every trans-Planckian step in collapse. Instead, it asks a narrower question: what does the inherited GG-6/GDOS ledger already force us to say?

Three inherited ingredients are enough to make the central move.

1. The bulk ontology is globally closed and unitary, while the boundary observer sees an open-system projection.

2. Physical state counting must respect DGI; in gravity this means one cannot naively factorize across a gauge-redundant horizon split.

3. The correct entropy ledger for an open system is GFF, with a geometric cost term and an entropic-curvature term.

Everything else in the section follows from working out the consequences of these three facts for the particular case of an evaporating Schwarzschild-like black hole.

48.2 Closed-system versus operational formulations

The usual closed-system statement of the paradox can be made precise in four steps.

First, one begins with a pure state $|\Psi_{\rm in}\rangle$ of collapsing matter together with the vacuum of quantum fields on the exterior geometry. If the entire process is governed by unitary quantum mechanics, the total final state should remain pure:

$$|\Psi_{\rm tot}(t)\rangle = U(t)|\Psi_{\rm in}\rangle, \qquad U^\dagger U = \mathbb{1}. \tag{287}$$

Second, semiclassical quantum field theory on a background with an event horizon predicts approximately thermal particle creation. For a non-rotating uncharged black hole of mass M, the Hawking temperature is

$$T_H = \frac{1}{8\pi G_N M}, \tag{288}$$

while the Bekenstein-Hawking entropy is

$$S_{\rm BH}(M) = \frac{A_H}{4G_N} = 4\pi G_N M^2. \tag{289}$$

Here $A_H = 16\pi G_N^2 M^2$ is the horizon area of the Schwarzschild geometry. The evaporation law may be written schematically as

$$\frac{{\rm d}M}{{\rm d}t} = -\frac{\alpha_H}{M^2}, \tag{290}$$

where $\alpha_H > 0$ packages the species count and graybody factors. Integrating gives

$$M^3(t) = M_0^3 - 3\alpha_H t, \qquad t_{\rm evap} = \frac{M_0^3}{3\alpha_H}, \tag{291}$$

for an initial mass M_0.

Third, if one insists that the outgoing radiation be described by an exactly thermal density matrix throughout the evaporation process, the coarse-grained radiation entropy increases monotonically. A convenient schematic estimate is

$$\dot{S}^{\rm th}_{\rm rad} \approx \eta\, \frac{-\dot{M}}{T_H}, \tag{292}$$

with η an order-unity coefficient. Using Eq. (288) gives

$$\dot{S}^{\rm th}_{\rm rad} \approx 8\pi\eta\, G_N M(-\dot{M}). \tag{293}$$

Integrating with Eq. (290) yields the familiar monotonic behavior

$$S^{\rm th}_{\rm rad}(t) \approx \eta\left[S_{\rm BH}(M_0) - S_{\rm BH}(M(t))\right]. \tag{294}$$

If the black hole evaporates away completely and the radiation remains exactly thermal, then the late-time radiation state is mixed and carries entropy of order the initial area.

Fourth, one now collides this with Eq. (287). A pure state should not evolve into a truly mixed state under exact unitary evolution. This is the paradox in its sharpest textbook form.

Stated this way, the tension appears irresistible. Yet it conceals a crucial shift of level. Equation (287) is a statement about a *total* closed system. Equation (294), by contrast, is a statement

about a *subsystem* described by an effective thermal channel. The contradiction is therefore only forced if one identifies the accessible subsystem with the total ontology.

That identification is precisely what OGGEP rejects.

In the operational formulation, the outside observer does not directly possess $|\Psi_{\rm tot}\rangle$. The observer possesses a boundary-accessible state,

$$\rho_\partial(t) = \Pi_t(\rho_6(0)) = \sum_k K_k(t)\,\rho_6(0)\,K_k^\dagger(t), \tag{295}$$

where $\rho_6(0)$ is the global 6D state, Π_t is a CPTP bulk-to-boundary channel, and the Kraus operators K_k encode the operational projection together with environmental tracing and horizon-mediated openness. The point of Eq. (295) is not that black holes are somehow special among open systems. The point is the opposite: once the outside observer is correctly described as holding only ρ_∂, pure-to-mixed behavior in that reduced state is completely unsurprising.

This is the first clean GDOS move. The paradox is not "solved" by erasing Hawking thermality. It is solved at the constitutional level by refusing to impose total-system purity on a boundary-accessible subsystem.

48.3 The black hole as a System Instance Packet

The second clean GDOS move is to define the system before deriving its entropy law. The black hole is treated as a System Instance Packet,

$$\mathcal{S}_{\rm BH} = (\mathcal{M}_{\rm BH}, \mathcal{G}_{\rm BH}, \mathcal{P}_{\rm in}, \mathcal{P}_{\rm out}, F_{\mathsf{GFF}}, \mathcal{O}_{\rm BH})\,. \tag{296}$$

Each entry has a concrete meaning.

- $\mathcal{M}_{\rm BH}$ is the coarse-grained state manifold. For a Schwarzschild-like evaporating hole one can take coordinates such as $(M, S_{\rm BH}, S_{\rm rad}, \chi_{\rm env}, \ldots)$, where $\chi_{\rm env}$ stands schematically for environment and measurement records.
- $\mathcal{G}_{\rm BH}$ is the induced metric and connection data on that manifold, including the gauge-constrained directions associated with diffeomorphism redundancy and horizon edge data.
- $\mathcal{P}_{\rm in}$ denotes infalling matter and measurement interventions that inject energy or information into the accessible description.
- $\mathcal{P}_{\rm out}$ denotes outgoing Hawking quanta and other exterior fluxes.
- F_{GFF} is the energy-entropy ledger that ranks admissible coarse-grained histories.
- $\mathcal{O}_{\rm BH}$ is the observable dictionary: what an exterior observer can actually register.

This explicit SIP formulation clarifies a common confusion. In a closed Hamiltonian derivation one often oscillates between three levels without warning: the exact total state, the coarse-grained outside description, and the detector records. GDOS forbids that slippage. Each must be placed in

the packet explicitly. Once this is done, the two entropy notions that are usually conflated separate cleanly:

1. the fine-grained entropy of the total 6D bulk state, which remains constant under global unitary evolution, and
2. the coarse-grained or operational entropy of the exterior-accessible subsystem, which may increase, saturate, or decrease depending on the dominant branch in the GFF ledger.

In this sense the SIP formalism already explains why Hawking's original monotonic argument was so compelling and yet incomplete. It was a valid statement about one reduced description, not a theorem about the total ontology.

48.4 GFF and the correct entropy ledger

The third clean GDOS move is to identify the correct entropy functional. In the ordinary semiclassical account, one separately discusses horizon area and radiation entropy, and then later discovers that the physically relevant quantity is their sum. In the present framework that sum is not a surprise but the direct form of the GFF ledger.

The basic object is

$$F_{\mathsf{GFF}}[X;t] = F_{\mathsf{GEP}}[X;t] + F_{\mathsf{CEP}}[X;t]. \tag{297}$$

For black-hole entropy accounting the geometric term is the area contribution,

$$F_{\mathsf{GEP}}[X;t] = \frac{A[X;t]}{4G_N}, \tag{298}$$

while the entropic-curvature term is the bulk entanglement or state-volume contribution,

$$F_{\mathsf{CEP}}[X;t] = S_{\mathrm{bulk}}[X;t]. \tag{299}$$

Thus

$$F_{\mathsf{GFF}}[X;t] = \frac{A[X;t]}{4G_N} + S_{\mathrm{bulk}}[X;t]. \tag{300}$$

This is exactly the structure that modern semiclassical gravity writes as generalized entropy. The difference is interpretive. In GG-6/GDOS it is not a late-discovered miracle of gravity. It is the only admissible scalar ledger once one accepts the open-system constitutional layer.

The physically accessible radiation entropy is then not the naive thermal entropy of quanta emitted to infinity. It is the minimum over admissible branches:

$$S_R(t) = \min_{X \in \mathfrak{B}_{\mathrm{adm}}} F_{\mathsf{GFF}}[X;t], \tag{301}$$

where $\mathfrak{B}_{\mathrm{adm}}$ is the set of boundary-admissible saddles or coarse-grained branches. In ordinary island language one might call these branches "no-island" and "island" extremal surfaces. In the present language they are simply competing GFF saddles inside the GDOS kernel.

This move is decisive because it replaces the question

"Why does the entropy of thermal Hawking quanta stop growing?"

by the better-posed question

"Which branch minimizes the operational GFF ledger at a given time?"

Once the question is posed this way, a Page-like transition becomes not mysterious but expected.

48.5 DGI and the factorization puzzle

The standard derivation of the paradox also quietly assumes that one may factorize the total Hilbert space into interior and exterior pieces,

$$\mathcal{H}_{\text{tot}} \stackrel{?}{=} \mathcal{H}_{\text{in}} \otimes \mathcal{H}_{\text{out}}. \tag{302}$$

In ordinary non-gauge quantum mechanics such a split is frequently harmless. In gravity it is not.

The reason is simple. Physical gravitational states obey constraint equations. Schematically,

$$\mathcal{C}_\alpha |\Psi_{\text{phys}}\rangle = 0, \tag{303}$$

where $\mathcal{C}_\alpha$ stands for the Hamiltonian, momentum, and other gauge constraints. These constraints tie together what one naively calls "inside" and "outside" data. The true physical state space is therefore a constrained quotient, not a naive tensor product. In symbol,

$$\mathcal{H}_{\text{phys}} = \frac{\ker \mathcal{C}_\alpha}{\text{gauge redundancy}} \neq \mathcal{H}_{\text{in}} \otimes \mathcal{H}_{\text{out}}. \tag{304}$$

This is the DGI content of the factorization puzzle. The physical microstate volume must be counted only after quotienting by gauge redundancy. Horizon edge data and constraint matching mean that the entangling cut is not an ordinary partition of independent subsystems. Naive factorization therefore overcounts states and misidentifies what counts as independently accessible information.

From the standpoint of GG-6/GDOS, this is not an exotic correction added specifically for black holes. It is exactly what DGI always says. Equivalent microscopic descriptions related by the relevant gauge symmetry are not distinct physical states. In the black-hole problem that principle becomes exceptionally consequential because the horizon is the very place where naive subsystem splits are most tempting.

This also explains a virtue of the modern island story without requiring its ontology as fundamental. Once one accepts that the entanglement wedge of the radiation need not coincide with the naive geometrical exterior, one is already acknowledging that the operational subsystem is subtler than Eq. (302). GG-6/GDOS simply interprets this as the natural consequence of DGI plus OGGEP rather than as the primary ontology.

48.6 The Page curve as a GSSB saddle switch

We may now derive the Page-like behavior in the cleanest GDOS language. Consider two competing branches in Eq. (301).

The first branch, which we may call the "Hawking" or "no-bridge" branch, is the one relevant at early times when the operational radiation subsystem is well described by approximately thermal emission:

$$S_R^{(1)}(t) \approx \eta \left[S_{\rm BH}(M_0) - S_{\rm BH}(M(t)) \right]. \tag{305}$$

The second branch, which we may call the "bridge" or "island-like" branch, caps the accessible entropy by the remaining geometric area ledger:

$$S_R^{(2)}(t) \approx S_{\rm BH}(M(t)). \tag{306}$$

The operational entropy is the smaller of the two:

$$S_R(t) = \min \left\{ S_R^{(1)}(t),\, S_R^{(2)}(t) \right\}. \tag{307}$$

If one sets $\eta \simeq 1$ for a schematic estimate, the crossing occurs when

$$S_{\rm BH}(M_0) - S_{\rm BH}(M_{\rm Page}) = S_{\rm BH}(M_{\rm Page}), \tag{308}$$

hence

$$S_{\rm BH}(M_{\rm Page}) = \frac{1}{2} S_{\rm BH}(M_0). \tag{309}$$

Using Eq. (289) gives

$$M_{\rm Page} = \frac{M_0}{\sqrt{2}}. \tag{310}$$

Substituting into the evaporation law (291) yields

$$t_{\rm Page} = \left(1 - 2^{-3/2}\right) t_{\rm evap} \approx 0.65\, t_{\rm evap}, \tag{311}$$

again up to order-unity corrections from species counts and the exact radiation entropy coefficient.

What matters physically is not the exact numerical prefactor. What matters is the structural conclusion. The Page turnover is the moment at which the dominant operational entropy branch changes. In GDOS language the kernel admits multiple attractor branches, and the realized branch is selected by the GFF minimum. The branch switch is therefore a GRIP event whose middle stage is geometric spontaneous symmetry breaking. The full kernel allows multiple saddles; the realized entropy history selects one.

This is a sharper statement than the loose slogan "information comes out after the Page time." It says that the late-time exterior-accessible algebra is reorganized because the minimizing saddle changes. Information recovery is therefore not a miracle and not an afterthought. It is the thermodynamic selection of the correct operational branch.

48.7 OGGEP and boundary-accessible information

The open-system lesson can now be stated more carefully. OGGEP says that one must never confuse what exists in the bulk with what is reconstructible from a specific operational channel. This point is easy to say abstractly and easy to forget in black-hole discussions, because Hawking

quanta are so seductive as concrete objects. Yet an outside observer never has direct access to the full bulk microstate. The observer has a stream of finite detector clicks, finite timing windows, finite angular resolution, and a history of prior coarse-graining.

It is therefore useful to distinguish three nested notions of information.

a) **Global information.** This is the fine-grained information in the full 6D state $\rho_6(t)$, conserved by exact unitary evolution.

b) **Boundary-accessible information.** This is the information reconstructible from the state $\rho_\partial(t)$ supplied by the CPTP map in Eq. (295).

c) **Detector-record information.** This is the information actually registered by a finite observational protocol, a further coarse-graining of $\rho_\partial(t)$.

The standard paradox often treats (b) and (c) as if they were equal to (a). OGGEP says they are not. This does not trivialize the problem. On the contrary, it sharpens the scientific demand. A genuine resolution must say how these three levels are related and under what circumstances a record protocol can recover the correlations implied by the global unitary state.

The correct statement is therefore the following.

> The information is never fundamentally lost because the bulk evolution is unitary. What can appear lost is only boundary-accessible or detector-accessible information under an incomplete operational channel.

This distinction is exactly the one any open-system theorist would make for a noisy quantum channel or a dissipative mesoscopic device. The novelty in black holes is not the logic but the stakes: the inaccessible environment is partly encoded by geometry and horizon structure rather than by a laboratory heat bath.

We can turn this into a useful entropy inequality. Let ρ_6 denote the exact bulk state and $\rho_\partial = \Pi(\rho_6)$ the accessible boundary state. Then the data-processing inequality implies

$$I_{\mathrm{acc}} \leq I_{\mathrm{bulk}}, \tag{312}$$

for any suitable monotone information measure. In words: operational accessibility cannot increase under the projection channel. If an outside observer appears to lose access to purification data, that is not evidence that the total purification has vanished. It is evidence that the observer has applied or been forced into a channel that throws away correlations.

This is the cleanest way to state what the present framework means by "boundary-accessible information." It is not whatever one would like the outside radiation to encode. It is whatever the operational projection allows one to recover without violating DGI and causal openness.

48.8 Why the firewall is not forced in the GDOS formulation

The firewall argument gains force when three statements are treated as jointly mandatory: smoothness of the horizon, purity of the late radiation, and ordinary monogamy of entanglement

under a sharp interior/exterior split. GG-6/GDOS relaxes the sharpness of the subsystem split rather than the smoothness condition itself.

Because the physical state space is DGI-constrained and because the operational subsystem is boundary-defined rather than ontologically total, the entanglement bookkeeping that drives the firewall contradiction is already modified before one invokes any violent horizon drama. The horizon does not need to become a wall of Planckian drama merely because the outside observer changes branches in the entropy ledger. The operational algebra reorganizes; the bulk need not lose its local smoothness in every semiclassical regime.

This does not mean that all horizon-scale drama is forbidden. It means only that the existence of a paradox is not enough by itself to force one. A putative firewall is not the default consequence of GG-6/GDOS. It would require independent positive evidence that smoothness fails, not merely the existence of an entropy bookkeeping puzzle that has already been reformulated more cleanly by OGGEP and DGI.

48.9 What would count as a genuine resolution

Once the problem is posed correctly, the standards for a genuine resolution become much clearer. A satisfactory resolution should meet five requirements.

(1) Global unitarity. The fine-grained ontology of the full closed system must remain unitary. A theory that simply declares non-unitary destruction of information fundamental has not unified quantum theory and gravity; it has abandoned one of the pillars to save the other.

(2) Operational honesty. The theory must distinguish clearly between the total state, the boundary-accessible subsystem, and actual measurement records. It is not enough to say that "information is in the radiation" if the statement relies on a subsystem identification that violates the theory's own gauge constraints.

(3) Gauge-consistent state counting. The resolution must explain why naive interior/exterior factorization is either valid or invalid. If it is invalid, the theory must provide a principled replacement rather than a hand-waving slogan.

(4) Correct entropy ledger. The resolution must reproduce the early Hawking-like growth of operational entropy and the late Page-like turnover without artificial patching. It must explain why the relevant entropy is area plus bulk entanglement rather than one or the other alone.

(5) No new contradiction worse than the original. A successful resolution should not buy purity by introducing state duplication, uncontrolled nonlocality, or a breakdown of the operational calculus everywhere else in the theory.

These requirements can be summarized in a compact comparison table.

Table 30: What counts as a genuine resolution of the black-hole information paradox in the present paper. The table is intentionally qualitative. The point is not to critic the full literature here, but to state the standards GG-A10-Cosmos adopts. The columns are the four criteria and the rows the approaches judged against them, this framework included and not exempted.

Approach	**Global unitarity**	**Operational honesty**	**Gauge-consistent factorization**	**How the Page curve is obtained**
Pure Hawking thermality	No	Partial	No	Not obtained; radiation entropy grows monotonically
Black-hole complementarity	Intended yes	Partial	Partial	Heuristic; subsystem identification remains delicate
Firewall logic	Intended yes	Partial	No/unclear	Purity enforced at the price of horizon smoothness
Island / QES bookkeeping	Yes	Better	Better	Entropy obtained by minimizing generalized entropy over admissible saddles
GG-6 / GDOS / OGGEP	Yes	Yes	Yes	Page turnover is the GFF-minimizing saddle switch in an open-system ledger

The last row contains the specific claim of the present paper. It does not deny the utility of island technology. Rather, it reinterprets island-like late saddles as effective branches of the GFF ledger rather than as the primary ontology of the theory. The fundamental object is the closed 6D bulk state together with the OGGEP channel. Island-like objects enter only as the most efficient late-time way of writing the operational entropy minimum.

48.10 The clean GDOS resolution in one sentence

The clean GDOS resolution can now be stated in one sentence:

> The black-hole information paradox arises when one mistakes the thermal entropy of a boundary-accessible open subsystem for the fine-grained entropy of the globally closed 6D bulk state, and it is cured when the operational entropy is computed from the DGI-respecting GFF minimum rather than from naive subsystem factorization.

That is the theorem-level statement the section is meant to establish.

48.11 How the GDOS resolution compares with the island/QES program

Because the contemporary literature now phrases the Page-curve solution in the language of islands and quantum extremal surfaces, it is worth stating explicitly how the present section relates

to that program. The answer is that the two descriptions are not enemies, but they assign primacy differently.

In the island/QES formulation one begins from semiclassical gravity and then discovers that the entropy of the radiation is obtained by extremizing and minimizing the generalized entropy functional,

$$S_R(t) = \min_X \operatorname{ext}_X \left[\frac{A[X;t]}{4G_N} + S_{\text{bulk}}[X;t] \right]. \tag{313}$$

This is a major conceptual advance because it shows that the correct entropy is not the naive thermal entropy of the emitted quanta alone. The present framework fully agrees with that conclusion, which is the island result of Penington (2020) and Almheiri et al. (2019), built on the holographic entanglement formula of Ryu and Takayanagi (2006) and its quantum generalization (the holographic background is Maldacena, 1998, 1999). Equation (313) is mathematically the same object as the GFF minimum already written in Eq. (301). What changes is the interpretation of why this object should be the right one.

In the island story, the generalized-entropy rule can look like a remarkable repair made specifically for black holes. In GG-6/GDOS it appears instead as the direct consequence of the universal open-system ledger. The area term is simply the GEP contribution, while the bulk entanglement term is the CEP contribution. The "island" is then not the fundamental ontology but the late-time support of the minimizing operational branch. Put differently: the island program tells us *how* to compute the correct entropy saddle, whereas the GDOS language tells us *why* a branch-minimized generalized-entropy functional had to be the correct answer in the first place.

This distinction can be organized in a compact dictionary.

Table 31: Parallel reading of the modern island/QES program and the GG-6/GDOS open-system reformulation. The table is not meant as a polemic. It identifies the level at which each description places the primary ontology. Each row poses one question and answers it in both languages. The table is a translation and not a claim of priority: where the two accounts agree the agreement is the point, and where they differ the difference is stated rather than resolved.

Question	**Island / QES language**	**GG-6 / GDOS language**
What is fundamental?	Boundary-unitary or semiclassical-plus-QES description	Globally unitary 6D bulk plus OGGEP projection
What is the relevant entropy?	Generalized entropy extremized over QES candidates	GFF ledger minimized over admissible operational branches
Why is naive radiation entropy wrong?	Entanglement wedge of radiation is larger than naive exterior region	Boundary-accessible subsystem is an open CPTP image and need not coincide with total ontology
Why is factorization subtle?	Interior/exterior split is replaced by a nontrivial entanglement wedge	DGI forbids naive gauge-breaking factorization across the horizon
What happens at the Page time?	A different QES dominates	A different GFF saddle dominates; branch switch is GSSB

The practical consequence is important for the logic of the paper. We do not need to deny the validity of island calculations in the semiclassical regime. We only need to deny that their geometric bookkeeping should be mistaken for the deepest layer of ontology. The deepest layer, in the present program, is the closed six-dimensional state together with the causal projection rules that determine what a boundary observer can actually access.

48.12 Scrambling, Hayden–Preskill logic, and operational information return

A second place where operational clarity matters is the question of *when* information should begin to become recoverable from the outgoing channel. The modern literature often phrases the answer in the Hayden–Preskill language: an "old" black hole past the Page time behaves as a rapidly scrambling mirror, so additional information thrown in can in principle be reconstructed from a modest extra amount of radiation after a scrambling delay. The present framework can absorb that statement, but only after translating it into OGGEP terms.

The relevant time scale is the scrambling time,

$$t_{\rm scr} \sim \frac{\beta_H}{2\pi} \ln S_{\rm BH} = \frac{1}{2\pi T_H} \ln S_{\rm BH} \approx 4G_N M \, \ln S_{\rm BH}, \tag{314}$$

where $\beta_H = 1/T_H$ and we used Eq. (288). The logarithm is the key feature: scrambling is fast compared with the full evaporation time, but it is not instantaneous. In operational language this means that the boundary channel has a finite mixing time before newly inserted information can appreciably affect the outgoing accessible record.

One may formalize the distinction with three nested mutual informations,

$$I_{\rm det}(R : {\rm Ref}) \le I_{\partial}(R : {\rm Ref}) \le I_6(R : {\rm Ref}), \tag{315}$$

where Ref denotes an abstract reference system entangled with the infalling matter, I_6 is the exact bulk mutual information, I_∂ is the mutual information in the boundary-accessible state, and $I_{\rm det}$ is the information actually extracted by a finite measurement protocol. Equation (315) is simply the data-processing inequality applied twice, but it cleanly states what "information return" can and cannot mean.

Information return does *not* mean that the detector record instantaneously reproduces the exact microscopic infalling state. It means that after the Page transition and after the additional scrambling delay, the operational channel is no longer locked into the purely thermal branch; correlations with the reference system can begin to reappear in principle in the boundary-accessible radiation. This is the open-system version of the Hayden–Preskill idea.

The reason this matters for cosmology is not that we expect to decode astrophysical black holes any time soon. The reason is conceptual. A theory that cannot state clearly the difference between global purification, boundary accessibility, and detector-level reconstruction will speak ambiguously not only about black holes but also about horizon physics, inflationary perturbations, and late-time relics. Scrambling therefore serves here as a diagnostic of the framework's operational discipline.

48.13 Why remnants, baby universes, and final bursts are not needed in the minimal module

Another virtue of the GDOS reformulation is that it removes the pressure to use a remnant-like cure. In many older discussions one tried to preserve purity by suggesting that a tiny final object survives evaporation and secretly stores the information. The problem is immediate. If a remnant of mass $M_{\rm rem} \sim M_{\rm Pl}$ had to store the entire initial Bekenstein–Hawking entropy, the number of remnant microstates would have to scale as

$$\mathcal{N}_{\rm rem} \sim \exp[S_{\rm BH}(M_0)]\,, \tag{316}$$

which becomes catastrophically large for macroscopic M_0. That suggests a huge low-mass state degeneracy and the usual pair-production catastrophes that plague remnant scenarios.

In the present framework this problem is unnecessary because the information was never required to be parked in a final tiny object. The global bulk state remains unitary throughout, and the boundary-accessible entropy is governed by the branch-minimized GFF ledger throughout. A remnant is therefore not needed as an information warehouse. Nor is it necessary to invoke baby-universe branching as the primary bookkeeping device. Such possibilities are not logically excluded by the constitution, but they are not demanded by the minimal module, and the present paper has no need to purchase consistency by introducing them.

Exactly the same remark applies to final explosive-burst cures. If one imagines that a nearly thermal evaporation history remains featureless until the last instant and then somehow returns the full information content in a spectacular burst, one has again waited too long. The Page-curve logic says that the reorganization of accessible correlations begins around $t_{\rm Page}$, not only at $t_{\rm evap}$. In the present language that means the minimizing saddle has already switched well before the final stages. The correction is distributed over the late channel structure, not postponed to a miraculous terminal event.

This is another way to see why the open-system formulation is cleaner. Once the correct branch structure is identified, one is no longer forced into desperate end-stage mechanisms. The paradox is resolved by the evolution of the entropy ledger itself.

48.14 A compact falsification dictionary for future black-hole observations

Although direct experimental access to black-hole information recovery is extremely limited, the framework should still state what sorts of observations would push on it. The right standard is not "measure the full Page curve of an astrophysical black hole tomorrow," which is unrealistic, but rather: identify which classes of observations would support or stress the assumptions entering the minimal module.

Table 32: Compact observational dictionary for the black-hole information block. The entries are deliberately phrased at the level of principle rather than as a near-term mission proposal. The columns give the observable class, what the minimal module expects, and what would stress it. The entries are phrased at the level of what an instrument could actually decide, which is coarser than the theory can state, and that coarseness is deliberate.

Observable class	**What the minimal module expects**	**What would stress the module**
Ringdown and exterior classical geometry	Ordinary GR-like exterior relaxation in the semiclassical regime	Robust evidence that the exterior classical regime fails far from the endpoint without an accompanying open-system explanation
Near-horizon drama	No firewall forced by entropy bookkeeping alone	Unambiguous evidence that smooth semiclassical horizons are generically impossible even when DGI and operational openness are respected
Late radiation correlations	In-principle recovery after the Page transition and scrambling delay, not exact early thermality forever	A theorem or observation implying exact monotonic thermal entropy of the accessible channel all the way to completion
Subsystem factorization	Gauge-constrained nonfactorization near the horizon	A consistent fully gravitational factorization proof that makes the naive split exact without changing the entropy result
Generalized entropy / GFF bookkeeping	Area plus bulk-entanglement ledger controls the branch minimum	A demonstrable contradiction between the generalized-entropy ledger and the observed or calculated Page behavior

This dictionary again shows the style of the program. The point is not to pretend that every item is experimentally settled. The point is to make clear what the framework would count as support and what it would count as trouble. Even in an area as hard as black-hole information, the theory should publish its kill criteria rather than hide behind abstraction.

48.15 Why the black-hole block matters for the rest of cosmology

It may be tempting to think that the black-hole information block is conceptually elegant but observationally disconnected from the earlier sections on inflation, dark matter, dark energy, BBN, and structure formation. That would be a mistake. In the present program the black-hole block is where the constitutional claims of the earlier parts are exposed to their hardest internal audit.

Part IV argued that the universe should be read as an open boundary subsystem of a higher-dimensional closed geometry. Part V argued that chemical and nuclear histories should be written as GDOS evolutions on constrained state manifolds with explicit ports and conserved ledgers. Part VI argued that the CMB and the growth of structure are not just collections of fit parameters but macroscopic readouts of a DGI-respecting open-system flow. The black-hole problem tests whether

these slogans remain coherent when the relevant subsystem is separated from the rest of the world by the strongest causal barrier in classical relativity.

If the framework were to fail here, the consequences would spill backward into the whole paper. For example, if one were forced to abandon the distinction between global closure and operational openness in the black-hole case, then the OGGEP language used earlier for inflationary perturbations and cosmological horizons would become much less convincing. If naive factorization survived even in gravity, then the special role assigned to DGI in the perturbation and measurement sectors would be weakened. If the correct entropy ledger for horizons turned out not to be the GFF minimum, then the earlier use of GFF as the scalar ranking functional for cosmological branches would lose one of its strongest nontrivial supports.

In that sense, black holes are the consistency check that keeps the rest of GUT-Cosmos honest. They force the paper to answer three foundational questions that might otherwise remain hidden under phenomenology.

Question 1: What counts as the system? In BBN the answer was the nuclear abundance network inside a cooling radiation bath. In the CMB it was the baryon–photon plasma inside a causal Hubble region. In a black hole the answer is trickier because the outside observer is causally prevented from accessing part of the geometry. If the theory cannot define the system there, it has not yet earned the right to claim a universal SIP/GDOS formalism.

Question 2: What counts as information? In the early universe, "information" often sounds harmless: it is encoded in curvature perturbations, transfer functions, and phase correlations. In black holes the word becomes dangerous because one is no longer allowed to slide casually between microscopic purity and macroscopic observability. The black-hole block therefore clarifies that information is always protocol-relative on the boundary and only globally absolute in the closed bulk ontology.

Question 3: What counts as entropy? In ordinary cosmological applications one can often get away with a simplified entropy language: comoving entropy conservation, photon entropy, or effective fluid entropy. Black holes remove that luxury. The correct entropy is not optional bookkeeping but the whole content of the paradox. If the theory identifies the wrong ledger here, it cannot be trusted when it later speaks about attractors, coincidence, or relic abundances.

These questions can be summarized in one compact balance law for the operational sector. Let S_∂ denote the boundary-accessible entropy, let $S_{\rm env}$ denote environment and horizon-hidden coarse-grained entropy, and let $I_{\rm corr}$ denote the cross-correlations discarded by a finite observational protocol. Then the schematic coarse-grained entropy balance reads

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(S_\partial + S_{\rm env} - I_{\rm corr}\right) \geq 0. \tag{317}$$

The point of Eq. (317) is not its exact microscopic derivation inside this paper; the point is the logic it captures. Accessible entropy can rise even while global purity is preserved, because the information hidden in discarded correlations and inaccessible sectors is not destroyed but merely

moved out of the chosen protocol. This is exactly the same logic that underlies decoherence in laboratory open systems, but horizon geometry makes the partition especially severe.

The black-hole block also matters because it disciplines language about "paradox." Earlier in the paper several observational tensions were discussed: the Hubble tension, the S_8 tension, and the JWST early-galaxy problem. Those are disagreements between inference chains, calibrations, or sectors of phenomenology. The information paradox is different. It is not a tension between two data sets but a tension between levels of description. Solving it therefore teaches a methodological lesson: many apparent contradictions in theoretical cosmology arise from asking one descriptive layer to carry obligations that belong to another layer.

That lesson flows back into the whole paper. The Hubble tension can then be reread as an inference-dictionary problem rather than as a literal contradiction of background dynamics. The CMB gauge issue can be reread as a DGI coarse-graining problem rather than as an ambiguity about physical observables. Even the dark-sector problem can be reread as a source-ledger problem rather than as a license to add arbitrary new fluids. Black holes sharpen this style of reasoning because the danger of mixing descriptive layers is maximal there.

For that reason the present section is not a decorative appendix. It is one of the places where the paper most clearly shows what it means to keep a single constitutional grammar from the beginning of the universe to the late fate of information-bearing gravitational systems.

48.16 What the present framework claims, and what it does not

It is important not to overclaim.

The framework does *claim* the following:

1. There is no fundamental information destruction at the GG-6 bulk level.
2. Pure-to-mixed behavior in the exterior description is expected because the boundary observer sees an open subsystem.
3. The correct coarse-grained entropy ledger is generalized entropy, re-read as GFF.
4. The Page turnover is a GDOS saddle switch and therefore a GSSB event.
5. Naive interior/exterior Hilbert-space factorization is not constitutionally admissible once DGI is taken seriously.

The framework does *not* yet claim the following:

1. a complete microscopic 6D basis for all black-hole microstates;
2. an exact numerical graybody-factor derivation inside the present paper;
3. a full detector-level reconstruction protocol for practical information retrieval from an evaporating astrophysical black hole.

These are not evasions. They are the correct boundaries of the present cosmology paper. The goal here is to show that the paradox is no longer structurally fatal once the constitutional layer

of GG-6 is enforced. The dedicated companion (Arisaka, 2026, C3-BH) pushes the microstate and detector questions much further — deriving the exact $A/4$ coefficient from induced gravity (its Theorem *T-BH-1*), the EHT-shadow and ringdown predictions, and the four laws of black-hole mechanics as GDOS statements.

48.17 Interim conclusion of the information-paradox section

The black-hole information paradox is not an embarrassment left over after the rest of cosmology has been solved. It is a precision test of what the framework means by information, entropy, gauge redundancy, and measurement. The GDOS resolution is clean because it does not require special pleading. The same open-system principles used elsewhere in the program apply here too. A black hole is a particularly severe open system, not an exception to the theory.

This conclusion also gives the correct bridge to the next section. Once one understands black holes as extreme information-processing objects rather than as isolated curiosities, primordial black holes become a natural cosmological laboratory. They are not only compact objects; they are possible fossilized relics of the primordial perturbation sector itself. That is why the next section turns from the information paradox to primordial black holes.

49 Primordial black holes

49.1 Historical background

Primordial black holes were introduced by Hawking in 1971 (Hawking, 1971), who proposed that large density perturbations in the very early universe could collapse directly into black holes without proceeding through stellar evolution. The resulting mass spectrum spans many orders of magnitude, from below evaporation thresholds at $\sim 10^{15}$ g to supermassive scales relevant to galactic dynamics. The present section asks which slice of that mass spectrum, if any, is consistent with the inherited ledger and could play a non-trivial role in dark-matter, gravitational-wave, or seed-black-hole physics without conflicting with the rest of the cosmic audit.

Proposition 49.1 (Minimal PBH scarcity theorem (T-Cosmos-14))**.** *If the smooth GUT-adjacent inflationary branch (the derived Starobinsky/ξ plateau) and the Kaxion dark-matter branch are both retained, then the primordial black-hole abundance remains at most a trace component unless an additional narrow small-scale enhancement can be derived from the same inherited scalar ledger. In the minimal smooth module,*

$$\beta(M) \sim \exp\left[-\frac{\delta_c^2}{2\sigma^2(M)}\right]$$

is exponentially tiny in the relevant mass windows — the result depends on the smoothness of the plateau spectrum, not on the absolute inflationary scale, so it is unchanged by the GUT-adjacent reading.

Primordial black holes are attractive because they compress several cosmological questions into one object. They probe the small-scale primordial spectrum, the collapse threshold of the

early universe, the thermal background history, dark-matter accounting, and late compact-object phenomenology. In many frameworks they are used opportunistically: if particle dark matter becomes difficult, PBHs are brought in as an alternative. In the present framework the logic is different. The Kaxion already occupies the primary cold-dark-matter slot. PBHs are therefore not introduced as a rescue sector. They are introduced as a stress test of the perturbation ledger already developed in Parts II through IV.

This change of logic matters. A framework that already has a constrained dark-matter candidate does not need PBHs to solve the dark-matter problem. That means it can ask a much cleaner question: does the inherited GG-6/GDOS cosmology naturally produce a substantial primordial black-hole population, a tiny trace population, or essentially none at all?

The answer this paper suggests is sharp but nuanced. The minimal smooth GG-6 inflationary module, together with the red scalar spectrum and the absence of a large small-scale enhancement, predicts that PBHs are *not* a dominant dark-matter component. In the minimal ledger their abundance is exponentially suppressed. At the same time, the framework does not forbid subdominant narrow windows associated with specific transition scales. Those windows become useful because their mass scale maps directly to formation temperature and hence to specific epochs in the inherited cosmic chronology.

49.2 Formation windows

The standard logic of primordial-black-hole formation begins when an overdense region re-enters the horizon during the radiation era. If the overdensity exceeds a threshold δ_c, pressure support fails to prevent collapse and a black hole forms. The relevant mass is tied to the horizon mass at re-entry:

$$M_H = \frac{4\pi}{3}\rho\, H^{-3}. \tag{318}$$

During radiation domination,

$$\rho = \frac{\pi^2}{30} g_* T^4, \qquad H = 1.66\, g_*^{1/2} \frac{T^2}{M_{\rm Pl}}, \tag{319}$$

so Eq. (318) becomes

$$M_H(T) \simeq 9.4 \times 10^{25}\,\mathrm{g} \left(\frac{106.75}{g_*}\right)^{1/2} \left(\frac{\mathrm{TeV}}{T}\right)^2. \tag{320}$$

The PBH mass is usually a fraction γ_c of the horizon mass,

$$M_{\rm PBH} \simeq \gamma_c M_H, \qquad \gamma_c \sim 0.1\text{–}0.2. \tag{321}$$

Equation (320) is extremely informative because it translates a cosmic temperature into a PBH mass window. At a TeV-scale exit feature — the lock crossing, were it not prompt — one expects

$$M_{\rm PBH}^{(\rm lock)} \sim 10^{24}\,\mathrm{g}, \tag{322}$$

whereas the QCD epoch around $T \sim 0.1$–$0.2\,\mathrm{GeV}$ maps to

$$M_{\mathrm{PBH}}^{(\mathrm{QCD})} \sim \mathcal{O}(M_\odot). \tag{323}$$

This is why PBH phenomenology is so sensitive to the precise epoch that generates any small-scale enhancement.

The next step is the collapse fraction. If the primordial density contrast smoothed on mass scale M is Gaussian with variance $\sigma^2(M)$, the formation fraction at re-entry is

$$\beta(M) = \int_{\delta_c}^{\infty} \frac{\mathrm{d}\delta}{\sqrt{2\pi}\sigma(M)} \exp\left[-\frac{\delta^2}{2\sigma^2(M)}\right]. \tag{324}$$

For the tail relevant to PBH production, $\sigma \ll \delta_c$, one may use the asymptotic form

$$\beta(M) \approx \frac{\sigma(M)}{\delta_c\sqrt{2\pi}} \exp\left[-\frac{\delta_c^2}{2\sigma^2(M)}\right]. \tag{325}$$

The present-day PBH fraction inherits the standard radiation-era growth factor,

$$f_{\mathrm{PBH}}(M) \propto \beta(M)\, M^{-1/2}, \tag{326}$$

up to the usual g_* and equality-era prefactors. The crucial point is the exponential tail in Eq. (325). PBHs are not produced efficiently unless the primordial variance on the relevant scale is dramatically larger than on CMB scales.

That single mathematical fact already indicates what the minimal GG-6 prediction must be. The scalar spectrum inherited from Part IV and audited in Part VI is red and smooth. It does not contain the large small-scale spike required to make $\beta(M)$ appreciable. The natural expectation is therefore exponentially small PBH production.

49.3 A minimal quantitative estimate

To see how sharp the suppression is, one need not run a full numerical collapse code. A tail estimate suffices. Take a representative radiation-era threshold

$$\delta_c^{(\mathrm{rad})} \sim 0.4\text{–}0.5. \tag{327}$$

Now suppose the minimal smooth spectrum gives a small-scale variance $\sigma(M) \sim 10^{-2}$, which is already far above the ordinary CMB-scale amplitude but still nowhere near the values needed for strong PBH production. Then

$$\frac{\delta_c^2}{2\sigma^2} \sim \frac{0.45^2}{2\times 10^{-4}} \sim 10^3, \tag{328}$$

and Eq. (325) gives

$$\beta(M) \sim \exp(-10^3), \tag{329}$$

up to algebraic prefactors that are negligible compared with the exponential. This is effectively zero.

Even if one imagines a much larger variance, say $\sigma \sim 0.05$, the exponent becomes

$$\frac{\delta_c^2}{2\sigma^2} \sim 40, \tag{330}$$

so the formation fraction is still tiny. One therefore needs not just a modest enhancement but a genuinely strong feature in the small-scale spectrum to make PBHs important. That feature is absent from the theorem-level minimal module already written into this paper.

This simple estimate is sufficient to justify a central statement of Part VII:

> In the minimal smooth GG-6 cosmology, PBHs are not a dominant dark-matter component. Kaxion dark matter remains primary, and any PBH population is at most a small trace component unless an additional narrow small-scale enhancement can be derived from the same geometric ledger.

49.4 Formation windows tied to the inherited thermal history

Even though the minimal abundance is tiny, the horizon-mass formula remains informative because it maps hypothetical features to concrete epochs. Three windows are particularly useful conceptually.

(i) Lock-crossing window. If the prompt lock crossing were to carry a narrow spike or short-lived softening of the equation of state (the crossing is prompt, so none is expected — this is the contingency, not the prediction), the corresponding PBHs would lie near the TeV horizon mass in Eq. (322). These are asteroid-scale objects. They are well above the evaporation threshold and would survive to the present epoch.

(ii) Electroweak-like window. Any narrow feature tied to physics near the electroweak-to-TeV interval would shift the mass upward by roughly two orders of magnitude in temperature and hence four orders in mass. Such objects would still be very far from stellar masses and would fall primarily into lensing, dynamical, and accretion-style searches.

(iii) QCD softening window. The QCD crossover lowers the effective sound speed of the plasma and is the classic place where stellar-mass PBHs become easiest to form. Equation (323) explains why. The relevant horizon mass is near the solar-mass range, precisely where merger catalogs and microlensing constraints become powerful.

The reason to list these windows is not to claim that the framework predicts all of them. It is to show how any future PBH hint could be translated back into the thermal and dynamical history of the early universe. The PBH mass is a timestamp written in compact-object form.

49.5 The integer-ledger mass ladder

The horizon-mass map of Eq. (320) acquires extra structure from the inherited ledger. If a narrow small-scale feature does arise, it does not sit at an arbitrary scale: the same OGN-5 integer

$J = 7$ that fixes the recombination switch $z_* = e^7$ of Part VI may also space the throat resonances, so a primordial-black-hole mass spectrum, if one exists, would fall on a geometric ladder

$$M_n^{\rm PBH} \;\propto\; e^{-7n}, \qquad n = 0, 1, 2, \ldots,$$

the gravitational-sector image of the very integer that organizes the cosmological corridor. This spectral ladder is carried by the companion black-hole paper (Arisaka, 2026, C3-BH) at the grade that paper gives it — a model-level conjecture, not a theorem of the cohort, whose exponent is not derived there and whose absolute anchor floats for want of a cosmology-side input — and it is registered there as the named falsifier F-BH-4, which fires only on a well-populated primordial-black-hole mass function whose spacing is inconsistent with the e^7 ladder; a single detection does not trigger it. (An earlier version of this sentence called the ladder the companion's "dedicated prediction"; the companion's own grade is the one carried here.) The cosmological point here is narrower but fully consistent: the minimal module makes such a feature scarce, with the abundance exponentially suppressed, so the ladder predicts *where* any rare PBH population would sit, not that the population is large.

49.6 Connection to early background dynamics

The next question is whether anything already derived in Parts IV and VI can generate the strong small-scale enhancement PBH formation requires.

The answer is mostly no in the minimal module, and the reasons are important.

First, the slow-smooth inflationary sector is deliberately smooth. The entire point of the geometric plateau in Part IV was to avoid a jagged inflaton potential landscape with arbitrary features. A smooth plateau gives a red and almost scale-invariant scalar spectrum, not a narrow blue spike.

Second, no early-time source survives in Part VI. The early-dark-energy lobe an earlier version carried, with characteristic redshift

$$z_c \approx 3500, \tag{331}$$

is withdrawn ($f_{\rm EDE} = 0$; a detection at $f_{\rm EDE} \gtrsim 0.02$ is a registered kill), and even had it stood it sat well after the inflationary exit, on scales relevant to the sound horizon and the matter transfer function, not to horizon-scale collapse of ultra-small modes: it could not have functioned as a PBH factory.

Third, the same growth logic that helps the JWST problem points away from PBHs as the main explanation of early supermassive black holes. Part VI argues that earlier first light and direct-collapse baryonic seeds arise from a lower effective collapse threshold for high-redshift gas — the EDE small-scale boost that used to head that list is withdrawn (§44.3), which strengthens rather than weakens the present point, since what remains is precisely the baryonic part. Those mechanisms are later-time and baryonic. They therefore offer a cleaner route to early supermassive seeds than invoking a large primordial black-hole abundance.

This is the important cross-consistency statement. The same framework that helps early galaxies does not need to use PBHs as a second rescue story. It prefers direct-collapse baryonic seeds plus a

modest acceleration of early structure assembly.

49.7 PBHs versus Kaxion dark matter

The dark-matter accounting sharpens the same conclusion. Part V already occupied the cold-dark-matter slot with the Kaxion, whose relic abundance is tied to the compact-angle sector and the inherited misalignment window. If PBHs were also to supply an order-unity fraction of dark matter, the total matter ledger would be overbooked unless the Kaxion abundance were reduced correspondingly. But reducing the Kaxion abundance would destroy the clean compact-angle interpretation that motivated it in the first place.

This leads to a simple hierarchy of hypotheses:

1. **Minimal hypothesis (Derived-conditional):** Kaxions supply the dominant cold-dark-matter abundance; PBHs are negligible.
2. **Allowed but nonminimal hypothesis:** a narrow primordial feature creates a tiny PBH subcomponent in a specific mass window, without disturbing the primary Kaxion ledger.
3. **Disfavored hypothesis:** PBHs supply most or all of dark matter. This is not impossible abstractly, but it is not the natural outcome of the smooth GG-6 module already developed.

This ranking is more scientifically useful than a vague statement that PBHs are "possible." It tells the reader what the framework actually expects.

49.8 Observational searches

The observational program for PBHs may be organized by mass.

Evaporation regime. Black holes lighter than about

$$M_{\rm evap} \sim 5 \times 10^{14}\,{\rm g} \tag{332}$$

have evaporated by the present epoch. This follows from the cubic evaporation-time law

$$t_{\rm evap}(M) \approx 6.2 \times 10^{-27}\,{\rm s}\left(\frac{M}{\rm g}\right)^3. \tag{333}$$

Any PBH population formed far below this mass contributes today only through relic evaporation products and the constraints attached to them.

Asteroid-to-planetary regime. The lock-crossing window of Eq. (322) lies in the broad asteroid-mass corridor. Such objects survive, do not radiate appreciably today, and are primarily tested by lensing, dynamical heating, accretion bounds, and indirect compact-object searches.

Stellar-mass regime. The QCD window lies near the solar-mass scale and is therefore constrained by microlensing, merger catalogs, and stochastic compact-binary backgrounds. Because Part VI already offers a direct-collapse route to early heavy seeds, a strong observational case for a dominant stellar-mass PBH population would place pressure on the minimal cosmological module of this paper.

These qualitative regimes can be summarized compactly.

Table 33: Illustrative PBH mass windows relevant to GG-A10-Cosmos. The first column is the early-universe epoch or feature; the last column states the role the present framework assigns to the corresponding PBH population. The columns run from the formation epoch through the characteristic mass and the search handles to the role the window would play here. The windows are illustrative: they say where a primordial black hole population would have to sit to matter, not that the framework predicts one.

Formation epoch / feature	**Characteristic mass**	**Primary search handles**	**Role inside minimal GG-6/GDOS**
TeV lock-crossing feature	$\sim 10^{24}\,\mathrm{g}$	lensing, dynamics, accretion-style bounds	not required; at most narrow subcomponent if a transition feature exists
Electroweak-scale feature	10^{27} to $10^{29}\,\mathrm{g}$	lensing and dynamical bounds	same: not primary dark matter in the smooth module
QCD softening window	$\mathcal{O}(M_\odot)$	microlensing, merger catalogs, stochastic backgrounds	useful test window, but not the preferred explanation of early SMBH seeds
Sub-$5\times10^{14}\,\mathrm{g}$ evaporation regime	evaporated by today	high-energy evaporation products	not relevant to present dark matter; only indirect historical constraints

49.9 Bottom line of the PBH section

The present paper should therefore state the PBH conclusion in deliberately asymmetric language. Primordial black holes are not the engine that makes GG-6 cosmology work. They are an auxiliary, high-leverage test of the small-scale primordial spectrum and the early thermal history. The minimal smooth module predicts that they are scarce. A tiny subcomponent is not excluded a priori, especially in a narrow transition window, but a dominant PBH dark-matter universe is not the natural expectation once the Kaxion sector and the smooth inflationary ledger are accepted.

That conclusion also clarifies the bridge to the next section. The same primordial perturbation sector that could in principle seed PBHs also seeds the primordial tensor spectrum. In the minimal module both are featureless for the same reason: the inflationary source ledger is smooth, so PBHs are exponentially scarce and the tensor spectrum is a plain power law at the fixed amplitude $r = 0.0038$. Primordial gravitational waves are therefore the cleanest final audit of the same logic.

50 Primordial gravitational waves

50.1 Historical background

The proposal that the early universe leaves behind a present-day stochastic background of gravitational waves with a roughly scale-invariant spectrum is due to Starobinsky in 1979 (Starobinsky, 1979), who computed the amplitude of relict tensor modes in a de Sitter phase and argued that they would survive to the present as a background detectable through CMB B-mode polarization or future direct interferometry. The amplitude is set by the inflationary energy scale, and the derived GUT-adjacent plateau fixes it: the Starobinsky relation at the reheating-matched count. The present section converts the derived launch into a sharp, two-sided falsification target.

Theorem 50.1 (Tensor target (T-Cosmos-15, revised))**.** *The same derived plateau that computes the scalar tilt fixes the primordial tensor sector at a firm central value:*

$$r = 0.0038 \quad (12/N^{*2} = 0.0043 \text{ at leading order}).$$

This is one of the cleanest near-term falsifiers of the minimal cosmology branch, and it is two-sided: a tensor measurement inconsistent with this value beyond the launch's spread 0.0037–0.0040*, in either direction — a detection well above it, or an exclusion well below it — rules out the derived launch rather than merely requiring a small parameter adjustment.*

The primordial tensor sector is the sharpest observational kill switch in the entire cosmology paper. Scalar-sector anomalies can sometimes be patched by late-time source terms, transfer-function subtleties, or environmental systematics. Primordial tensors are much less forgiving. Once the inflationary energy scale is fixed, the tensor amplitude is effectively fixed. This is why Part IV invested its effort in deriving the plateau class and its placement — one e-fold below the GUT edge, with the amplitude the sector's one declared input.

The consequence is sharp. If the minimal GG-6 inflationary module is correct, primordial tensor modes sit at exactly $r = 0.0038$ — a factor of nine below the present bound (BICEP/Keck Collaboration, 2021), and several times above the projected Simons-Observatory/LiteBIRD floor ($\sigma(r) \simeq (1\text{–}2) \times 10^{-3}$; Ade et al., 2019; LiteBIRD Collaboration, 2023). This is one of the framework's strongest falsification handles, because it ties a concrete, two-sided observational statement directly to the derived launch.

50.2 Tensor modes from the derived plateau

For a transverse-traceless metric perturbation h_{ij} on a spatially flat background,

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + a^2(t)\,[\delta_{ij} + h_{ij}]\,\mathrm{d}x^i \mathrm{d}x^j, \qquad \partial_i h_{ij} = 0, \qquad h_{ii} = 0, \tag{334}$$

each polarization mode obeys the standard equation

$$u_k'' + \left(k^2 - \frac{a''}{a}\right) u_k = 0, \tag{335}$$

where u_k is the canonically normalized tensor variable and primes denote conformal-time derivatives.

The primordial tensor power spectrum at horizon exit is

$$\mathcal{P}_t(k) = \frac{2H_*^2}{\pi^2 M_{\rm Pl,red}^2}, \tag{336}$$

while the tensor-to-scalar ratio is

$$r = \frac{\mathcal{P}_t}{\mathcal{P}_\mathcal{R}} = \frac{2H_*^2}{\pi^2 M_{\rm Pl,red}^2 A_s}. \tag{337}$$

Using $H_*^2 \simeq V_*/(3M_{\rm Pl,red}^2)$ gives the equivalent expression already derived in Part IV,

$$r = \frac{2V_*}{3\pi^2 A_s M_{\rm Pl,red}^4}. \tag{338}$$

The key input is the derived plateau of Part IV — class forced by the boundary law, height fixed by the one declared input:

$$V_0^{1/4} = 8.1 \times 10^{15}\ \text{GeV} \quad (V_*^{1/4} = 8.0 \times 10^{15}\ \text{GeV at the pivot}), \qquad N^* = 52.6, \tag{339}$$

and substituting into Eq. (338), together with $A_s \simeq 2.1 \times 10^{-9}$, yields

$$r = 0.0038 \quad \big(12/N^{*2} = 0.0043 \text{ at leading order}\big), \tag{340}$$

in agreement with Eq. (73). (Evaluated at the retired lock-scale height, this chain would give $r \sim 10^{-52}$.)

It is useful to express the same result via the inflationary Hubble scale. From $3M_{\rm Pl,red}^2 H_*^2 \simeq V_*$ one obtains

$$H_* = H_1 \simeq 1.5 \times 10^{13}\,\text{GeV} \tag{341}$$

(the lock-scale combination $M_{\rm lock}^2/\sqrt{3}\,M_{\rm Pl,red} \simeq 3.5 \times 10^{-12}\,\text{GeV}$ survives as the curvature scale of the middle basin — the terminus seesaw, not the launch). The tensor target is therefore not a side remark. It is a direct algebraic consequence of the same derived plateau that computes the scalar tilt.

50.3 Why the primordial signal is exactly one signal: vacuum tensors, and silence from the transitions

There are in fact two statements rather than one.

The first is the vacuum tensor amplitude of Eq. (339): the derived plateau fixes $H_* \simeq 1.5 \times 10^{13}$ GeV, hence $r = 0.0038$ — modest against the old GUT-scale folklore, decisive against the coming instruments.

The second concerns the nature of the exits. The graceful exit of Part IV is *smooth*, and the lock crossing below it is *prompt* and barrier-free. Neither is a violent first-order transition with highly relativistic bubble collisions, so neither generates an additional stochastic gravitational-wave

background beyond the vacuum tensor spectrum — and the absence of any TeV-epoch transition background is itself a named falsifier: a detected background at the lock scale would break the prompt-crossing bound.

This is an important internal consistency point. One cannot claim both that the descent is smooth-and-prompt and that it produces a large extra primordial gravitational-wave signal. The minimal paper does not try to have both: one clean vacuum tensor line at $r = 0.0038$, and silence from the transitions.

50.4 B-mode polarization

At the observable level primordial tensor modes are most cleanly sought through CMB B-mode polarization. The total B-mode signal may be written schematically as

$$C_\ell^{BB} = C_{\ell,\mathrm{lens}}^{BB} + C_{\ell,\mathrm{prim}}^{BB}(r) + C_{\ell,\mathrm{fg}}^{BB}, \tag{342}$$

where the three terms are lensing-converted B-modes, primordial tensor B-modes, and foreground residuals.

In ordinary high-scale inflation, the primordial term can generate a recombination bump near $\ell \sim 80$ and a reionization bump at very low multipoles. In the present framework r is small but fixed:

$$\begin{aligned} C_{\ell,\mathrm{prim}}^{BB}(r = 0.0038) &: \text{a recombination bump at } \ell \sim 80, \\ &\quad \text{several times above the Simons-Observatory/LiteBIRD floor} \end{aligned} \tag{343}$$

— small, but decisively measurable. The scientific meaning is very sharp:

> After delensing, the primordial bump at the $r = 0.0038$ level is the minimal GG-6 inflationary module's signature: the Simons Observatory and LiteBIRD must find it — and must find it at that amplitude, not another.

This statement is stronger than the generic phrase "small r." It means that successful delensing must reveal a primordial bump at this level, within the launch's own spread ($r = 0.0037$–0.0040 across the reheating range of the matching that feeds Eq. (75), $N^* \in [51.4, 53.1]$): a robust primordial tensor signal outside $[10^{-3}, 10^{-2}]$, or its absence once the instruments reach $\sigma(r) \sim 10^{-3}$, rules the minimal module out, and a signal inside that window but inconsistent with 0.0038 at the instrument's precision refutes the count. (An earlier version of this sentence, written for the retired TeV-scale reading in which $r \sim 10^{-52}$, said the opposite — that delensing should reveal no bump — and stood here beside its replacement; it is withdrawn.)

To keep the paper crisp, it is worth stating the kill rule in plain terms:

$$\begin{aligned} &\text{A tensor measurement inconsistent with } r \simeq 0.004 \text{ in either direction} \\ &\quad \Longrightarrow \quad \text{the minimal derived GG-6 launch is false.} \end{aligned} \tag{344}$$

That is not rhetorical flourish. It follows because $r = 12/N^{*2}$ carries no dial: moving r means moving N^*, and N^* is pinned by the reheating match and the promptness of the lock crossing — so

a discrepant tensor amplitude breaks the derived chain itself, not a plateau parameter, and unravels the minimal module back at Part IV.

50.5 Interferometer and CMB cross-checks

The tensor audit should not stop at the CMB, because the primordial tensor background is often discussed in frequency language rather than multipole language. For an approximately scale-invariant tensor spectrum, the present-day stochastic background is roughly

$$\Omega_{\mathrm{GW},0}h^2 \sim \frac{\Omega_{\mathrm{r}}h^2}{24}\, A_s\, r \simeq 1.7\times 10^{-6}\, A_s\, r, \tag{345}$$

with $\Omega_{\mathrm{r}}h^2 \simeq 4.1\times 10^{-5}$ the present relativistic density, up to transfer-function and frequency-dependent corrections. With $A_s \simeq 2.1\times 10^{-9}$ and $r = 0.0038$, one obtains

$$\Omega_{\mathrm{GW},0}h^2 \sim 1.4\times 10^{-17}, \tag{346}$$

far below any interferometric reach — the primordial background is a CMB-polarization target, not an interferometer target. (A separate null *is* staked at interferometer-adjacent frequencies: the lock crossing is prompt, so no strongly supercooled TeV-scale transition background exists; a loud stochastic signal attributed to one would refute the prompt-crossing lemma.)

Equation (346) provides a powerful cross-check. The same module that fixes the CMB primordial B-mode bump at $r = 0.0038$ also predicts null interferometric primordial backgrounds. If a future stochastic background is seen in the millihertz, decihertz, or ground-based band, the first interpretation inside this framework should be *astrophysical* rather than primordial. Compact binaries, stellar collapse, and late-time structure formation can all produce stochastic backgrounds. The minimal primordial tensor sector cannot.

This distinction matters because otherwise a future gravitational-wave detection could be misread as confirming inflation in general. The correct statement for GG-A10-Cosmos is subtler:

> Detection of a stochastic gravitational-wave background does not by itself falsify the framework. Detection of a *primordial* inflationary-amplitude background does.

The CMB and interferometric channels therefore form a consistency pair. If the CMB bump appears at $r = 0.0038$ and the interferometers stay null for primordial signals, the minimal tensor module survives. If either reveals a robust primordial amplitude inconsistent with Eqs. (340) and (346), the module fails.

50.6 What the primordial-gravitational-wave section contributes to the full paper

This section does more than simply add one more observational test. It closes the loop between Parts IV, VI and VII.

From Part IV it inherits the derived plateau, one e-fold below the GUT edge with its height fixed by the declared amplitude. From Part VI it inherits the CMB rule that delensed near-term

B-mode data must reveal a primordial component at that amplitude. In Part VII it translates those inherited results into the strongest falsification handle available in the entire paper.

That is why the tensor sector belongs at the end of the present part. It is the final audit of the primordial source ledger before the paper turns, in a later part, to the chronological assembly of the full cosmic timeline.

51 Integrated summary of the black-hole and primordial-signature block

1. The black-hole information paradox is reformulated as a mismatch between global and operational descriptions (Bekenstein, 1973; Hawking, 1975; Penrose, 1994). It is dissolved at the theorem level by OGGEP and quantified at the entropy level by the GFF minimum.

2. PBHs are not needed as the main dark-matter component because the Kaxion already occupies that role and because the smooth scalar spectrum suppresses PBH formation exponentially.

3. Primordial tensors are the cleanest kill switch because the derived plateau implies $r = 0.0038$ — a background that the Simons Observatory and LiteBIRD must find at that amplitude within the launch's own spread, and in any case inside $[10^{-3}, 10^{-2}]$.

The whole block may be summarized in one table.

Table 34: Integrated summary of Part VII. The first column gives the open problem, the second the standard difficulty, the third the GG-6/GDOS interpretation, and the last the decisive empirical or conceptual test. The fourth column is the one that can fail. The first three set out the problem, the standard difficulty and the reading offered; the fourth names the test that would decide between them, and a row with a weak test in the fourth column is a weak row however strong the third looks.

Problem	**Standard difficulty**	**GG-6/GDOS interpretation**	**Decisive test**
Black-hole information paradox	thermal Hawking radiation versus global unitarity; naive factorization across the horizon	bulk is closed and unitary, boundary is open; generalized entropy is the GFF ledger; Page turnover is a GSSB saddle switch	whether a gauge-consistent operational entropy accounting can preserve global unitarity without firewalls
Primordial black holes	need a large small-scale scalar enhancement to get appreciable abundance	minimal smooth spectrum gives exponentially tiny $\beta(M)$; PBHs at most a trace subcomponent; Kaxion remains dominant dark matter	discovery of a dominant PBH dark-matter abundance would challenge the minimal scalar ledger
Primordial gravitational waves	amplitude tied directly to inflationary energy scale	derived GUT-adjacent plateau gives $r = 0.0038$ and $\Omega_{\mathrm{GW},0}h^2 \sim 10^{-16}$	a tensor measurement inconsistent with $r = 0.0038$ beyond the launch's spread 0.0037–0.0040, in either direction, is fatal

51.1 Cross-consistency and falsifiability inside the extreme-gravity block

Because Part VII contains some of the strongest claims in the paper, it is worth making the falsifiability logic explicit.

Black holes. The framework would be weakened if it could not produce a coherent operational entropy ledger that matches the modern Page-curve lesson. It would also be weakened if it required naive factorization after having built DGI into the constitution. The black-hole section therefore rises or falls with the clarity of its operational formulation.

PBHs. The PBH section is falsifiable in a simpler way. The minimal module says PBHs are not the dominant dark matter and that any substantial abundance requires an extra small-scale feature not yet derived. Discovery of a dominant PBH dark-matter population would therefore signal that either the scalar sector contains a hidden feature the present theory missed or the minimal module is false.

Primordial tensors. The tensor kill switch is the most brutal of all. A robust primordial tensor signal inconsistent with $r = 0.0038$ beyond the launch's spread immediately falsifies the derived launch (a detection well above $r = 0.0038$, or a null pushing r below $\sim 10^{-3}$, both fatal).

This asymmetry between the three tests is scientifically healthy. Not all parts of a theory fail in the same way. Some fail conceptually, some parametrically, and some observationally. Part VII contains all three kinds.

51.2 Bridge to the chronology of the universe

The correct bridge to the next major part of the paper is now clear. The preceding parts have constructed the beginning, the background expansion, the matter inventory, the growth history, and the extreme-gravity audits. What remains is to assemble them into a chronological timeline: the first microsecond, the first minutes, recombination, first light, and the route to the present epoch.

That assembly belongs to the next part. Part VII has done its job if it has shown that the extreme ends of the theory – horizons, primordial collapse, and tensor relics – do not require a separate constitution. They are governed by the same one.

The full worked-derivation notebook for the Part VII observables — the black-hole entropy ledger, the Page-curve branch switch, the primordial-black-hole tail estimate, and the tensor amplitude from the plateau height — is collected in Appendix N.11, so the body stays readable, as for Parts IV, V, and VI.

In plain words: the most extreme objects in the sky turned out to be the gentlest test of the constitution. A horizon is not a paradox factory; it is a port — a place where an open system does what open systems always do here, exchange and record. The entropy ledger closes at exactly $A/4$ because the same boundary bookkeeping that counts the cosmos counts the hole, and information is not lost because nothing in this framework was ever sealed. Even the tensor sky is load-bearing: the bump must sit exactly where the derived launch demands, and a detection well above or below it would be fatal.

The story so far has been told sector by sector — launch, matter, growth, horizons. What it still lacks is a single timeline: one ruler on which the first microsecond, the first three minutes, the first atoms, and the present afternoon all sit as milestones of one descent. That ruler is Part VIII.

52 Part VII in plain English: what is resolved here, and what is only reformulated

This Part makes a smaller claim than its subject usually invites. It does not offer a complete microscopic account of black-hole interior states, an exact graybody-factor derivation, or a full ultraviolet completion of collapse, and it says so in the criteria table before it says anything else. What it offers is a reformulation: the Page curve read as an open-system consequence of the same projection that produces quantum mechanics at the boundary, with the island prescription of the modern literature set beside it as a translation rather than a rival.

The honest position on the translation is that where the two accounts agree, the agreement is the point, and where they differ the difference is stated and not resolved. A reader who wants to know whether this framework has solved the information paradox will find the answer in the criteria table: it satisfies some of the criteria there and not others, and the ones it does not satisfy are named.

The primordial material — black holes and gravitational waves — is different in kind and is the more testable half of the Part. The mass windows are illustrative rather than predicted; they say where a population would have to sit to matter here. The gravitational-wave signature is the one place in this Part where an instrument could return a number that the framework has committed to in advance, and that is why it is carried at length despite being the least developed section.

Part VIII
Evolutionary chronology of the universe

The chronology dictionary translating between t, μ, r, a, and z; the master chronology table from the GUT boundary to the asymptotic de Sitter future; and the requested milestone ladder in one glance

53 The chronology as the GRIP descent

Part VIII reads the whole of cosmic history through one organizing coordinate and one organizing idea, shared with Parts VI and VII: the logarithmic radial coordinate $r = \ln(E_{\rm GUT}/\mu)$ *is* the GRIP-In descent coordinate of GG-A6-GRIP, and the cosmic timeline is the GRIP cascade read as a single finite descent in r. Each milestone of the history is a stage of that descent; the qualitative switches that punctuate it are GRIP-Out events; and the history terminates on a Ground Geometric Attractor.

Four landmarks fix the reading. The graceful exit from slow-smooth inflation is the first cosmological GRIP-Out event — the failure of slow roll that ends the plateau and hands the boundary to reheating; the prompt lock crossing that seeds the dark sector is the second. Recombination is the third and cleanest GRIP-Out event: the GSSB freezing of the baryon–photon plasma into the CMB geometric attractor, fixed by the integer ledger at $z_* = e^7 \approx 1097$ (Parts VI–VII). And the asymptotic de Sitter future is the Ground Geometric Attractor of the gravitational stratum, the basin into which the descent finally relaxes with $w_\Lambda = -1$. Between these switches the universe runs as GRIP-In relaxation — the long, quiet stretches of radiation, matter, and vacuum domination that occupy almost all of the proper-time budget but only a modest interval of r.

The practical payoff of this framing is the subject of the rest of the Part: read against t the history is violently nonuniform, but read against r it is a compact, evenly organized ladder — exactly as a genuinely geometric GRIP descent should be. The remaining sections build the change-of-variables dictionary, walk the milestones from the GUT edge to the present, and close with the master chronology table and the radial bookkeeping that make the descent quantitative.

54 Why a chronology block is required

$$r = \ln\!\left(\frac{E_{\rm GUT}}{\mu}\right), \tag{347}$$

where μ is the characteristic boundary energy scale of the process under discussion.

Figure 28 lays the whole of this section on one landscape strip — every milestone read simultaneously in t, T, z, and r. This is more than a change of notation. The ordinary clock t is excellent for discussing durations, causal horizons, and observational dating. The variables a and z are excellent

for the background expansion seen by the boundary observer. But the GG-6 program claims that the most economical organizing coordinate of the whole physical descent is the logarithmic coordinate r. Chronology therefore becomes a dictionary problem: every cosmic event should have a lawful image in t, a, z, μ, and r.

Part VIII is written with three purposes in mind.

First, it gives a continuous narrative from the ultraviolet boundary at E_{GUT} to the present ΛCDM-like attractor and beyond, instead of leaving the reader to assemble the story manually from the previous Parts.

Second, it sharpens the theory's falsifiability. If the chronology demands incompatible placements for inflation, baryogenesis, BBN, recombination, or the first galaxies, then the framework fails not in a specialized subproblem but in its most elementary promise: to tell one coherent history.

Third, it clarifies an important conceptual feature of the program: the universe looks very different depending on whether it is plotted against t or against r. In ordinary proper time, the early universe flashes by almost instantaneously and the late universe takes forever. In the logarithmic GG-6 coordinate the story is far more evenly distributed. The first microsecond, the first three minutes, recombination, first-star formation, and the present all sit on one relatively compact radial ladder. That is exactly the kind of compression one should expect if RG descent is genuinely geometric.

The rest of the Part is organized in the same spirit as the outline just given, but with one editorial refinement. Before walking through the first microsecond, the first three minutes, the first 3×10^5 years, and the first 10^8 years, we first write the change-of-variables dictionary that lets these landmarks be translated into the GG-6 coordinate system. Only then do we state the history itself. The Part ends with a long master summary table and several worked conversions, so that a reader can immediately locate any epoch on the same common map.

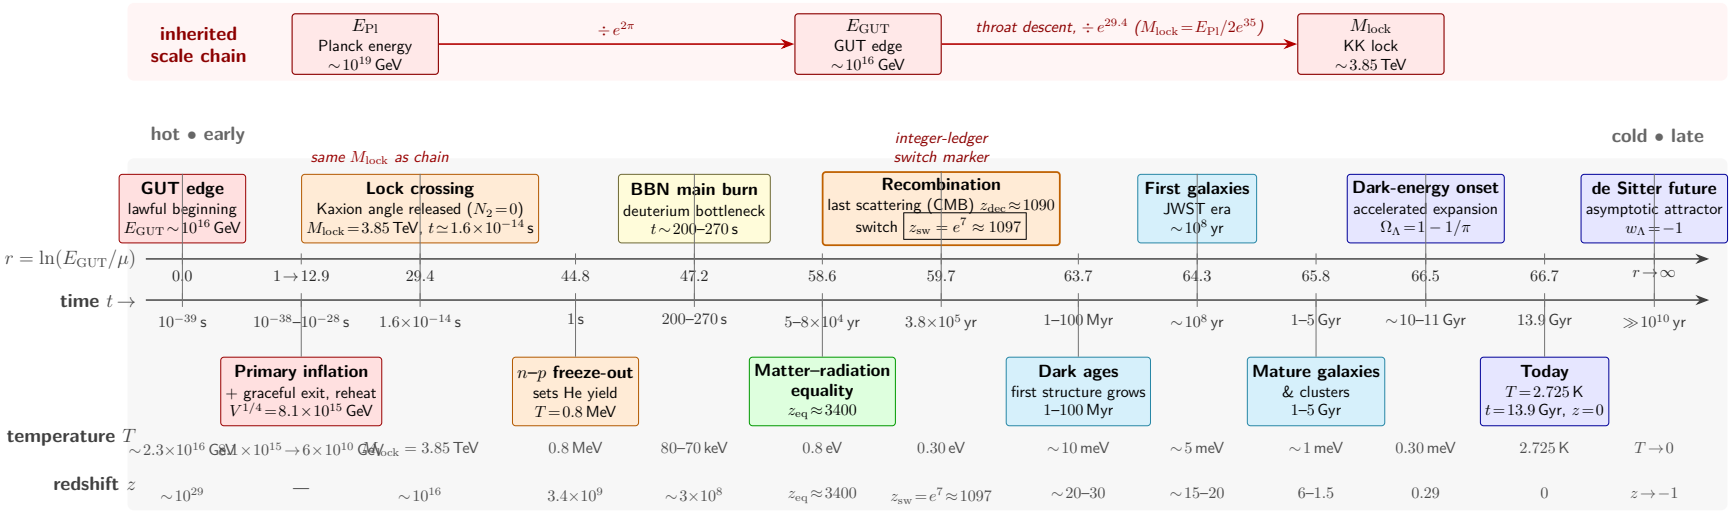

Figure 28: Cosmic-history timeline of GG-A10-Cosmos. The horizontal axis runs from the hot, early universe on the left to the cold, late universe on the right; every milestone is read off simultaneously in cosmic time t, photon temperature T, and redshift z along the three aligned axis rows. The milestones and their times/temperatures are the *standard* cosmic-history values (identical to the master chronology, Table 36): the GUT-edge lawful beginning, the primary inflation one e-fold below it with its graceful exit and Planck-strength reheating, the prompt lock crossing at $M_{\rm lock} \sim 3.85$ TeV with the release of the Kaxion's angle, neutron–proton freeze-out at $T = 0.8$ MeV, the BBN deuterium bottleneck, matter–radiation equality at $z_{\rm eq} \approx 3400$, recombination, the dark ages, the first (JWST-era) galaxies, mature galaxies and clusters, dark-energy onset with $\Omega_\Lambda = 1 - 1/\pi$, the $\Omega_m = 1/\pi$ crossing at $t = 13.8$ Gyr, today ($t = 13.9$ Gyr, $T = 2.725$ K, $z = 0$), and the asymptotic de Sitter future. The boxed corridor-switch redshift $z_{\rm sw} = e^7 \approx 1097$ is the headline *integer-ledger marker* discussed in the body (not a fit): the recombination-era GSSB switch descends from the OGN-5 integer $J = 7$, while the atomic decoupling derived from the framework's own constants lies 0.4% below it, on the measured value (§42.1). The top band shows the inherited ultraviolet scale chain $E_{\rm Pl} \rightarrow E_{\rm GUT} \rightarrow M_{\rm lock}$ with its integer gaps, the dominant one being the e^{35} lock factor $M_{\rm lock} = E_{\rm Pl}/(2e^{35})$ that fixes the middle crossing and, through the seesaw, the terminus; the dashed link ties that lock scale to the lock-crossing milestone on the timeline below. A fourth aligned row gives the radial / renormalization-group coordinate $r = \ln(E_{\rm GUT}/\mu)$ at each milestone, running from $r = 0$ at the GUT edge to $r \rightarrow \infty$ in the de Sitter future.

55 The chronology dictionary: translating between t, μ, r, a, and z

55.1 Five variables, one caveat

The paper will use five chronological variables repeatedly:

$$\{t, \mu, r, a, z\}. \tag{348}$$

Their meanings are not identical.

- t is ordinary cosmic proper time measured by comoving observers.
- μ is the characteristic boundary energy scale of the process being tracked.
- $r = \ln(E_{\rm GUT}/\mu)$ is the logarithmic radial coordinate inherited from GG-6.
- $a(t)$ is the scale factor, normalized to $a_0 = 1$ today.
- z is the corresponding redshift, $1 + z = a^{-1}$.

The only important caveat is the status of μ. In a tightly coupled relativistic plasma, μ is naturally identified with the temperature T. During reheating that same role is played by the reheating temperature. During the inflationary plateau, where a thermal plasma is not the correct description, μ should instead be read as the characteristic vacuum scale $V^{1/4}$. During the late matter and dark-energy eras, one may again use the photon temperature T_γ as a convenient operational clock even though it is no longer the dominant energy density. Thus μ is the characteristic *process* scale, not necessarily the energy density of the dominant cosmic component.

With that caveat stated, the entire chronology can still be presented on one shared map.

55.2 Exact algebraic relations

The first two relations are exact:

$$1 + z = a^{-1}, \tag{349}$$

$$r = \ln\!\left(\frac{E_{\rm GUT}}{\mu}\right). \tag{350}$$

The second equation implies the useful differential identity

$$\mathrm{d}r = -\mathrm{d}\ln\mu. \tag{351}$$

Therefore equal steps in r correspond to equal multiplicative drops in the characteristic energy scale. That is why the r-coordinate is much better than t for displaying the hierarchy structure of the cosmic history.

The next relation is inherited directly from Part III. During adiabatic expansion of a coupled relativistic plasma,

$$T(a) \propto g_{*s}^{-1/3}(T)\, a^{-1}, \tag{352}$$

so that, after normalizing to the present photon temperature $T_{\gamma,0}$,

$$T_\gamma(z) = T_{\gamma,0}(1+z)\left(\frac{g_{*s,0}}{g_{*s}(T)}\right)^{1/3}. \tag{353}$$

Substituting Eq. (353) into Eq. (350) gives the practical redshift–radial dictionary for epochs that are well described by a photon temperature:

$$r(z) = \ln\left(\frac{E_{\rm GUT}}{T_{\gamma,0}}\right) - \ln(1+z) + \frac{1}{3}\ln\left(\frac{g_{*s}(T)}{g_{*s,0}}\right). \tag{354}$$

If the change in g_{*s} can be ignored across the interval of interest, this reduces to the extremely simple law

$$r(z) \simeq r_0 - \ln(1+z), \qquad r_0 := \ln\left(\frac{E_{\rm GUT}}{T_{\gamma,0}}\right). \tag{355}$$

Using the inherited numerical value $E_{\rm GUT} \simeq 2.28\times10^{16}$ GeV together with $T_{\gamma,0} \simeq 2.35\times10^{-4}$ eV, one finds

$$r_0 \simeq 66.7. \tag{356}$$

This number is worth pausing over. It means that the entire observable history of the boundary universe, from the GUT edge to the current photon bath, spans only about 67 natural logarithmic units in energy. That is a huge ratio in ordinary linear units, but it is a compact interval in the native GG-6 coordinate.

55.3 Radiation-era time-temperature law

For early epochs the most useful relation is again inherited from Part III:

$$t \simeq 0.301\, g_*^{-1/2}\frac{M_{\rm Pl}}{T^2}. \tag{357}$$

In practical units this is often written as

$$t \simeq 2.42\, g_*^{-1/2}\left(\frac{T}{\rm MeV}\right)^{-2}\ {\rm s}. \tag{358}$$

Solving for T gives

$$T \simeq \sqrt{0.301}\, g_*^{-1/4}\sqrt{\frac{M_{\rm Pl}}{t}}. \tag{359}$$

Every early-universe clock in this Part is read from Eq. (358) with $g_* = 106.75$ above the electroweak seam — so the lock, 3.85 TeV, sits at $t \simeq 1.6\times10^{-14}$ s — and with $g_* = 3.36$ after electron-positron annihilation, so that 80 keV is $t \simeq 200$ s and 70 keV is $t \simeq 270$ s. Combining this with Eq. (350) yields the radiation-era radial clock

$$r_{\rm rad}(t) = \ln\left[\frac{E_{\rm GUT}}{\sqrt{0.301}\, g_*^{-1/4}\sqrt{M_{\rm Pl}/t}}\right]. \tag{360}$$

The virtue of Eq. (360) is conceptual as much as numerical. It shows that in radiation domination the logarithmic coordinate grows only as $\frac{1}{2}\ln t$. A gigantic multiplicative stretch in ordinary time corresponds to only a modest additive displacement in r.

That is why events that look unimaginably far apart in seconds or years can be only a few radial units apart on the GG-6 map.

55.4 Matter-era and vacuum-era clocks

Once the universe is matter dominated, the simple radiation-era temperature law is no longer the best way to read the background. In the ordinary matter era,

$$H(z) \simeq H_0\sqrt{\Omega_m}\,(1+z)^{3/2}, \tag{361}$$

so integrating $\mathrm{d}t = -\mathrm{d}z/[(1+z)H(z)]$ gives the standard approximation

$$t(z) \simeq \frac{2}{3H_0\sqrt{\Omega_m}}\,(1+z)^{-3/2}, \qquad z_{\mathrm{eq}} \gg z \gg \mathcal{O}(1). \tag{362}$$

This relation is not exact all the way to the present because dark energy later becomes important, but it is perfectly adequate for the interval from recombination through the first stars and much of galaxy assembly.

At late times, once the residual vacuum term dominates, the background approaches

$$a(t) \propto e^{H_\Lambda t}, \qquad H_\Lambda \simeq \text{const.} \tag{363}$$

In this era the photon bath continues to cool as $T_\gamma \propto a^{-1}$, so the radial coordinate still increases according to Eq. (355). The dominant vacuum component does not cool, but the operational thermometer used by observers does.

55.5 A compact dictionary table

Before moving to the actual history, it is helpful to summarize the change-of-variables map in one place.

Table 35: Coordinate dictionary used throughout Part VIII. The point is not to replace standard cosmological variables, but to display how the GG-6 logarithmic coordinate r can be read alongside the ordinary chronology. The columns give the variable, its definition, and its meaning here. The dictionary does not replace the standard cosmological variables; it sets the radial coordinate beside them so that a reader can move between the two descriptions without translating in their head.

Variable	Definition	Meaning in Part VIII
t	cosmic proper time	The ordinary chronological variable used by astrophysics and thermal history.
μ	characteristic process scale	Usually the plasma temperature, but during inflation the characteristic vacuum scale, and during the late universe often the photon temperature used as an operational thermometer.
r	$r = \ln(E_{\rm GUT}/\mu)$	The logarithmic GG-6 radial coordinate; equal steps correspond to equal multiplicative drops in energy.
a	scale factor with $a_0 = 1$	Global expansion variable.
z	$1 + z = a^{-1}$	Observationally friendly version of the scale factor.
Radiation law	$t \simeq 0.301\, g_*^{-1/2} M_{\rm Pl}/T^2$	Main early-universe conversion between t and $\mu \sim T$.
Redshift law	$r(z) \simeq r_0 - \ln(1+z)$ for slowly varying g_{*s}	Practical late-time conversion between redshift and the radial coordinate.

The rest of the chronology now follows directly from walking down this dictionary.

56 The first microsecond: from the GUT boundary to the electroweak seam

56.1 Why the first microsecond already includes several logically distinct stages

In a standard textbook chronology the phrase "the first microsecond" usually refers to the time window during which the thermal bath cools from the electroweak or multi-GeV range toward the QCD crossover. In the present framework that language is still useful, but it must be interpreted more carefully because the true beginning of the boundary cosmology has already been placed at $E_{\rm GUT}$, while the inflationary plateau sits one e-fold below it and the lock scale $M_{\rm lock} \sim 3.85$ TeV is crossed on the way down.

This means that the first microsecond of the universe is not one homogeneous physical stage. It contains at least four logically distinct chapters:

1. the ultraviolet initialization of the boundary theory at $\mu = E_{\rm GUT}$, $r = 0$;
2. the primary slow-smooth plateau one e-fold below the GUT edge, and its graceful exit;

3. Planck-strength reheating to 6×10^{10} GeV and the radiation descent to the promptly crossed lock;

4. resonant leptogenesis at the TeV scale, sphaleron conversion, and the electroweak-to-GeV thermal descent.

In linear time these stages are compressed into a fantastically small interval. In r-space, however, the structure is easier to read. The ultraviolet start sits at $r = 0$. The lock seam sits at

$$r_{\rm lock} = \ln\!\left(\frac{E_{\rm GUT}}{M_{\rm lock}}\right) \simeq 29.4. \tag{364}$$

The one-microsecond milestone, by contrast, sits near the GeV scale and therefore near

$$r_{\mu\rm s} \sim \ln\!\left(\frac{2.28 \times 10^{16}\,\text{GeV}}{1\,\text{GeV}}\right) \simeq 37.7. \tag{365}$$

Thus everything from the lock crossing to the one-microsecond mark occupies only about eight radial units:

$$\Delta r_{\text{lock}\to\mu s} \simeq 8.3. \tag{366}$$

This is the first place where the GG-6 map becomes genuinely useful: it makes the early chronology visually compact.

56.2 Ultraviolet initialization and the descent to the plateau

The first statement of the chronology is not a guess about a pre-geometric singularity. It is the already-inherited Part III statement that the boundary description starts at the GUT surface, not at $M_{\rm Pl}$. The chronology therefore begins on a finite ultraviolet seam carrying finite initial data. The role of this starting point is constitutional: it tells us where the effective 4D cosmological language is first allowed to be used.

From there the system descends in r. At this stage it is misleading to picture the universe as an ordinary thermal plasma with a single temperature and an observationally meaningful redshift. The better picture is that the six-dimensional throat is selecting the low-energy seam on which the later 4D history will be exported. In other words, the relevant evolution is geometric before it is thermal.

This clarifies where the inflationary stage sits in the descent. The boundary description turns on at $E_{\rm GUT}$, and the primary inflationary plateau selected by the GDOS boundary law lies just below it — one e-fold under the GUT edge, at $V^{1/4} = 8.1 \times 10^{15}$ GeV. Inflation is therefore a high-scale, GUT-adjacent event, not a deep-infrared one; the lock scale $M_{\rm lock}$ enters the later history (the electroweak seam and the dark sector), not the inflationary plateau.

56.3 Slow-smooth inflation on the GUT-adjacent plateau

The inflationary stage occupies the GDOS-forced Starobinsky/ξ plateau one e-fold below the GUT edge, at $V^{1/4} = 8.1 \times 10^{15}$ GeV. Evaluated at the reheat-matched e-fold count $N^* = 52.6$, the plateau fixes the tilt $n_s = 0.9634$ exact ($1 - 2/N^* = 0.9620$ at leading order) and the tensor ratio

$r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order) — small but nonzero, a factor of nine below the present bound and a two-sided target (Part IV). The Hubble scale of inflation is not identical to the ultraviolet boundary scale $E_{\rm GUT}$, but it sits just below it, so the inflationary event is high-scale.

The intuitive chronology is therefore the following. The boundary description turns on at the GUT surface, and the modulus is captured on the plateau immediately below it. There it executes a quasi-de Sitter slow descent, stretching the comoving universe, erasing curvature and horizon problems, and generating the scalar perturbations later fossilized in the CMB. The duration in ordinary time can be tiny while the effect on $a(t)$ is enormous, because the defining quantity is the number of e-folds, not the absolute size of t.

An earlier reading of this paper placed the inflationary plateau at the lock scale $M_{\rm lock}$, predicting a nearly vanishing tensor ratio; that reading is retired (it gave a blue tilt and is refuted in Part IV), and is replaced by the derived GUT-adjacent plateau above, whose tensor prediction $r = 0.0038$ is a sharp, falsifiable target rather than an unobservable null.

56.4 Graceful exit, reheating, and the two crossings that stock the sectors

Exit is slow-roll failure at the plateau's edge (Part IV); the lock GSSB is the later, separate crossing. The important chronological point is that several sectors are born along one descent.

First, the stored vacuum energy is transferred to a hot thermal bath with the derived reheating temperature $T_{\rm rh} \simeq 6 \times 10^{10}$ GeV (fourteen e-folds of modulus decay; Appendix G), the lock scale $M_{\rm lock}$ then crossed promptly as the bath cools. Second, the compact-angle sector is released misaligned at the lock crossing, giving the Kaxion dark matter its initial condition (single-patch by conjecture, Part V). Third, the residual vacuum remnant is left behind as the seed of the later dark-energy ledger. In ordinary model building one often treats these as unrelated sectors: one field for inflation, another for dark matter, another constant for dark energy. In the present chronology they are three descendants of one descent, read at two of its crossings.

Once reheating has completed, the language of temperature, reaction rates, and equilibrium becomes operationally appropriate. The universe is now a radiation bath with

$$T_{\rm rh} \simeq 6 \times 10^{10}\ {\rm GeV}, \qquad r_{\rm rh} \simeq 12.9, \qquad 1 + z_{\rm rh} \sim \frac{T_{\rm rh}}{T_{\gamma,0}} \sim 10^{23}. \tag{367}$$

This is the cleanest chronological landing point for the first genuine thermal chapter of the universe.

56.5 TeV leptogenesis and the electroweak corridor

The next stage is the most important matter-origin step of the entire early chronology. The same bath that the modulus decay produced contains the heavy quasi-degenerate Majorana states responsible for the inherited resonant leptogenesis corridor. Because the lock scale is already fixed, the baryogenesis story does not ask us to insert a separate ultra-heavy seesaw sector. It plays out where the thermal bath already is.

Chronologically the logic is straightforward. Heavy neutrinos with masses of order $M_{\rm lock}$ fall out of equilibrium and decay while electroweak sphalerons are still active. The holonomy-fixed CP asymmetry and the washout factor then feed the final baryon asymmetry. The result is that by

the time the universe cools through the electroweak seam, the sign and magnitude of the future baryonic inventory are already effectively fixed.

This matters for the later history because every later stage—BBN, recombination, first stars, galaxies, and even the late-time Ω_b share—assumes the existence of an already-fixed small baryon asymmetry. Part VIII makes the dependence explicit: ordinary baryons are not an arbitrary late-time abundance parameter. They are the surviving residue of a process that happened near the very beginning of the thermal descent.

The electroweak seam itself sits around $\mu \sim 100$ GeV, at a cosmic age of order 10^{-11} seconds and at

$$r_{\rm EW} = \ln\left(\frac{E_{\rm GUT}}{100\,\text{GeV}}\right) \simeq 33.1. \tag{368}$$

By the one-microsecond mark the universe has cooled to the GeV corridor, so the whole electroweak and leptogenesis story is already over. This is why the first microsecond is a natural milestone. It marks the end of the matter-origin prelude and the handoff to the hadronic and nuclear chapters.

57 The first three minutes: nucleosynthesis and the fixation of ordinary matter

57.1 From the GeV corridor to the MeV corridor

Once the electroweak chapter is closed, the universe continues as an ordinary radiation bath. In linear time this stage is still extremely brief, but in physics content it is rich. The thermal bath must pass through the QCD crossover, neutrino decoupling, neutron-proton freeze-out, electron-positron annihilation, and the deuterium bottleneck before the light-element abundances are fixed.

The one-microsecond milestone sits near $\mu \sim 1$ GeV and $r \sim 37.7$. The QCD crossover follows soon after, around 150 to 200 MeV, which corresponds to

$$r_{\rm QCD} \sim \ln\left(\frac{2.28 \times 10^{16}\,\text{GeV}}{0.15\,\text{GeV}}\right) \simeq 39.6. \tag{369}$$

Already one sees the convenience of the radial map. From the GeV corridor to hadronization is less than two radial units. Yet in ordinary language that step changes the natural degrees of freedom from quarks and gluons to protons, neutrons, mesons, and leptons.

57.2 Hadronization and the operational meaning of $U(1)_B$

Part VIII is a good place to make a conceptual point that is easy to miss in purely particle-theory presentations. The exact $U(1)_B$ symmetry of the bulk becomes most operationally meaningful once hadrons exist. Before the QCD transition, baryon number is carried by a quark plasma. After hadronization, it is carried by identifiable long-lived baryonic objects. This is the stage at which the protected nature of the bulk baryon charge begins to look like the ordinary stability of protons and neutrons in an expanding thermal background.

The chronology therefore has a hierarchy of carriers:

holonomy and anomaly data → lepton asymmetry → baryon asymmetry → hadronic baryon carriers.

That sequence is conceptually important because it prevents a common confusion. The exactness of $U(1)_B$ does not mean that the low-energy baryonic content was always transparently visible. It means that once the bath cools enough for hadronic carriers to be the right variables, they inherit a charge that was protected all along.

57.3 Neutrino decoupling, freeze-out, and the first-second milestone

Near $T \sim 1$ MeV the weak interaction rates finally lose the race against the Hubble expansion. This single statement controls several later facts at once. Neutrinos decouple from the electromagnetic plasma. The neutron-to-proton ratio ceases to track its equilibrium value exactly. And the future helium abundance becomes calculable from the frozen n/p ratio plus neutron decay during the remaining delay before efficient nuclear burning begins.

In the coordinate dictionary this is the moment

$$\mu \sim 1\,\mathrm{MeV}, \qquad t \sim 1\,\mathrm{s}, \qquad r \sim 44.6, \qquad 1+z \sim 4\times 10^9. \tag{370}$$

Several features of the history are worth emphasizing.

First, the universe is still overwhelmingly radiation dominated.

Second, the baryon asymmetry generated near the TeV corridor is now simply a background parameter for a nuclear-reaction network. In other words, the chronology has changed language: the matter-origin problem is over, and the nuclear-abundance problem has begun.

Third, the difference in r between the electroweak seam and the one-second milestone is only about 11.5. This is precisely why the early universe is so efficiently summarized by a logarithmic coordinate: an enormous range of proper time and reaction content is compressed into a modest interval of r.

57.4 The deuterium bottleneck and the ignition of BBN

A universe can contain plenty of protons and neutrons without immediately synthesizing helium. The reason is the deuterium bottleneck: at sufficiently high temperature any newly formed deuteron is quickly photodissociated by the energetic photon tail. Only once the temperature has dropped to the ~ 0.1 MeV corridor does deuterium survive long enough to act as a stepping stone to helium and trace heavier light elements.

This is the technical reason the light-element story is concentrated not at one second but at tens to hundreds of seconds. The thermal bath needs time to cool into the narrow nuclear corridor where the reactions can run faster than photodissociation while still faster than the expansion. Once this happens, the universe burns quickly.

At the opening of the bottleneck,

$$T \sim 80\,\mathrm{keV}, \qquad t \simeq 200\,\mathrm{s}, \qquad r \sim 47.1. \tag{371}$$

By the end of the main burn, $T \simeq 70$ keV at $t \simeq 270$ s — the "three minutes" of the textbook name, stretched to four and a half by the post-annihilation $g_* = 3.36$ — the main helium synthesis is essentially done. This is why the checkpoint of "the first three minutes" is a natural one. It is the interval over which ordinary visible matter, as a chemical inventory, becomes recognizable.

57.5 The three-minute milestone

Setting $t \simeq 270$ s in the radiation-era relation ($g_* = 3.36$) gives a characteristic temperature near 70 keV, the end of the main burn, corresponding to

$$r_{3\mathrm{min}} = \ln\left(\frac{E_{\mathrm{GUT}}}{70\,\mathrm{keV}}\right) \simeq 47.2. \tag{372}$$

This single number is worth remembering because it anchors the whole visible-baryon chronology. By $r \sim 47$,

- the baryon asymmetry is already inherited from the TeV corridor,
- hadronic carriers already exist,
- weak freeze-out is already over,
- the deuterium bottleneck has opened,
- and the helium yield has nearly frozen in.

In ordinary time four and a half minutes have passed. In radial distance the universe has already traveled almost 18 logarithmic units beyond the lock seam, and 34 beyond reheating.

57.6 The lithium window inside the chronology

The standard BBN calculation is famously successful for ^{4}He and deuterium but overpredicts ^{7}Li. Part V introduced the GG-6/GDOS repair as an open-system port term acting only on the A=7 subspace. In Part VIII the role of that module is purely chronological: it tells us exactly where the repair sits and why it does not disturb earlier or later stages.

The relevant window is the $t \simeq 200$ to 1500 s corridor, 80 down to 30 keV on the clock of Eq. (358). This is late enough that the gross helium yield is already insensitive to the detailed A=7 adjustment, and early enough that the A=7 network can still be influenced through the proposed log-enhanced port. In the chronology this means the lithium repair is not a loose astrophysical afterthought applied gigayears later in stars. It is a sharply localized deformation of the original nuclear freeze-out history.

That localization is one of the theory's cleanest virtues. The BBN correction is not allowed to be broad. It must sit in a narrow time-energy corridor, preserve D/H and Y_p, and then hand off a fixed baryonic inventory to all later eras. The chronology therefore turns a phenomenological requirement into a consistency requirement.

57.7 What is fixed by the end of the first three minutes

By the end of the first three minutes, three things are fixed for the rest of the paper.

First, the visible baryonic composition is fixed to high accuracy.

Second, the baryon number density entering recombination, structure growth, and the later π-locked inventory is fixed.

Third, the dark sector has already been present since the lock crossing, so the universe now carries both the visible matter inventory and the compact-angle dark matter inventory into the long radiation-dominated interval that precedes equality.

This is exactly why the early chronology cannot be split into disconnected particle and cosmology narratives. By the time one reaches the end of the first few minutes, the particle ledger and the cosmic ledger are already inseparable.

58 The first 3×10^5 years: equality, recombination, and the release of the CMB

58.1 The long radiation-dominated bridge

After the BBN window closes, the universe enters a long interval that is much less dramatic chemically but enormously important dynamically. Radiation still dominates the expansion, the photon bath keeps cooling as a^{-1}, the baryon density continues diluting as a^{-3}, and subhorizon perturbations evolve under the competition between radiation pressure and gravitational attraction.

In ordinary time, the bridge from a few minutes to tens of thousands of years is enormous. In radial distance it is more modest. The universe slides from $r \sim 47$ at the end of BBN toward $r \sim 58$ by equality. This illustrates again the usefulness of the r-coordinate: the early thermal history is a ladder of only moderate logarithmic length.

58.2 Matter-radiation equality and why it is a pivotal checkpoint

Equality is a decisive milestone because it changes which component controls the Hubble rate and therefore which perturbation modes can grow efficiently. In the baseline ΛCDM story equality occurs near $z \sim 3400$. In the GG-6 chronology that statement is retained without an overlay: Part VI once placed a brief early-dark-energy lobe and a small dark-radiation excess here, and has withdrawn both — the residue is forbidden by the module's own kill criterion and the radiation term is a derived null — so the sound horizon and the equality-adjacent transfer history are the standard ones.

Using the ordinary photon temperature as the process scale, equality sits near

$$\mu_{\rm eq} \sim 0.8\,{\rm eV}, \qquad r_{\rm eq} \simeq 58.6, \qquad t_{\rm eq} \sim 5 \times 10^4 \text{ to } 8 \times 10^4\,{\rm yr}. \tag{373}$$

Chronologically equality matters for two reasons.

First, it is the pivot that separates a radiation-shaped perturbation history from a matter-shaped growth history.

Second, it is where an early-dark-energy lobe would have to sit to move the sound horizon — which is why its absence is a test: the framework carries none, and a lobe detected near $z \sim 3500$ at $f_{\rm EDE} \gtrsim 0.02$ counts against it (Part VI).

58.3 Recombination as a GSSB event

The next great checkpoint is recombination. In standard cosmology one says that the Thomson scattering rate falls below the Hubble rate, neutral hydrogen forms, and photons decouple. All of that remains true here. But GG-6/GDOS sharpens the interpretation by reading the event as a GSSB transition followed by a geometric attractor.

Before recombination the baryon-photon plasma is a tightly coupled fluid. Many microscopic scattering histories are operationally equivalent at the level of macroscopic perturbation variables. During the transition the free-electron fraction plummets, the gauge orbit of available scattering histories collapses, and one realized phase pattern of the acoustic oscillations is frozen in. After the transition the photons free stream and the pattern becomes the primary CMB sky.

This reinterpretation changes neither the observed acoustic peaks nor the ordinary decoupling physics. What it changes is the status of the event in the theory. Recombination is not just a passive loss of opacity. It is one of the cleanest cosmological examples of DGI, GSSB, and GA in one place.

58.4 The recombination milestone

Numerically recombination lies near

$$\mu_{\rm rec} \sim 0.26\,{\rm eV}, \qquad r_{\rm rec} \simeq 59.7, \qquad t_{\rm rec} \simeq 3.8 \times 10^5\,{\rm yr}, \qquad z_* = e^7 \approx 1097. \tag{374}$$

Here the epoch is quoted at the integer-ledger corridor-switch value $z_* = e^7 \approx 1097$ of Parts VI–VII — the topological landmark, not the measured decoupling; the rounded standard empirical value is $z_{\rm rec} \simeq 1100$. The radial separation between equality and recombination is only about one unit:

$$\Delta r_{\rm eq\to rec} \simeq \ln\!\left(\frac{0.8\,{\rm eV}}{0.26\,{\rm eV}}\right) \approx 1.1. \tag{375}$$

That single line is a beautiful example of why the logarithmic chronology is so powerful. Equality and recombination are separated by tens to hundreds of thousands of years in ordinary time, but in energy terms they are neighboring steps on the same radial ladder.

58.5 The CMB as fossil record and chronological anchor

Once released, the CMB becomes the cleanest fossil record of the earlier chronology. It remembers the scalar tilt generated during inflation, the baryon loading fixed by the visible matter inventory, the damping tail shaped by diffusion, the peak spacing determined by the sound horizon, and any small equality-adjacent shifts produced by EDE-like contributions.

This means that Part VIII is not merely summarizing older results. It is also showing how the paper's chronology is constrained. One cannot place the inflationary stage, the baryon fraction, or the equality shift arbitrarily and still keep the CMB. In that precise sense the CMB is the strongest

chronological anchor of the whole paper.

59 The first 10^8 years: dark ages, first stars, and early galaxy formation

59.1 Dark ages and residual baryon-dark-matter streaming

After recombination the universe becomes neutral and transparent, but it is not yet luminous. This is the era usually called the dark ages. The background is matter dominated, the photon bath continues to cool passively, and structure grows in the dark matter potentials while baryons begin to fall in after being released from photon drag.

In the minimal 4D ΛCDM picture, the baryons inherit a residual streaming velocity relative to cold dark matter. That relative motion suppresses gas collapse in the shallowest minihalos and delays first-star formation. An earlier reading of Part VI modified the story here — a small drag from the gauged $U(1)_B$ sector damping the streaming velocity — and that reading is withdrawn: the sector has no vector below the lock scale, so the wind is the ΛCDM wind (§44.4).

Chronologically this remains the place where a small environmental change becomes a large astrophysical timing effect: because the first luminous objects form on the exponentially sensitive high-mass tail of the halo distribution, the one surviving lever — a lowered effective collapse threshold in the baryonic environment (Part VI) — can shift the onset of collapse visibly in the timing of first stars and early galaxies.

59.2 The one-hundred-million-year milestone

The 10^8-year milestone is therefore not arbitrary. It lies exactly where the dark ages are ending and the first luminous objects are expected to appear. Taking $z \sim 20$ as a representative value gives

$$a \sim 0.048, \qquad \mu \sim T_\gamma \sim 5\,\mathrm{meV}, \qquad r \sim 63.7. \tag{376}$$

The crucial comparison is with recombination:

$$\Delta r_{\mathrm{rec}\to 100\,Myr} = \ln\!\left(\frac{0.26\,\mathrm{eV}}{5\,\mathrm{meV}}\right) \approx 4.0. \tag{377}$$

Thus the entire march from photon decoupling to the edge of first-star formation occupies only about four radial units. In ordinary time the interval is close to a hundred million years. In logarithmic energy it is a very short walk.

59.3 Population III stars and the end of the dark ages

The first stars are usually placed in the corridor $z \sim 20$ to 15, corresponding to approximately 0.1 to 0.2 Gyr. In the GG-6 reading their formation is encouraged by three previously inherited facts:

(i) the compact-angle dark matter is already present as a cold background,

(ii) the baryon-dark-matter streaming wind was once claimed to be modestly damped — *withdrawn*, no mediator (§44.4),

(iii) and the small-scale transfer function was said to have received a mild early boost from the equality-adjacent EDE lobe — an inheritance withdrawn (§44.3), where the lobe's effect on this channel is computed and is negative;

None of these changes by itself is dramatic. Together they move the onset of first light to slightly higher redshift and allow more massive halos to participate earlier than in the most conservative 4D baseline.

This is exactly the type of chronology change one wants from a unification program: a small number of tightly connected upstream modifications, not a long menu of disconnected late-time fixes.

59.4 The first galaxies and the JWST corridor

By $z \sim 15$ to 9, roughly 0.2 to 0.5 Gyr after the hot beginning, the first long-lived galaxies and heavy black-hole seeds appear. This is the corridor in which JWST has already provided the most dramatic observational pressure on conservative ΛCDM structure forecasts.

In the chronology of this paper, the explanation is not that galaxy formation is a new problem requiring a new sector invented specifically for JWST. Rather, the explanation is that the same cold compact-angle dark matter and the lowered baryonic collapse threshold shift the formation timeline coherently. (An early transfer-function modification and a baryonic drag damping once stood in that list; the first's derivation is withdrawn and its computed effect on this channel is negative, §44.3; the second has no mediator, §44.4.) The first galaxies therefore appear neither miraculously nor by decree, but as the integrated consequence of the earlier history.

This coherence is crucial. If a model solved the JWST timing problem only by wrecking the CMB, BBN, or the Hubble audit, the chronology would fail. Part VIII is precisely where one sees that the framework is at least attempting something sharper: the same early modifications are being asked to survive every later checkpoint.

59.5 Reionization as the first large-scale feedback era

Once stars and accreting black holes exist, the chronology enters a new regime. The universe is no longer merely expanding and cooling. It is being actively reprocessed by luminous structure. Ultraviolet photons reionize the intergalactic medium, supernovae inject metals, and black holes seed energetic feedback. This is the first cosmic era in which the chronology of matter becomes inseparable from astrophysical ecology.

Part VIII does not try to solve every detail of reionization. It records the structural point. Reionization is the first major feedback era built on top of the earlier attractor sequence. Its onset depends on when the first stars and galaxies form, and those times depend on the full earlier ledger.

60 From the first gigayear to the present attractor

60.1 From young galaxies to the mature cosmic web

After the first gigayear the universe enters the long middle age in which most familiar structure grows. Halos merge, galaxies become chemically enriched, clusters assemble, and the large-scale cosmic web visible to present astronomy is built. The universe is still matter dominated for most of this interval, so the background continues to decelerate while structure grows efficiently.

The chronology of this stage is usually written entirely in terms of z, because the observational data are galaxy surveys, weak lensing, BAO, and cluster counts. In the GG-6 map the same era is also a slow drift from $r \sim 63$ toward $r \sim 67$. The smallness of that radial span is striking. Gigayears of ordinary astrophysics happen within only a few logarithmic energy units of the present CMB temperature.

60.2 Why the late universe is narrow in r

This is one of the cleanest global insights yielded by the chronology block. The late universe occupies most of the proper-time budget of cosmic history but only a very small fraction of the radial energy budget. The consequence is subtle but important. A small mistake in the late-time background or growth law can consume a large fraction of the remaining radial corridor and therefore show up clearly in precision observables such as H_0, S_8, BAO, or the age-redshift relation.

Put differently: the late universe is dynamically slow but geometrically tight. That is why precision tensions are such powerful probes of the program. There is not much room left in r-space for arbitrary late-time modifications if one wants to keep the rest of the chronology coherent.

60.3 Peak star formation and mature structure

The peak cosmic star-formation epoch near $z \sim 2$ sits around $t \sim 3$ Gyr and $a \sim 0.33$. By then the universe has become an efficient machine for converting baryons into stars and black-hole accretion, but matter still dominates the expansion. Later cluster growth, large-scale lensing, and galaxy peculiar velocities all inherit the same matter content and same earlier transfer history.

This is the era in which the S_8 problem becomes observationally concrete. The Part VI claim was that the same EDE and dark-radiation ingredients that help the Hubble audit also suppress structure growth mildly, while a small geometric lensing enhancement partly compensates. Part VIII records the chronological location of that claim: it is a late-universe consistency statement about the mature web, not an arbitrary insertion at the beginning.

60.4 Matter-dark-energy equality and the end of deceleration

Eventually the vacuum term catches up. In the π-locked inventory language,

$$\Omega_m = \frac{1}{\pi}, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi}. \tag{378}$$

Matter-vacuum equality occurs when

$$\Omega_\Lambda = \Omega_m(1+z)^3, \tag{379}$$

so that

$$1 + z_{\Lambda m} = \left(\frac{\Omega_\Lambda}{\Omega_m}\right)^{1/3} = (\pi - 1)^{1/3}, \qquad z_{\Lambda m} \approx 0.288. \tag{380}$$

The present paper already emphasized in Part IV that this is the cleanest geometric answer to the cosmic coincidence problem. In the chronology language it means something very simple: the recent transition to vacuum domination is not a lucky accident inserted by hand. It is where the already-fixed geometric shares force the late-time expansion to pivot.

Chronologically this transition occurs at a cosmic age of order 10 to 11 Gyr. From then onward the expansion accelerates and the background begins its drift toward the late-time vacuum attractor.

60.5 The present epoch

Today sits at

$$t_0 \simeq 13.9\,\mathrm{Gyr}, \qquad a_0 = 1, \qquad z_0 = 0, \qquad T_{\gamma,0} \simeq 2.35 \times 10^{-4}\,\mathrm{eV}, \qquad r_0 \simeq 66.7. \tag{381}$$

Two observations matter.

First, the present epoch is only about seven radial units below recombination:

$$\Delta r_{\mathrm{rec}\to 0} = \ln\!\left(\frac{0.26\,\mathrm{eV}}{2.35 \times 10^{-4}\,\mathrm{eV}}\right) \approx 7.0. \tag{382}$$

Second, the present epoch is only about three radial units below the 100 Myr milestone:

$$\Delta r_{100\,\mathrm{Myr}\to 0} \approx \ln\!\left(\frac{5\,\mathrm{meV}}{2.35 \times 10^{-4}\,\mathrm{eV}}\right) \approx 3.1. \tag{383}$$

These are extraordinarily compact distances. They tell us that most of late cosmic history unfolds in a narrow logarithmic corridor, which is exactly why the late-time parameter audits are so unforgiving.

61 Beyond the present: the asymptotic future

$$w_\Lambda = -1. \tag{384}$$

Therefore the late universe approaches an asymptotically de Sitter-like attractor rather than a big rip or a future recollapse. Matter and radiation continue to dilute, the photon bath cools exponentially with the scale factor, and the redshift variable for any future observer tends formally to

$$z \to -1. \tag{385}$$

From the point of view of the radial coordinate, the future is simple:

$$\mu \to 0, \qquad r \to +\infty. \tag{386}$$

The chronology is thus open-ended in r, even though the observable patch of any one worldline tends toward a finite event-horizon geometry.

This does not make the far future irrelevant to the present paper. The black-hole and information results of Part VII live precisely on this late-time background. Any complete chronology must therefore end on the same type of asymptotic attractor that Part VII assumes when discussing evaporation and horizon bookkeeping. Part VIII closes that logical loop.

62 Master chronology table

The chronology assembled in the preceding sections is collected here in two complementary forms. Figure 29 draws the descent to scale on the radial coordinate $r = \ln(E_{\rm GUT}/\mu)$, reading each key epoch simultaneously in energy, time, and the scale factor, with the late universe magnified below a break so that recombination, structure, and the present epoch remain legible. Table 36 then lists every milestone numerically — the same (t, μ, r, a, z) values — together with its standard cosmological role and its GG-6/GDOS reading. The two are redundant by design: the figure shows the *shape* of the one open-system descent at a glance, and the table gives the numbers behind it, so that a reader can move between the picture and its ledger without leaving the page-spread.

Figure 29: The cosmic GRIP descent, drawn to scale on $r = \ln(E_{\rm GUT}/\mu)$. One open-system descent from the GUT edge ($r = 0$) to the de Sitter future, read simultaneously in four quantities — energy μ, time t, and the scale factor a, each against the radial depth r. The key epochs of the master chronology below are drawn to scale and colored, with the late universe magnified beneath the break so that recombination, structure, and the present epoch remain legible. Highlighted rows carry the framework's derived markers: the plateau tilt and tensor ratio ($n_s = 0.9634$, $r = 0.0038$), the lock $M_{\rm lock} = E_{\rm Pl}/(2e^{35}) = 3.85$ TeV, the corridor switch $z_{\rm sw} = e^7$ (integer-ledger marker $J = 7$), and today's π-locked inventory $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\ 1 - 1/\pi,\ 1/2\pi^2)$. Every value is a standard cosmological reading of the tabulated chronology that follows; nothing on the axis is fitted.

Table 36: Master chronology of the universe in GG-A10-Cosmos. The table is intentionally redundant: every milestone is shown simultaneously in cosmic time t, characteristic energy scale μ, radial coordinate r, scale factor a, redshift z, its standard cosmological role, and its GG-6/GDOS interpretation.

Epoch / milestone	**Cosmic time** t	**Characteristic scale** μ	$r = \ln(E_{\rm GUT}/\mu)$	a	z	**Standard cosmological meaning**	**GG-6/GDOS reading**
UV boundary surface	10^{-40}–10^{-39} s	$E_{\rm GUT} = 2.28 \times 10^{16}$ GeV	0.0	$\sim 10^{-29}$	$\sim 10^{29}$	Operational beginning of the 4D boundary description; effective fields are initialized at the GUT edge rather than at a formal Planck singularity.	The bulk-to-boundary cosmology turns on at $r = 0$. The initial datum is a finite ultraviolet surface of the anomaly-free throat, not an uncontrolled $a \to 0$ singularity.
Primary inflation (the derived plateau)	$\sim 10^{-39}$–3×10^{-36} s	$V_{\rm inf}^{1/4} = 8.1 \times 10^{15}$ GeV, $H_1 \simeq 1.5 \times 10^{13}$ GeV	$\simeq 1.0$	$a \times e^{59}$	—	Quasi-de Sitter stage one e-fold below the GUT edge; shrinks the comoving Hubble radius and supplies $N \simeq 60$ e-folds ($N^* = 52.6$ observable).	The inflaton is the Σ_r scale modulus on the GDOS-forced plateau; the tilt and tensor ratio are computed here ($n_s = 0.9634$, $r = 0.0038$).
Modulus oscillation and reheating	$\sim 3 \times 10^{-36}$–10^{-28} s	bath $T_{\rm max} \simeq 7 \times 10^{12} \to T_{\rm rh} \simeq 6 \times 10^{10}$ GeV (3×10^{10} at completion, $\rho_r = 3\rho_\phi$)	$\simeq 1.3 \to 12.9$	—	—	The modulus oscillates ($w \simeq 0$) and decays at Planck strength ($\Gamma = M^3/M_{\rm Pl,red}^2$) over $\sim$14 e-folds; the hot Big Bang begins in earnest at 6×10^{10} GeV.	Graceful exit is slow-roll failure at the plateau's edge; the total energy (6×10^{15} GeV) leaks into a bath that peaks at $T_{\rm max}$ and never re-enters the modulus's own scale.
Radiation descent to the lock	$\sim 10^{-28}$–1.6×10^{-14} s	6×10^{10} GeV $\to$ $M_{\rm lock} = 3.85$ TeV	$12.9 \to 29.4$	—	—	One monotone thermal descent; the heavy Majorana sector equilibrates on the way, near 10–20 TeV.	The middle basin is approached from above by ordinary cooling; no second inflation intervenes ($N_2 = 0$).
Lock crossing / Kaxion angle released	$\simeq 1.6 \times 10^{-14}$ s	$M_{\rm lock} = 3.85$ TeV	29.4	—	—	The lock GSSB is crossed promptly (the ledger-action bound: no supercooling, no relic dilution); cold dark matter is seeded by the compact-angle sector at this second crossing rather than by a late ad hoc freeze-out.	The tensor's dual, released here from the pinning the single-patch conjecture presumes, is the Kaxion; its relic density is inherited from the same geometric transition whose residue seeds dark energy.

Continued on next page

Epoch / milestone	**Cosmic time** t	**Characteristic scale** μ	$r = \ln(E_{\rm GUT}/\mu)$	a	z	**Standard cosmological meaning**	**GG-6/GDOS reading**
Leptogenesis window	$\sim 6\times10^{-16}$–2×10^{-13} s	20 TeV $\to$ 1 TeV	27.8 $\to$ 30.8	—	—	Heavy Majorana modes, brought to equilibrium by their own inverse decays near 10–20 TeV, fall out of equilibrium and decay between $z = M_N/T \simeq 1$ and 4, while sphalerons are still active, building a net $B-L$ asymmetry (zero dilution: the lock crossing is prompt).	The CP phase is inherited from holonomy rather than tuned Yukawa phases; the standard thermal window, closed in the retired TeV-reheating reading, is open.
Electroweak seam / sphaleron conversion	$\sim 2\times10^{-11}$ s	100 GeV	33.1	2.3×10^{-15}	4.3×10^{14}	The electroweak plasma remains in the symmetric or near-symmetric regime; sphalerons convert part of the lepton asymmetry into baryon number.	This is the last major stage in which the matter asymmetry can be reprocessed before ordinary radiation domination takes over as the main narrative.
One microsecond milestone	10^{-6} s	~ 1 GeV	37.7	2.3×10^{-13}	4.3×10^{12}	By the first microsecond the universe is a hot radiation bath at roughly the GeV scale; the electroweak story is over and the QCD corridor is next.	In the radial language this corresponds to $r \simeq 37.7$, so the first microsecond already sits almost eight logarithmic units below the lock seam.
QCD crossover / hadronization	10^{-5}–10^{-4} s	150–200 MeV	39.6	1.6×10^{-12}	6.4×10^{11}	Quark-gluon degrees of freedom confine into hadrons; protons and neutrons become the appropriate long-lived baryonic variables of the thermal bath.	The exact $U(1)_B$ symmetry now becomes especially transparent operationally: baryons are stable carriers of the protected bulk charge.
Neutrino decoupling	~ 1 s	1 MeV	44.6	2.3×10^{-10}	4.3×10^{9}	Weak rates drop below the Hubble rate and neutrinos decouple from the electromagnetic plasma.	Any later extra-radiation or dark-radiation effect must be measured against this inherited decoupling ledger; in the chronology this is the reference point for $N_{\rm eff}$ audits.
Neutron-proton freeze-out	~ 1 s	0.8 MeV	44.8	3.0×10^{-10}	3.4×10^{9}	The n/p ratio stops tracking equilibrium exactly and begins to set the helium yield of standard BBN.	The chronology of ordinary matter is now controlled by the usual freeze-out competition Γ/H, but the allowed asymmetry and baryon density were already fixed upstream by the TeV leptogenesis module.

Continued on next page

Epoch / milestone	**Cosmic time** t	**Characteristic scale** μ	$r = \ln(E_{\rm GUT}/\mu)$	a	z	**Standard cosmological meaning**	**GG-6/GDOS reading**
Electron-positron annihilation	1–10 s	0.5 MeV	45.3	4.7×10^{-10}	2.1×10^{9}	Entropy is dumped into the photon bath, changing the relation between photon and neutrino temperatures and locking in the familiar T_ν/T_γ ratio.	This is one of the places where the g_{*s} correction matters in the conversion from z to r; the chronology is logarithmic, not naively linear in temperature.
Deuterium bottleneck opens	$\simeq 200$ s	80 keV	47.1	2.9×10^{-9}	3.4×10^{8}	Photodissociation becomes inefficient enough for deuterium to survive, opening the network that rapidly builds ^{4}He and trace A=7 nuclei.	This is the first truly precision nuclear checkpoint of the framework. Any cosmology module that fails here should be rejected before later structure claims are entertained.
Three-minute milestone / main BBN burn	$\simeq 270$ s	70 keV	47.2	3.4×10^{-9}	2.98×10^{8}	The helium fraction and the deuterium residue are essentially fixed by the end of the first few minutes.	In the radial language the first three minutes correspond to $r \approx 47.2$. By then the universe has traversed almost eighteen additional logarithmic units beyond the lock seam.
Lithium-port window	$\simeq 200$–1500 s	80–30 keV	47.1–48.1	5×10^{-9} (order of magnitude)	2×10^{8} (order of magnitude)	The standard A=7 network is most sensitive in this window; ^{7}Be production and the later lithium problem originate here.	The GDOS $U(1)_B$ port term acts only on the A=7 subspace, depleting ^{7}Be without disturbing the successful D/H and ^{4}He plateaus.
Light-element ledger frozen	10^3–10^6 s	10–1 keV	51.5	2.3×10^{-7}	4.3×10^{6}	Nuclear reaction rates collapse relative to the expansion; the light-element abundances become effectively frozen and are later sampled by astrophysical archaeology.	After this stage the matter content entering recombination and structure growth is already locked, so later cosmology must respect the BBN ledger.
Matter-radiation equality	5×10^4–8×10^4 yr	0.8 eV	58.6	2.9×10^{-4}	3.4×10^{3}	Radiation ceases to dominate; the sound horizon and transfer function become sensitive to the exact equality timing.	The framework places nothing here beyond the standard history: its early residue is forbidden by its own kill criterion and its radiation term is a derived null, so the sound horizon is the standard one.
Recombination onset	3×10^5 yr	0.30 eV	59.6	9×10^{-4}	1.1×10^{3}	The Thomson rate falls toward the Hubble rate; neutral hydrogen begins to form and the photon mean free path grows rapidly.	In GDOS language this is the start of a GSSB transition in the baryon-photon plasma: a continuous scattering orbit begins to freeze into one realized acoustic pattern.
Last scattering / CMB geometric attractor	3.8×10^5 yr	0.26 eV	59.7	9.1×10^{-4}	1090 (atomic; switch e^7 0.4% above)	Primary temperature and polarization anisotropies are released; acoustic phases are frozen and propagate almost freely to us.	The CMB sky is the first great GA of cosmology. Peak positions, baryon loading, and damping are read as deterministic shadows of the locked background and the radial BFFT coordinate.

Continued on next page

Epoch / milestone	**Cosmic time** t	**Characteristic scale** μ	$r = \ln(E_{\rm GUT}/\mu)$	a	z	**Standard cosmological meaning**	**GG-6/GDOS reading**
Baryon drag release / streaming wind	0.4–10 Myr	$0.26 \to 0.01$ eV	63.0	2.3×10^{-2}	4.3×10^{1}	Baryons are freed from photon drag but retain a residual streaming velocity relative to dark matter; the dark ages begin.	The wind is the ΛCDM wind: the gauged-baryon drag once proposed here has no mediator below the lock scale (Part VI).
Dark ages	1–100 Myr	$0.07 \to 0.005$ eV	63.7	4.8×10^{-2}	2.1×10^{1}	No stars yet shine; structure grows linearly and then nonlinearly inside dark-matter potentials.	Most of the observable universe is now matter dominated, but the cold compact-angle dark sector and the baryonic collapse threshold have already set the timing of small-scale collapse.
One-hundred-million-year milestone	10^8 yr	~ 5 meV	63.7	4.8×10^{-2}	~ 20	The earliest massive minihalos are expected to host the first Population III stars or their immediate precursors.	In the radial dictionary this milestone sits near $r \approx 63.7$. Notice how little radial distance separates the dark ages from the present compared with the enormous physical time involved.
Population III ignition	0.1–0.2 Gyr	5–3 meV	64.2	0.05–0.07	20–15	Molecular-hydrogen cooling triggers the first star formation in minihalos; metals and ionizing photons are born.	A lowered effective collapse threshold in the baryonic environment allows this stage to occur earlier than in minimal 4D ΛCDM; the streaming wind is unchanged.
First galaxies / JWST corridor	0.2–0.5 Gyr	3–2 meV	64.3	0.07–0.10	15–9	The first long-lived galaxies and heavy black-hole seeds assemble; UV luminosity functions and stellar-mass functions become directly observable.	The framework attributes the early onset to a lower effective collapse threshold alone (the small-scale transfer-function boost and the reduced baryon wind that used to accompany it are withdrawn, Part VI).
Reionization era	0.5–1 Gyr	2–1.6 meV	64.8	0.10–0.14	9–6	Ionizing photons from stars and accreting objects progressively reionize the intergalactic medium.	This is a boundary-era radiative feedback problem, but its onset and duration inherit the earlier small-scale timing shifts produced by the same locked background.
Mature galaxy and cluster growth	1–5 Gyr	$1.6 \to 0.6$ meV	65.8	0.14–0.4	6–1.5	Nonlinear structure growth builds the familiar cosmic web of halos, clusters, and large galaxies.	This is the era in which the Hubble and S_8 audits must remain consistent with the earlier CMB and galaxy-timing modules rather than being adjusted independently.
Peak cosmic star-formation epoch	~ 3 Gyr	0.7 meV	65.7	0.33	~ 2	The universe reaches its peak rate of converting baryons into stars, metals, and black-hole accretion luminosity.	By this time the background is already close to the late-time π-locked inventory, but matter still dominates the expansion rate.

Continued on next page

Epoch / milestone	**Cosmic time** t	**Characteristic scale** μ	$r = \ln(E_{\mathrm{GUT}}/\mu)$	a	z	**Standard cosmological meaning**	**GG-6/GDOS reading**
Matter-dark-energy equality	$\sim$ 10–11 Gyr	0.30 meV	66.5	0.777	0.287	The matter and vacuum densities become comparable; the decelerating history gives way to accelerating expansion.	Because $\Omega_m = 1/\pi$ and $\Omega_\Lambda = 1 - 1/\pi$ are fixed geometrically, the transition epoch is no longer a coincidence but a derived stage on the same ledger.
$\Omega_m = 1/\pi$ crossing	13.8 Gyr	$\approx T_{\gamma,0}$	66.7	0.995	0.005	The matter fraction falls through exactly $1/\pi$: the full π-locked inventory $(1/\pi,\ 1-1/\pi,\ 1/(2\pi^2))$ is exact at this instant, 75 Myr before the present (Planck central values; timing blur $\pm$160 Myr).	The chart epoch: this is the snapshot the codebook's π-locked fractions refer to (Table 12).
Present stationary attractor	13.9 Gyr	$T_{\gamma,0} = 2.35 \times 10^{-4}$ eV	66.7	1	0	The universe is vacuum dominated but still contains the fossil record of earlier radiation, matter, and structure-building stages.	Today lies at $r \simeq 66.7$, 75 Myr past the crossing: the matter fractions sit fractionally below their π-forms ($\Omega_m = 0.315$ against $1/\pi = 0.318$).
Asymptotic de Sitter future	$\gg 10^{10}$ yr	$\mu \to T_\gamma \to 0$	$r \to +\infty$	$a > 1$ and growing exponentially	$z \to -1$	Matter dilutes away relative to vacuum; the observable patch approaches a de Sitter-like event-horizon state.	The chronology ends on a GA selected by the residual GFF vacuum functional. No big rip is required in the minimal module because the equation of state remains locked at $w = -1$.

63 Worked conversion examples

63.1 Example A: the one-microsecond milestone

Take $t = 10^{-6}\,\mathrm{s}$. In the radiation era Eq. (358) gives

$$10^{-6} \simeq 2.42\, g_*^{-1/2} \left(\frac{T}{\mathrm{MeV}}\right)^{-2}. \tag{387}$$

Using $g_* \sim 60$ as an order-of-magnitude QCD-plasma value, one finds

$$\left(\frac{T}{\mathrm{MeV}}\right)^{2} \sim 2.42\, 60^{-1/2} \times 10^{6} \sim 3.1 \times 10^{5}, \tag{388}$$

hence

$$T \sim 5.6 \times 10^{2}\,\mathrm{MeV} \sim 0.6\,\mathrm{GeV}, \tag{389}$$

which is comfortably in the expected GeV corridor. Taking $\mu \sim 1$ GeV as the rounded milestone scale,

$$r_{\mu\mathrm{s}} = \ln\!\left(\frac{2.28 \times 10^{16}}{1}\right) \simeq 37.7. \tag{390}$$

The formal redshift associated with a GeV-scale bath is

$$1 + z_{\mu\mathrm{s}} \sim \frac{1\,\mathrm{GeV}}{2.35 \times 10^{-13}\,\mathrm{GeV}} \sim 4.3 \times 10^{12}, \qquad a_{\mu\mathrm{s}} \sim 2.3 \times 10^{-13}. \tag{391}$$

So the first microsecond sits at

$$(t, \mu, r, a, z) \sim \left(10^{-6}\,\mathrm{s},\ 1\,\mathrm{GeV},\ 37.7,\ 2.3 \times 10^{-13},\ 4.3 \times 10^{12}\right). \tag{392}$$

63.2 Example B: the deuterium bottleneck (the first three minutes)

Take $t = 200\,\mathrm{s}$, the opening of the deuterium bottleneck. Again using Eq. (358),

$$200 \simeq 2.42\, g_*^{-1/2} \left(\frac{T}{\mathrm{MeV}}\right)^{-2}. \tag{393}$$

For $g_* \sim 3.36$ after electron-positron annihilation, this gives

$$\left(\frac{T}{\mathrm{MeV}}\right)^{2} \sim \frac{2.42}{\sqrt{3.36} \times 200} \sim 6.6 \times 10^{-3}, \tag{394}$$

so

$$T \sim 8 \times 10^{-2}\,\mathrm{MeV} = 80\,\mathrm{keV}. \tag{395}$$

Then

$$r_{\mathrm{D}} = \ln\!\left(\frac{2.28 \times 10^{16}\,\mathrm{GeV}}{8 \times 10^{-5}\,\mathrm{GeV}}\right) \simeq 47.1. \tag{396}$$

The corresponding formal redshift is

$$1 + z_{\mathrm{D}} \sim \frac{8 \times 10^{-5}\,\mathrm{GeV}}{2.35 \times 10^{-13}\,\mathrm{GeV}} \sim 3.4 \times 10^{8}, \qquad a_{\mathrm{D}} \sim 3 \times 10^{-9}. \tag{397}$$

Thus

$$(t, \mu, r, a, z) \sim \left(200\,\mathrm{s},\ 80\,\mathrm{keV},\ 47.1,\ 3 \times 10^{-9},\ 3.4 \times 10^{8}\right). \tag{398}$$

This single computation explains why the first three minutes are such a universal checkpoint in cosmology: visible matter is being fixed at a point only 17 or 18 radial units below the lock seam.

63.3 Example C: the recombination milestone

Take the standard recombination values (the GG corridor-switch marker is $z_{\mathrm{sw}} = e^{7} \approx 1097$, distinct from the decoupling value; §42.1)

$$t_{\mathrm{rec}} \simeq 3.8 \times 10^{5}\,\mathrm{yr}, \qquad z_{\mathrm{rec}} \simeq 1100. \tag{399}$$

Then

$$a_{\mathrm{rec}} = \frac{1}{1 + z_{\mathrm{rec}}} \simeq 9.1 \times 10^{-4}. \tag{400}$$

The photon temperature is

$$T_{\mathrm{rec}} \simeq T_{\gamma,0}(1 + z_{\mathrm{rec}}) \simeq 2.35 \times 10^{-4}\,\mathrm{eV} \times 1100 \simeq 0.26\,\mathrm{eV}, \tag{401}$$

up to the usual small g_{*s} bookkeeping corrections. Therefore

$$r_{\mathrm{rec}} = \ln\left(\frac{2.28 \times 10^{16}\,\mathrm{GeV}}{2.6 \times 10^{-10}\,\mathrm{GeV}}\right) \simeq 59.7. \tag{402}$$

So

$$(t, \mu, r, a, z) \sim \left(3.8 \times 10^{5}\,\mathrm{yr},\ 0.26\,\mathrm{eV},\ 59.7,\ 9.1 \times 10^{-4},\ 1100\right). \tag{403}$$

The striking fact is that recombination lies only about 12.6 radial units below the BBN corridor:

$$\Delta r_{\mathrm{D}\to\mathrm{rec}} \approx 59.7 - 47.1 \approx 12.6. \tag{404}$$

63.4 Example D: the hundred-million-year milestone

Take $t = 10^{8}\,\mathrm{yr}$ and use $z \sim 20$ as a representative first-star-era redshift. Then

$$a_{100\mathrm{Myr}} \sim \frac{1}{21} \sim 4.8 \times 10^{-2}, \tag{405}$$

and the photon temperature is

$$T_{\gamma}(100\,\mathrm{Myr}) \sim 2.35 \times 10^{-4}\,\mathrm{eV} \times 21 \sim 5 \times 10^{-3}\,\mathrm{eV}. \tag{406}$$

Hence

$$r_{100\mathrm{Myr}} = \ln\left(\frac{2.28\times10^{16}\,\mathrm{GeV}}{5\times10^{-12}\,\mathrm{GeV}}\right) \simeq 63.7. \tag{407}$$

Therefore

$$(t,\mu,r,a,z) \sim \left(10^{8}\,\mathrm{yr},\ 5\,\mathrm{meV},\ 63.7,\ 4.8\times10^{-2},\ 20\right). \tag{408}$$

From recombination to the edge of first-light formation the universe thus moves only about four radial units. That is the geometrically concise version of the first-structure era.

63.5 Example E: the present epoch

Today is the easiest example because the scale factor and redshift are normalized by definition:

$$a_0 = 1, \qquad z_0 = 0. \tag{409}$$

Using $T_{\gamma,0} = 2.35\times10^{-4}\,\mathrm{eV}$, one finds

$$r_0 = \ln\left(\frac{2.28\times10^{16}\,\mathrm{GeV}}{2.35\times10^{-13}\,\mathrm{GeV}}\right) \simeq 66.7. \tag{410}$$

Thus

$$(t_0,\mu_0,r_0,a_0,z_0) \sim \left(13.9\,\mathrm{Gyr},\ 2.35\times10^{-4}\,\mathrm{eV},\ 66.7,\ 1,\ 0\right). \tag{411}$$

The distance from the 100 Myr milestone to today is therefore

$$\Delta r_{100\mathrm{Myr}\to0} \simeq 3.0, \tag{412}$$

again showing that most of the universe's visible proper-time age occupies only a very small fraction of the radial hierarchy.

63.6 Example F: matter-dark-energy equality

Finally, use the π-locked inventory. Setting $\rho_m(z) = \rho_\Lambda$ gives Eq. (380),

$$z_{\Lambda m} \approx 0.288, \qquad a_{\Lambda m} \approx 0.777. \tag{413}$$

The corresponding photon temperature is

$$T_{\gamma,\Lambda m} \approx (1+z_{\Lambda m})T_{\gamma,0} \approx 0.30\,\mathrm{meV}, \tag{414}$$

and therefore

$$r_{\Lambda m} = \ln\left(\frac{2.28\times10^{16}\,\mathrm{GeV}}{3\times10^{-13}\,\mathrm{GeV}}\right) \simeq 66.5. \tag{415}$$

So even the recent pivot from deceleration to acceleration occurs in the same narrow r-neighborhood as the present epoch.

64 Radial bookkeeping: how much history is packed into each segment

The seven checkpoint scales and their radial steps are collected in the ladder of Figure 30.

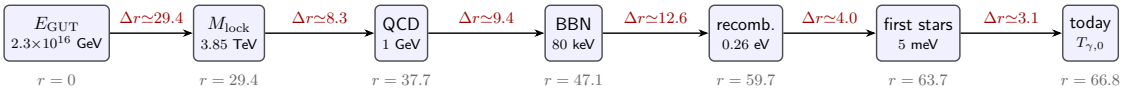

Figure 30: The radial ladder of cosmic history. The seven checkpoint scales of the descent, with the radial step $\Delta r = \ln(\mu_i/\mu_{i+1})$ on each arrow and the running coordinate r beneath each node. The six steps sum to 66.8, reproducing $r_{\rm now} \simeq 66.6$–66.7 (segment rounding): the whole of cosmic history, from the GUT edge to today, is sixty-seven e-folds of scale packed into steps of only a few to thirty — the quantitative face of the front-loaded/back-loaded duality developed below.

What is remarkable is not that the early universe spans the largest total range in r; that is obvious. What is remarkable is that the main visible checkpoints of the history are separated by radial steps of order only a few to ten. The enormous disparity in proper time is therefore partly a coordinate artifact of plotting the history against t instead of against $\ln \mu$.

This is especially helpful when one asks a question such as "where in the total history does recombination sit?" On a linear time axis recombination looks much closer to the beginning than to the present, but still unimaginably later than the first seconds. On a redshift axis it looks remote but not impossible. On the GG-6 axis the answer is cleaner: recombination sits at $r \sim 60$, only about seven units above the present and only about thirteen units above the BBN corridor. That picture is far more compatible with the actual structure of the equations, because the equations organize themselves by powers and logarithms of scale, not by linearly spaced seconds.

64.1 The universe is front-loaded in energy and back-loaded in clock time

This observation can be stated more sharply.

During the earliest stages, a tiny change in proper time corresponds to a vast multiplicative change in the dominant energy scale:

$$\mu \propto t^{-1/2} \qquad \text{in radiation domination,} \tag{416}$$

so a factor of 10^{12} increase in time only lowers the characteristic energy by a factor of 10^6. Because r is logarithmic, that same interval becomes only

$$\Delta r = \tfrac{1}{2}\ln(10^{12}) \approx 13.8. \tag{417}$$

This is why the history from the first second to the first few hundred thousand years is substantial but not huge in the radial language.

Conversely, once the universe enters the late matter and vacuum eras, the energy scale tracked by the photon bath drops very slowly. Billions of years can then pass while the characteristic process scale falls by only a factor of a few or a few tens. In r that means a displacement of only a few units. The late universe is therefore back-loaded in proper time but thin in the radial coordinate.

This duality between front-loading in energy and back-loading in clock time is one of the central conceptual outcomes of Part VIII. It explains why the early universe is where almost all of the hierarchy is generated, while the late universe is where almost all of the precision diagnostics live.

64.2 A second summary table: intervals rather than milestones

The next table complements the master chronology. Instead of listing absolute milestones, it lists the interval between them and the corresponding physical work done in that interval.

Table 38: Radial bookkeeping for the main segments of cosmic history. The table shows that the late universe is huge in years but narrow in r, while the early universe is tiny in years but wide in r. The columns give the interval, its scale drop, the radial span, the approximate elapsed time, and the physical work done in it. Reading the third and fourth columns together is the point: the late universe is vast in years and narrow in radial coordinate, and the early universe is the reverse.

Interval	Characteristic scale drop	Δr	Approximate time interval	Main physical work done
GUT edge → lock seam	$2.3 \times 10^{16}\,\mathrm{GeV} \to 3.85\,\mathrm{TeV}$	29.4	$10^{-39}\,\mathrm{s} \to 1.6 \times 10^{-14}\,\mathrm{s}$	Boundary description turns on; inflation runs on the plateau and ends at its edge; reheating completes; the radiation descent reaches the lock.
Lock seam → one microsecond	$3.85\,\mathrm{TeV} \to 1\,\mathrm{GeV}$	8.3	$1.6 \times 10^{-14}\,\mathrm{s} \to 10^{-6}\,\mathrm{s}$	The lock is crossed and dark matter is seeded, baryogenesis occurs, electroweak processing ends.
One microsecond → three minutes	$1\,\mathrm{GeV} \to 80\,\mathrm{keV}$	9.4	$10^{-6}\,\mathrm{s} \to 200\,\mathrm{s}$	QCD crossover, neutrino decoupling, n/p freeze-out, deuterium bottleneck, helium synthesis, lithium window.
Three minutes → recombination	$80\,\mathrm{keV} \to 0.26\,\mathrm{eV}$	12.6	$200\,\mathrm{s} \to 3.8 \times 10^{5}\,\mathrm{yr}$	Radiation-dominated bridge, equality, acoustic processing, last scattering.
Recombination → 100 Myr	$0.26\,\mathrm{eV} \to 5\,\mathrm{meV}$	4.0	$3.8 \times 10^{5}\,\mathrm{yr} \to 10^{8}\,\mathrm{yr}$	Dark ages, baryon settling, first-star preparation.
100 Myr → today	$5\,\mathrm{meV} \to 2.35 \times 10^{-4}\,\mathrm{eV}$	3.1	$10^{8}\,\mathrm{yr} \to 13.9\,\mathrm{Gyr}$	Galaxies, reionization, mature structure, matter-vacuum transition, present attractor.
Today → asymptotic future	$T_{\gamma,0} \to 0$	unbounded	$13.9\,\mathrm{Gyr} \to \infty$	Approach to a de Sitter-like vacuum attractor with a finite event horizon.

64.3 Why this bookkeeping is scientifically useful

This interval view does three kinds of work.

First, it prevents chronological confusion. It becomes immediately clear that BBN is not "almost simultaneous" with the electroweak story just because both happen early in absolute terms. There is a radial gap of about nine units between the GeV corridor and the 80 keV corridor, which is plenty of room for qualitatively new physics to happen—and in fact it does.

Second, it highlights where model freedom is actually dangerous. The largest radial block after

reheating is the radiation descent to the lock; the largest after the lock is the interval between BBN and recombination. That second interval is precisely the one most heavily audited by the CMB and by equality-related observables. Any hidden extra radiation, decays, or late entropy dump inserted there will show up prominently.

Third, it explains why late-time precision tensions are so constraining. From 100 Myr to today there are only about three radial units left. A theory has very little room to maneuver there without ruining the rest of the ledger.

64.4 The chronology in one radial sentence

One can compress the whole chronology into a single radial sentence:

$$0 \;\to\; 29.4 \;\to\; 37.7 \;\to\; 47.1 \;\to\; 59.7 \;\to\; 63.7 \;\to\; 66.7 \;\to\; \infty. \tag{418}$$

Each arrow is a multiplicative fall in energy, and each number can now be associated with a concrete cosmic stage:

$$\text{GUT edge} \to \text{lock seam} \to \text{one microsecond} \to \text{three minutes} \to \text{recombination} \to \text{first light} \to \text{today} \to \text{asymptotic future}.$$

This is arguably the most compact version of the history written in the present paper.

65 The milestone ladder in one glance

1. the end of the electroweak/TeV prelude and the entrance into the hadronic bath,
2. the fixation of the light-element inventory,
3. the freezing of the acoustic information in the CMB,
4. and the ignition of the first luminous structures.

What makes these four times so instructive is that each one marks the completion of a different kind of physical work.

At one microsecond, the universe has not yet built atoms, nuclei, or galaxies, but it has already completed the geometric tasks that dominate the very beginning: UV initialization, the inflationary attractor and its exit, reheating, the radiation descent to the lock seam, the lock crossing and the birth of the Kaxion dark sector, and the TeV leptogenesis corridor. In other words, by the time the universe is one microsecond old, the deep ultraviolet and matter-origin work of the theory is already behind it.

At three minutes, the universe has completed the fixation of the ordinary visible matter inventory. This is why every later observational era—from recombination to galaxies to stars to observers—inherits a baryonic ledger that was set astonishingly early. The late universe only rearranges matter whose total abundance and light-element composition were already decided.

At 3×10^5 years, the universe has completed the information-freezing work of recombination. The primary CMB sky is released and the acoustic phase pattern becomes a stable fossil. In the language of the paper, this is the first truly macroscopic geometric attractor seen directly by observers.

At 10^8 years, the universe has completed the long dark bridge between a transparent but dark background and a luminous structured cosmos. The first stars, and then the first galaxies, appear as the earliest large-scale nonlinear consequences of everything that came before.

Seen this way, the four milestones are not arbitrary dates. They are the four best "memory cells" of the whole chronology:

ultraviolet origin, visible matter, frozen information, first light.

Table 39: The four milestones plus the present epoch, displayed as a compact five-clock ladder. This table is a compressed version of the master chronology and is especially useful for presentation, teaching, and cross-checking the Part VIII narrative. The five clock columns are one milestone read five ways, and the last column says what is finished by then.

Milestone	t	μ	r	a	z	**What is completed by then?**
First microsecond	10^{-6} s	~ 1 GeV	~ 37.7	$\sim 2.3\times10^{-13}$	$\sim 4\times10^{12}$	Inflation is already over, reheating is complete, the dark sector is seeded, and the TeV baryogenesis corridor has already acted.
First three minutes	200–270 s	80 to 70 keV	~ 47.1	$\sim 3\times10^{-9}$	$\sim 3\times10^8$	The visible baryonic composition is essentially fixed; helium, deuterium, and the A=7 sector have reached their primordial ledger.
First 3×10^5 years	3 to 4×10^5 yr	~ 0.3 eV	~ 59.7	$\sim 10^{-3}$	~ 1100	The CMB is released; the primary acoustic information is frozen into a geometric attractor.
First 10^8 years	10^8 yr	~ 5 meV	~ 63.7	$\sim 5\times10^{-2}$	~ 20	The dark ages are ending and the first stars are turning on; the universe becomes luminous.
Present epoch	13.9 Gyr	2.35×10^{-4} eV	~ 66.7	1	0	The late-time π-locked inventory is visible directly; vacuum domination has already begun.

The compressed ladder exposes two points that are easy to miss when the history is narrated only qualitatively.

The first is that the distance from the one-microsecond milestone to the three-minute milestone is almost ten radial units, whereas the distance from the first 10^8 years to the present is only about

three. Therefore the early history still contains more hierarchy per unit milestone than the entire visible late history.

The second is that the redshift variable becomes observationally natural only rather late. A formal redshift can be assigned to the one-microsecond or three-minute stages, but there are no direct observers there and the more native variables are t, T, and r. By recombination and certainly by 10^8 years, however, the redshift language becomes indispensable because the history is now read by astronomy rather than by nuclear or particle physics.

This transition in which variable is operationally natural is itself part of the theory's measurement philosophy. OGGEP says that different eras are read out through different boundary instruments. A collider-style or thermal-bath variable is natural in the early universe; a redshift and luminosity variable is natural in the late universe. The chronology remains one chronology because all those instruments are linked by the same underlying coordinate dictionary.

66 Chronological synthesis

> The universe begins on a finite ultraviolet boundary at E_{GUT}, descends geometrically to the inflationary plateau one e-fold below the GUT edge, exits by the failure of slow roll at the plateau's edge, reheats through the modulus's Planck-strength decay, crosses the lock promptly to seed dark matter and the terminus, generates baryon asymmetry in the TeV corridor, fixes visible matter in the first three minutes, freezes the CMB as a geometric attractor at $r \sim 60$, ignites the first stars near $r \sim 64$, and reaches the present π-locked vacuum-matter attractor by $r \sim 67$.

That sentence is long, but it is logically tight. Every clause is inherited from Parts II through VII, and every clause now has an explicit chronological placement. That is the real purpose of Part VIII. It turns a collection of problem-by-problem modules into one trackable cosmic story.

The deeper moral is equally important. The universe is violently nonuniform in ordinary time, but it is comparatively well organized in the logarithmic GG-6 coordinate. The first microsecond, the first three minutes, recombination, first-light formation, and the present are not disconnected wonders; they are milestones on a single finite descent in r, followed by an asymptotic drift toward $r \to \infty$.

This is exactly the sort of synthesis a unification paper ought to deliver.

There is a quiet philosophical lesson in the logarithmic clock. Asked in ordinary years, "when did the universe do this?" scatters its answers across twenty orders of magnitude, and the history reads as a sequence of unrelated miracles. Asked in the geometry's own coordinate, the same events fall at even steps of a single descent: the universe was never rushing and never idle; it was walking down a ladder at the pace the ledger set. What a calendar makes miraculous, a ruler makes lawful.

The narrative is now complete, and precisely because it is complete it must submit to an audit. Part IX stops telling the story and starts grading it: every number this paper has used, classified — Derived, Derived-conditional, or Asserted — with nothing quietly upgraded and nothing hidden in prose. The reader who has trusted the tale this far is owed the ledger behind it.

67 Part VIII in plain English: why a chronology block earns its place

A chronology of the universe is standard material and this one would not need writing if it were only that. It is here because the framework's central bookkeeping device — the radial coordinate — has to be shown to reproduce the ordinary history before it can be used to say anything new about it, and a translation table between cosmic time, energy scale, radial coordinate, scale factor and redshift is the only honest way to show that.

The table is deliberately redundant. Every milestone appears in all five coordinates at once, which makes the block long and makes it checkable: a reader who suspects the radial reading of smuggling something in can convert any row back to standard variables and see that nothing has been added.

The one thing the Part does claim is the shape of the packing. The early universe occupies a vast radial interval and a negligible span of years; the late universe occupies almost no radial interval and almost all of the years. That is a statement about the geometry rather than about the physics of any epoch, and it is what makes the descent picture of Part IV a picture of cosmic history rather than a metaphor laid over it. The worked conversions at the end are there so the claim can be checked with a pencil.

Part IX
Cosmological codebook: Derived / Derived (cond.) / Asserted

The cosmological observables of the π-locked inventory, classified by their provenance under the MSMF/OGN principle; the per-parameter audit identifying which entries are derived, which are derived modulo a named lemma, and which remain asserted pending further constitutional analysis

68 Codebook overview and the MSMF/OGN principle on the cosmological-boundary stratum

This Part is the cosmological analog of the twenty-eight-parameter codebook of GG-A9-SM for the Standard-Model sector. Where GG-A9-SM classified each electroweak/flavor observable as Derived (fully derived from non-physics axioms), Derived (cond.) (derived modulo a single explicitly-named lemma), or Asserted (reported as the empirical input pending a structurally-named lemma), the present Part classifies twenty-one cosmological observables on the cosmological-boundary stratum. The MSMF/OGN principle — minimal-structure-of-minimal-form / Ostrogradsky-graded-natural-degree — selects, for each observable, the unique constitutional cost-minimum admissible under the inherited ledger.

The cosmological observables of interest are the π-locked density inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_{\rm DM}$; the one Hubble background H_0, the sound horizon, and the recombination switch redshift; the scalar perturbation amplitude and tilt A_s and n_s; the structure-formation parameters σ_8 and S_8; the radiation-content effective number $\Delta N_{\rm eff}$; the inflationary energy scale and the tensor-to-scalar ratio r; the Kaxion mass and decay constant; and any feature-dependent supplementary channel (PBH mass spectrum slice, primordial-GW band).

The remainder of this Part provides the codebook itself. Section 68.1 states the twenty-one-observable codebook table. Section 68.2 states the per-parameter MSMF/OGN audit. Section 68.3 states the Hubble two-stage corridor named theorem.

68.1 The cosmological codebook

Table 40 is the cosmological analog of the twenty-eight-parameter Codebook Provenance Table of GG-A9-SM. It lists twenty-one cosmological observables and classifies each by its provenance under the MSMF/OGN principle on the cosmological-boundary stratum. The provenance map of Figure 32 traces the same classification visually, carrying each observable back through its named conjecture to the single geometric origin.

The classification triage is identical to GG-A9-SM:

- **Derived (D)**: fully derived from non-physics axioms with every factor traceable to upstream theorems of the Ultra Eight GG-Theory cohort. No empirical content is required to obtain the value.
- **Derived (cond.) (D*)**: derived modulo a named lemma — a single one for seventeen entries; D4 rests on Lemma I together with the companion's pinning conjecture, and A15 on Lemmas IV and V$'$. The lemma is an open research target with a definite mathematical statement; once proved, the entry is promoted to D.
- **Asserted (A)**: reported as the empirical input pending a structurally-named lemma whose proof would close the entry. The lemma is identified by name in the per-parameter audit of Section 68.2.

Of the twenty-one entries, one is fully derived — $\Delta N_{\rm eff} \leq 0.03$, a null forced by the field content — nineteen are derived modulo a named lemma (the three π-locked fractions Ω_m, Ω_Λ, and Ω_b, the dark-energy equation of state w, $\Omega_{\rm DM}$, the baryon asymmetry η_B, the lithium-7 abundance, the one background Hubble rate against its two comparators, the decoupling redshift z_* (the atomic chain on the fractions; §42.1), σ_8 and S_8 against their two comparators, the damping multipole ℓ_D, the Kaxion mass and decay constant, the scalar tilt n_s and the tensor ratio r (Lemma IV, revised), and the inflationary energy scale $V_{\rm inf}^{1/4}$ (Lemma IV for the class and Lemma V$'$ for the height)), and one remains asserted: the scalar amplitude A_s, the launch sector's declared input, awaiting the heat-kernel computation of Lemma V$'$. No observable in the inventory is fitted to itself, and the continuous inputs are three, all declared: the amplitude A_s (row A8), the epoch t_0 at which the π-locked fractions are read (Conjecture VIII's declared input), and the bulk Green–Schwarz normalization $f_{\rm bulk}$, calibrated by the dark-matter abundance in GG-C1-Kaxion and entering rows D13–D14; the rest of the input ledger is the Geometric Universal Five Constants GU-5 $= \{c, G_N, \hbar; k_B, Q_B\}$, the structural constants (π, e), and the proved OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, per the zero-free-parameters discipline of GG-A4-Euclid. (Earlier versions said "no adjustable continuous parameter" and "two declared inputs"; the calibrated scale is the third, and saying so costs the paper nothing it could defend losing.)

Table 40: Cosmological codebook: twenty-one observables of the cosmological-boundary stratum classified by their provenance under the MSMF/OGN principle. Empirical values from Planck Collaboration (2018) where applicable; the ladder value from the H0DN consensus (Casertano et al., 2026), with Riess et al. (2022) the leading single-team determination; the lensing S_8 value from Asgari et al. (2021); the 1998–1999 acceleration-discovery papers (Perlmutter et al., 1999; Riess et al., 1998) are cited for the historical reintroduction of Λ, not for parameter values. Status D = Derived (all factors traceable to upstream theorems); D* = Derived (cond.), modulo a named lemma (two for D4 and A15); A = Asserted (pending a structurally-named lemma). The per-parameter audit in Section 68.2 states each named lemma explicitly. The three π-locked inventory entries (Ω_m, Ω_Λ, Ω_b) are Derived-conditional for their *coupled π-form* under Conjecture VIII; their absolute present-day values additionally inherit the observation epoch (T_0) as the single L4 calibration input (§32.11).

Observable	Symbol	Empirical value	Status	First-order derivation source	Named lemma (if D* or A)
Matter density fraction	Ω_m	0.315 ± 0.007	D*	$\Omega_m = 1/\pi \approx 0.3183$; closed-packet topology of the cosmological-boundary stratum (Theorem 74.1, Part X of the present paper; bulk geometry inherited from Arisaka, 2026, A2-GG6)	Closed-packet-integration lemma (Conj. VIII)
Cosmological constant fraction	Ω_Λ	0.685 ± 0.007	D*	$\Omega_\Lambda = 1 - 1/\pi \approx 0.6817$; residual F_{GFF} on the late-time slice (Part IV; Arisaka, 2026, A6-GRIP)	Same lemma (via unit-trace identity)
Dark-energy equation of state	w	-1.028 ± 0.032 (Planck Collaboration, 2018)	D*	$w = -1$ exactly; frozen residual vacuum export of the completed late-time geometric transition (Part IV)	Vacuum-residual-exactness lemma (Conj. IX)
Baryon density fraction	Ω_b	0.0493 ± 0.0006	D*	$\Omega_b = 1/(2\pi^2) \approx 0.0507$; baryonic gauge protection from the inherited Standard-Model ledger (Arisaka, 2026, A9-SM). *Note:* the parameter-free prediction 0.0507 deviates from the empirical 0.0493 ± 0.0006 at the 2.3σ level. The D* classification reflects that the structural derivation is closed with no remaining free parameter; the residual 2.3σ discrepancy is a candidate next-order correction; a common second-order normalization is one of the three readings of §84.5, and the paper commits to none of them.	Closed-packet-integration lemma

Observable	Symbol	Empirical value	Status	First-order derivation source	Named lemma (if D* or A)
Dark-matter density fraction	$\Omega_{\rm DM}$	0.2657 ± 0.0070 ($\Omega_m - \Omega_b$; $\Omega_c = 0.2645$ excludes the neutrino share)	D*	$\Omega_{\rm DM} = \Omega_m - \Omega_b = 1/\pi - 1/(2\pi^2) \approx 0.2676$; Kaxion compact-angle relic (Part V; one scale calibrated)	Kaxion-misalignment-completeness lemma, with the single-patch history it presupposes (C-Kaxion-d), necessary by the isocurvature bound
Baryon-to-photon ratio	η_B	$(6.12 \pm 0.04) \times 10^{-10}$ (Planck Collaboration, 2018)	D*	$\eta_B \approx 5 \times 10^{-10}$ (central 4.89; $\times 3^{\pm 1}$; the time-resolved integration lands a factor of about ten below the ledger road's 6.18, App. I.6) from resonant leptogenesis on the inherited GG-A9-SM neutrino sector (UV completion of the Weinberg operator); geometric throughput $\varepsilon^{\rm res}\kappa = e^{-K/2}$, $K = 35$ (Part V, App. I)	Resonant-leptogenesis-normalization lemma (Conj. VI; partial-resonance condition)
Lithium-7 abundance	$^7{\rm Li}/H$	$(1.6 \pm 0.3) \times 10^{-10}$ (Fields, 2011)	D*	$(^7{\rm Li}/H)_{\rm GG} = {\rm SBBN}/\pi \approx 1.6 \times 10^{-10}$; the closed-packet $1/\pi$ acting on the $A = 7$ port (Part V)	Lithium-port-π-suppression lemma (Conj. X)
Hubble constant, against the distance ladder	$H_0^{\rm late}$	73.50 ± 0.81 km/s/Mpc (Casertano et al., 2026)	D*	The same background 66.8 ± 0.2: one Friedmann background, one rate on every probe that reads it (CMB, BAO, sirens, time-delay lensing). The ladder's 73.50 is read as a measurement carrying a $\approx 9\%$ systematic that the framework *predicts*; the disagreement is not accommodated, and the ladder-free late probes are the discriminant (Part VI, GG-C2-Hubble)	Stratum-selection-of-the-vacuum-residue lemma (Lemma II)
Hubble constant, against the microwave background	$H_0^{\rm CMB}$	67.4 ± 0.5 km/s/Mpc (Planck Collaboration, 2018)	D*	Derived 66.8 ± 0.2 with zero sky input from $\zeta_\Lambda = \pi/(LK)$ (Route I, Section 29.8), and 66.6 ± 0.2 from the π-locked fractions read against the measured densities; -0.9%	Same lemma
Decoupling redshift	z_*	1089.92 ± 0.25 (Planck Collaboration, 2018)	D*	Two quantities (§42.1): the corridor switch $z_{\rm sw} = e^7 = 1096.6$ is topological (OGN-5 integer $J = 7$, Theorem 68.8); the operational decoupling from the framework's own atomic chain on the π-locked inventory lands on the Planck value at the 0.1% level, 0.4% below the switch	Closed-packet-integration lemma (Conj. VIII, via the π-locked inputs; switch ledger-fixed under Conj. II)

Observable	Symbol	Empirical value	Status	First-order derivation source	Named lemma (if D* or A)
Scalar perturbation tilt	n_s	0.9649 ± 0.0042 (Planck Collaboration, 2018)	D*	$n_s = 0.9634$ exact on the GDOS-forced plateau at the reheat-matched $N^* = 52.6$ ($1 - 2/N^* = 0.9620$ at leading order; Part IV)	Plateau-class-and-prompt-crossing (Lemma IV, revised)
Scalar perturbation amplitude	A_s	$(2.10 \pm 0.04) \times 10^{-9}$ (Planck Collaboration, 2018)	A	The launch sector's declared input; the GDOS heat-kernel normalization on the same plateau is the named closure target (Section 68.2.4)	GDOS-heat-kernel-amplitude (Lemma V′)
Tensor-to-scalar ratio	r	< 0.036 (95% CL) (BICEP/Keck Collaboration, 2021)	D*	derived plateau class + reheat-matched count $\Rightarrow r = 12/N^{*2} = 0.0043$ at leading order, 0.0038 exact on the plateau at $N^* = 52.6$ (Starobinsky, 1979) (Part IV, §28.10)	Plateau-class-and-prompt-crossing (Lemma IV, revised) — the same lemma as the tilt
Matter clustering amplitude	σ_8	0.811 ± 0.006 (Planck Collaboration, 2018)	D*	0.80 ± 0.01, computed from the declared amplitude A_s and the exact tilt on the π-locked background (§43.3); -1.4%	GDOS-heat-kernel-amplitude (Lemma V′): the row carries the amplitude's grade
S_8 clustering parameter, against the CMB	S_8	0.832 ± 0.013 (Planck Collaboration, 2018)	D*	0.825 ± 0.010: $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$, both factors derived above (§43.3); -0.8%	Same lemma
S_8 clustering parameter, against cosmic shear	S_8	$0.759^{+0.024}_{-0.021}$ (Asgari et al., 2021)	D*	The same derived 0.825 ± 0.010: one number, two comparators; $+8.7\%$, $+2.5\sigma$ — the framework shares the S_8 tension and does not resolve it (§43.3)	Same lemma
Damping-tail multipole	ℓ_D	$\approx 1345 \pm 15$ (inferred from the $100\theta_D$ of Planck Collaboration 2018; the ± 15 a convention spread)	D*	$\ell_D = 1350 \pm 3$ from the diffusion integral on the π-locked inventory (Part VI); the BFFT overlay $K(2\pi)^2 \approx 1382$ is carried separately as a conjecture ($+2.3\%$ above the Boltzmann value)	Closed-packet-integration lemma (Conj. VIII); the BFFT overlay gated by Conj. VII

Observable	Symbol	Empirical value	Status	First-order derivation source	Named lemma (if D* or A)
Effective neutrino number deviation	ΔN_{eff}	$N_{\mathrm{eff}} = 2.99 \pm 0.17$ (Planck Collaboration, 2018)	D	$\Delta N_{\mathrm{eff}} \leq 0.03$, computed: $N_{\mathrm{eff}} = 3.044$, the Standard Model's own value (Akita and Yamaguchi, 2020), plus a freeze-in residue bounded above by the full-thermalization value 0.027 of one real scalar decoupling at $g_{*s} = 106.75$. A derived null to within that ceiling, on the same footing as the fifth-force null ($\eta = 0$) of the cohort's fifth-force paper, GG-P7-5thF: the termination of $U(1)_B$ by a boundary condition (GG-A3-BI6) leaves no vector below the lock scale, and the Kaxion — relativistic at $T \sim 1$ MeV, as any 8.1 μeV particle is — is populated only by gluon-portal freeze-in at T_{rh}, at 0.1–1 of equilibrium. Two-sided: any $\Delta N_{\mathrm{eff}} \gtrsim 0.2$ falsifies the termination, seven times the ceiling. (An earlier version printed $\Delta N_{\mathrm{eff}} = 0.04$ "inherited from the particle paper"; that was the Standard Model's own 0.044 counted a second time, and GG-A9-SM contains no such residue. The legacy $\Delta N_{\mathrm{eff}} \approx 0.30$ of the retired early-dark-energy corridor is recorded in Appendix K.10 as history.)	*none*
Kaxion mass	m_{Kaxion}	8.1 μeV, $\pm 1.3\%$ at fixed f_K ($\nu_K = 1.95$ GHz; ladder $[1.59, 2.13]$ GHz)	D*	the QCD image of a decay constant calibrated through the π-locked abundance closure at the benchmark cap loading, the ladder an ansatz; pseudo-Nambu–Goldstone mass from (Peccei and Quinn, 1977; Weinberg, 1978; Wilczek, 1978) template	Compact-angle-mass-uniqueness lemma; closure executed in companion GG-C1-Kaxion (open remainder: the orientation convention behind the calibration, GG-A3-BI6 O14, and a defining computation for the cap loading; entry kept D* pending the lemma)

Observable	Symbol	Empirical value	Status	First-order derivation source	Named lemma (if D* or A)
Kaxion decay constant	f_{Kaxion}	7.1×10^{11} GeV ($f_{\mathrm{bulk}} = 3.04 \times 10^{13}$ GeV calibrated, $\Xi_{\mathrm{cap}} = 43$ benchmark)	D*	Abundance closure calibrates the bulk normalization at the benchmark loading: the one calibrated scale, declared rather than derived (GG-A3-BI6 O14 open); the row is D* because everything downstream of that scale is derived, and A only if the scale itself is the entry	Same lemma
Inflationary energy scale	$V_{\mathrm{inf}}^{1/4}$	— (not measured; the framework's value is 8.1×10^{15} GeV, the derived class at the declared amplitude)	D*	The GDOS-forced Starobinsky/ξ plateau one e-fold below E_{GUT}, with graceful exit and prompt lock crossing (Part IV; Starobinsky, 1980 for the plateau template)	Plateau-class-and-prompt-crossing (Lemma IV, revised) for the class; GDOS-heat-kernel-amplitude (Lemma V′) for the height, which scales as $A_s^{1/4}$

Table 41: The complete numerical bottom line of GG-A10-Cosmos: the data behind Figure 31. Prediction, observed value, relative deviation, pull (with the theory tolerance $\sigma_{\rm th}$ and its source where one applies) and provenance grade (D derived, D* derived-conditional, A asserted); conventions with the preview copy in Part I (Table 2).

Entry	GG prediction	Observed	Dev	Pull	Grade
The π-locked inventory					
Ω_m	$1/\pi = 0.3183$	0.315 ± 0.007	$+1.05\%$ (cal. input)	$+0.5$ (—)	D*
Ω_Λ	$1 - 1/\pi = 0.6817$	0.685 ± 0.007	-0.48% (cal. 0.00)	-0.5 (0.0)	D*
w	-1 (exact)	-1.028 ± 0.032	$+2.72\%$	$+0.9$	D*
Ω_b	$1/(2\pi^2) = 0.05066$	0.0493 ± 0.0006	$+2.76\%$ (cal. $+1.69$)	$+2.3$ ($+1.4$)	D*
$\Omega_{\rm DM}$	$1/\pi - 1/(2\pi^2) = 0.2676$	0.2657 ± 0.0070	$+0.73\%$ (cal. -0.31)	$+0.3$ (-0.1)	D*
Baryogenesis and nucleosynthesis					
η_B	$\approx 5 \times 10^{-10}$ (central 4.89, closed form on two states, K=35; the priced 1.9 is 2.2σ short, the time-resolved 0.58 a factor ten, App. I.6)	$(6.12 \pm 0.04) \times 10^{-10}$	-20%	-0.3 ($\sigma_{\rm th}$ 4×10^{-10}, asserted; $\ln 3$ in the log, the $\times 3^{\pm 1}$ band, gives -0.2)	D*
$^7{\rm Li}/H$	SBBN/$\pi = 1.59 \times 10^{-10}$	$(1.6 \pm 0.3) \times 10^{-10}$	-0.6%	0.0	D*
The Hubble constant, recombination, the acoustic spectrum					
$H_0^{\rm CMB}$	66.8 ± 0.2 (Route I, zero sky input)	67.4 ± 0.5	-0.89%	-1.1 ($\sigma_{\rm th}$ 0.2, Route I; 2025: -1.1)	D*
$H_0^{\rm late}$	66.8 ± 0.2 (one background; the ladder's excess a predicted systematic)	73.50 ± 0.81	-9.12%	-8.0 ($\sigma_{\rm th}$ 0.2, same)	D*
z_*	1091.5 (atomic chain at the derived h; switch e^7 annotated)	1089.92 ± 0.25	$+0.15\%$	$+0.8$ ($\sigma_{\rm th}$ 0.18%, modeling)	D*
ℓ_D	1350 ± 3 (Boltzmann; in CLASS's $k_D D_M$ convention 1956 vs. 1947 ± 5, $+0.5\%$, $+1.6$)	$\approx 1345 \pm 15$ (convention spread)	$+0.37\%$	$+0.3$ ($\sigma_{\rm th}$ 3, Boltzmann)	D*
ω_b/ω_m	$1/2\pi = 0.15915$ (h-free, epoch-free)	0.15643 ± 0.0016	$+1.74\%$	$+1.7$	D*
$100\,\theta_*$	1.0381 (r_s/D_M at the carried rate)	1.04109 ± 0.00030	-0.29%	-2.4 ($\sigma_{\rm th}$ 0.12% via z_*; -10 on $\sigma_{\rm obs}$ alone)	D*
peak ratio 2/1	0.449 (image of $1/2\pi$)	0.453 ± 0.008	-0.8%	-0.5	D*
peak ratio 3/1	0.439 (image of $1/2\pi$)	0.443 ± 0.009	-0.9%	-0.4	D*
The inflationary spectrum (the derived plateau; prompt lock crossing, $N_2 = 0$)					
n_s	0.9634 exact (N^*=52.6; 0.9620 at leading order)	0.9649 ± 0.0042	-0.16%	-0.4 ($\sigma_{\rm th}$ 0.00034, half e-fold)	D*
r	0.0038 exact (0.0043 at leading order)	< 0.036 (95%, BK18, BICEP/Keck Collaboration, 2021)	$\times 9$ below	—	D*
A_s	anchored (the declared input; warp cand. 75%)	$(2.10 \pm 0.04) \times 10^{-9}$	0	0	A
Clustering and radiation					
σ_8	0.80 ± 0.01 (from A_s on the π-locked background)	0.811 ± 0.006	-1.4%	-0.9 ($\sigma_{\rm th}$ 0.010, amplitude)	D*
S_8 (vs. CMB)	0.825 ± 0.010	0.832 ± 0.013	-0.8%	-0.4 ($\sigma_{\rm th}$ 0.010, same)	D*
S_8 (vs. cosmic shear)	the same 0.825 ± 0.010	$0.759^{+0.024}_{-0.021}$	$+8.7\%$	$+2.5$ ($\sigma_{\rm th}$ 0.010, same)	D*
$N_{\rm eff}$	3.044 ($\Delta N_{\rm eff} \leq 0.03$, derived null)	2.99 ± 0.17	$+1.8\%$	$+0.3$ ($\sigma_{\rm th}$ $+0.027$, one-sided)	D

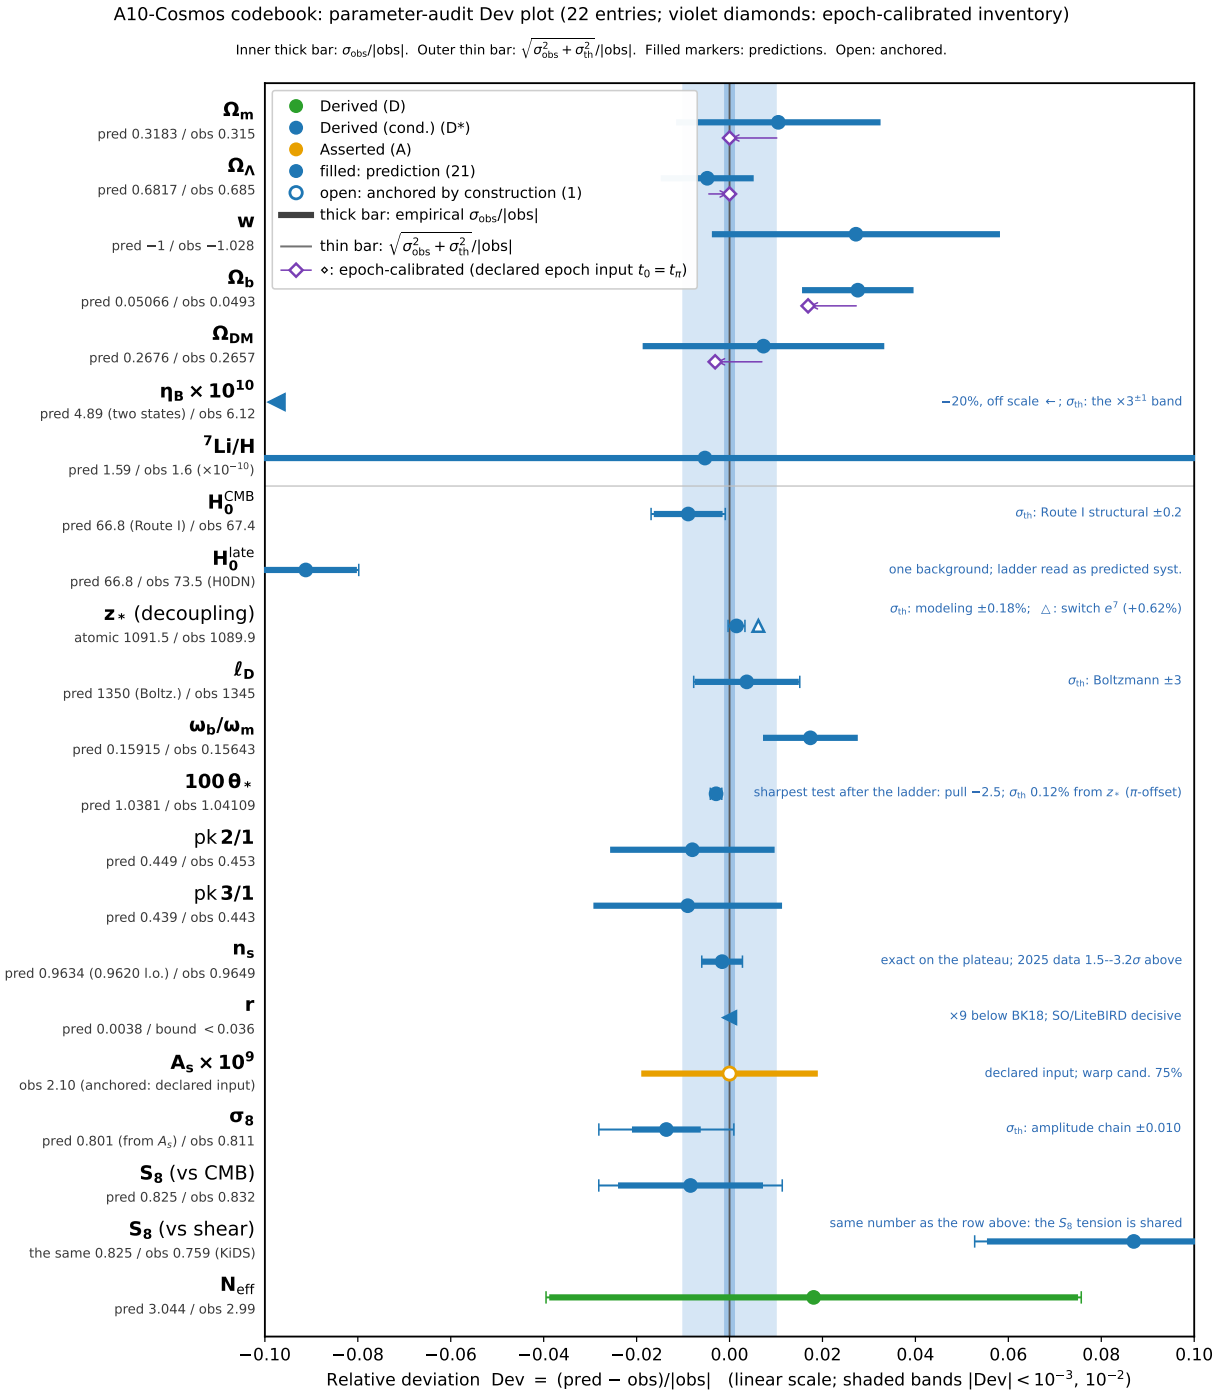

Figure 31: A10-Cosmos parameter-audit Dev plot. Fractional deviation (pred − obs)/|obs| for the twenty-two rows of Table 41, with nested precision bands at $|\mathrm{Dev}| < 10^{-3}, 10^{-2}$. *Filled* markers are predictions (twenty-one, of which about eight are independent); *open* markers are anchored by construction (Dev = 0); color encodes status (green = Derived, blue = Derived (cond.), gold = Asserted); *violet open diamonds* mark the epoch-calibrated positions of the four π-inventory rows. Error bars, in-panel annotations, and the z_*/corridor-switch marker follow the *Reading conventions* paragraph below; full data, sources, and uncertainties: Table 41.

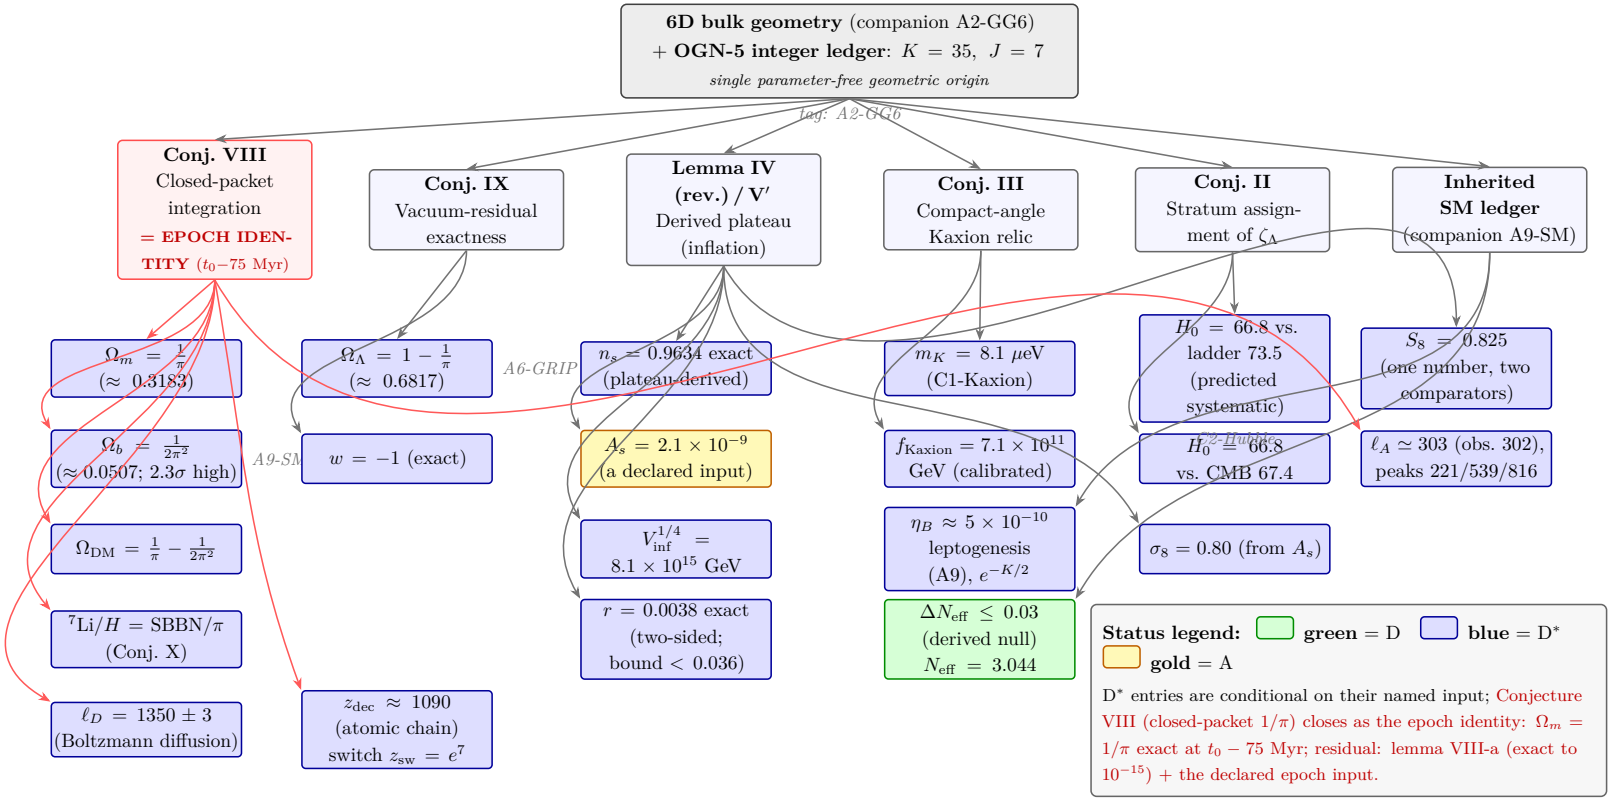

Figure 32: Provenance map of the cosmological codebook of GG-A10-Cosmos. A single parameter-free geometric origin (Stratum 1: the 6D bulk geometry of companion A2-GG6 together with the OGN-5 integer ledger $K = 35$, $J = 7$) flows downward through a row of named conjectures and inherited mechanisms (Stratum 2) to all 21 cosmological observables (Strata 3–4). Colors follow the status legend (lower right); Conjecture VIII (closed-packet $1/\pi$), flagged in red, is the load-bearing conditional. The color-coding, the epoch identity behind Conjecture VIII, the two-quantity decoupling entry, and the companion-paper tags are unpacked in the *Reading conventions* paragraph following Figure 31.

Reading conventions for Figures 32 and 31 and Table 41. In the provenance map (Figure 32) the colors are *not* cosmetic: every blue observable inherits the open status of its mechanism, and Conjecture VIII closes as the *epoch identity* ($\Omega_m = 1/\pi$ exact at $t_0 - 75$ Myr; residual lemma VIII-a exact to 10^{-15}), so the matter-density column is conditional on that single declared input. Only $\Delta N_{\rm eff} \leq 0.03$, the derived null of row D12, is genuinely Derived; the tilt n_s, the tensor ratio r and the scale $V^{1/4}$ are plateau-derived (D*, Lemma IV revised), and σ_8 and S_8 are computed from the amplitude (D*, Lemma V′), with A_s the launch sector's declared input. The lithium-7 abundance inherits the same closed-packet $1/\pi$: Conjecture X is its $A = 7$ instance. Companion-paper tags (A2-GG6, A6-GRIP, A9-SM, C1-Kaxion, C2-Hubble) mark where each result is developed in full. In Table 41, the columns are: the GG prediction in closed form; the observed value with its 1σ uncertainty; the relative deviation Dev = (pred − obs)/|obs|; the *pull* (pred − obs)/σ with $\sigma = \sqrt{\sigma_{\rm obs}^2 + \sigma_{\rm th}^2}$, the theory tolerance $\sigma_{\rm th}$ printed with its source in the pull cell of every row that carries one (exact closed forms carry none); and the provenance triage (**D** = Derived; **D*** = Derived-conditional, closed modulo a named conjecture or lemma; **A** = Asserted) with the marker type. Hubble constants are in km/s/Mpc; the background $H_0^{(0)} = 66.8 \pm 0.2$ is derived with zero sky input by the geometric Route I, its coefficient $\zeta_\Lambda = \pi/(LK)$ codebook-selected and pre-registered, and independently at 66.6 ± 0.2 from the π-locked inventory read against the measured densities (Section 29.8) — and scored against the CMB reading 67.4 (the coherent π-offset, Section 40.4); the local-ladder entry scores the same background against the H0DN consensus 73.50 ± 0.81, its −9.1% deviation the predicted ladder systematic (Section 42.6). The four acoustic rows are CLASS-verified images of the same fractions (correlated, not additional freedom); anchored entries sit at Dev = 0 by construction; η_B carries an order-of-magnitude transport band; ℓ_D is definition-dependent at the percent level. The z_* row scores the framework's audited atomic-chain decoupling prediction (1091.5, raw, zero fitted parameters, at the framework's own h) against the measured value, landing at the 0.15% level, with the topological corridor switch e^7 — a distinct quantity sitting 0.4% above it in e-folds (§42.1) — annotated in the plot, not scored. In the plot, the inner thick bar is the empirical $\sigma_{\rm obs}/|{\rm obs}|$ and the outer thin bar the combined $\sqrt{\sigma_{\rm obs}^2 + \sigma_{\rm th}^2}/|{\rm obs}|$, with entries carrying a theory tolerance $\sigma_{\rm th}$ annotated inside the panel; the radiation entry is shown as the total $N_{\rm eff}$. *Epoch calibration* (the 'cal.' entries in the table; violet diamonds in the plot): under Planck central values the exact crossing $\Omega_m = 1/\pi$ occurred 75 Myr before the present (±160 Myr, i.e. 0.5σ; §32.11), and all matter fractions (Ω_m, Ω_b, $\Omega_{\rm dm}$) decrease with time while Ω_Λ grows — so, using that single declared input once, the four inventory deviations collapse to (input, 0.00, +1.69, −0.31)%, every remaining inventory residual an epoch-independent ratio within about 1σ; the rows outside the inventory carry their own pulls, the distance ladder's −8 and $100\,\theta_*$'s −2.4 among them, and are not smoothed by this sentence. The clustering of two predictions (w and Ω_b) near +2.7% is examined in Section 84.5; an earlier version counted $S_8^{\rm late}$ as a third, and that row is withdrawn.

Counting summary. The twenty-one entries decompose as: one Derived (D, a derived null); nineteen Derived (cond.) (D*); one Asserted (A). Ten distinct named lemmas appear: (i) Kaxion-misalignment-completeness, (ii) stratum-assignment of the vacuum residue (formerly the two-stage corridor), (iii) compact-angle-mass-uniqueness, (iv) plateau-class-and-prompt-crossing (Lemma IV,

revised), (v) GDOS-heat-kernel-amplitude (Lemma V′, replacing the retired TeV-scale-inflation-uniqueness), (vi) resonant-leptogenesis-normalization, (vii) BFFT-damping-knee-identification, (viii) closed-packet-integration, (ix) vacuum-residual-exactness, (x) lithium-port-π-suppression. Each lemma is an open research target with a precise mathematical statement — the Roman numeral is its entry in the conjecture registry of Section 84.4, so "Lemma IV (revised)" and "Conjecture IV (revised)" name the same statement in its derivational and epistemic roles; once any one is proved, one or more table entries are promoted from D* or A to D. The status of the cosmological codebook is therefore: one entry fully closed, nineteen awaiting nine named lemmas, and one — the scalar amplitude, the launch sector's declared input — awaiting the single heat-kernel computation of Lemma V′. No entry requires an adjustable continuous parameter.

Post-hoc-fit defense. The reader who arrives at the π-locked inventory line for the first time will rightly ask whether $\Omega_m = 1/\pi$ is a post-hoc fit to the Planck data. The defense is structural and follows the same logic used in GG-A9-SM for the Standard-Model parameter sheet. First, the value $1/\pi$ for the matter fraction is forced by the closed-packet topology of the cosmological-boundary stratum, which is itself forced by the OGN-5 integer ledger of GG-A4-Euclid; the value is not chosen. Second, the same single topology forces *three* entries (Ω_m, Ω_Λ, Ω_b) simultaneously; a post-hoc fit explains one number, not three coupled numbers. Third, the Ω_b entry depends on $2\pi^2$, not on π; the closed-packet topology must produce the right *combination* of powers of π, and it does. Fourth, every other entry in the codebook is independently constrained by the one-background theorem, the inherited neutrino-and-CP ledger, or a named lemma; none of them is a free parameter that could have been tuned. The post-hoc-fit defense will be developed in full in the per-parameter audit below, drawing on the parallel defense in GG-A9-SM.

What the codebook does not yet contain. Three classes of entries remain outside the present codebook and will be added as the program matures. First, the feature-dependent supplementary channels (the primordial-black-hole mass-window slice; the primordial-gravitational-wave band in the Starobinsky channel) appear in Part VII but are not codebook entries because they are not single-parameter observables. Second, the remaining BBN light-element abundances (D/H, ^{3}He/H, ^{4}He/H) appear in Part V but are not codebook entries because they are coupled-output observables of a small set of inputs; the lithium-7 abundance is the exception, promoted to a codebook entry above because its closed-packet $1/\pi$ suppression is a single-number, parameter-free prediction. Third, the early-galaxy abundance statistics relevant to the JWST audit appear in Part VI and are coupled to the computed amplitude rather than to a separate parameter. The codebook as it stands is the headline parameter sheet; the coupled-observable audits live in the respective Parts and the consolidated benchmark audit of Appendix Q.

68.2 Per-parameter MSMF/OGN audit

This Section walks each row of Table 40 in turn and states, for each entry, the upstream-theorem chain that closes it (for the Derived entries) or the named lemma whose proof would close it (for the Derived (cond.) and Asserted entries). The principle that selects, for each entry, the unique

constitutional cost-minimum is the MSMF/OGN principle, first formalized in GG-A4-Euclid and applied parameter-by-parameter in GG-A9-SM for the Standard-Model sector. The present audit applies the same principle to the cosmological-boundary stratum.

68.2.1 Statement of the MSMF/OGN principle on the cosmological-boundary stratum

The MSMF/OGN principle in cosmology takes the same operational form as in the particle sector. On any stratum of the boundary moduli space, the constitutional cost functional $\mathcal{F} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ (the GFF ledger of GG-A6-GRIP) defines a finite-codimension critical manifold whose admissible saddles are the candidate Geometric Attractors. The MSMF/OGN principle then asserts:

Principle 68.1 (MSMF/OGN on the cosmological-boundary stratum)**.** On the cosmological-boundary stratum of the boundary moduli space, the constitutional cost functional $\mathcal{F} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ admits a unique cost-minimum saddle on every admissible sub-stratum corresponding to a cosmological observable. Equivalently: for each admissible cosmological observable, there is one and only one solution to $\nabla\mathcal{F} = 0$ on the relevant sub-stratum that satisfies the admissibility constraints. The uniqueness claim is the empirical content of the principle; it is in general non-trivial and is the open research target for those observables not yet fully closed.

The principle decomposes the cosmological observables into three operational closure modes, depending on the present state of evidence for the uniqueness claim:

(D) *Derived.* The uniqueness claim is proved: the cost-minimum on the relevant sub-stratum is unique, and the value is obtained from upstream theorems of the Ultra Eight GG-Theory cohort via a finite citation chain. No empirical content is required to obtain the value.

(D*) *Derived (cond.).* The uniqueness claim is conditional on an explicitly-named lemma about the sub-stratum's geometry (one lemma for seventeen entries; D4 adds the companion's pinning conjecture and A15 a second lemma). The lemma is an open research target with a precise mathematical statement; once proved, the entry is promoted to D.

(A) *Asserted.* The observable is reported as the empirical input pending a structurally-named lemma whose proof would identify the underlying cost-minimum on the relevant sub-stratum and thereby promote the entry to D* or D.

Thus the principle is a falsifiable uniqueness claim (the cost-functional saddle is unique on each admissible sub-stratum), and the three closure modes are operational summaries of the state of evidence for that claim.

The discipline that follows from Principle 68.1 is the zero-free-parameters discipline of GG-A4-Euclid: at no point in the audit is an adjustable continuous parameter introduced. The input ledger consists of the Geometric Universal Five Constants GU-5 $= \{c, G_N, \hbar; k_B, Q_B\}$, the structural constants (π, e), and the proved OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. Every cosmological observable is either a constructive function of these inputs (closure modes D and D*) or an honest empirical input awaiting such a function (closure mode A).

68.2.2 Audit of the one Derived (D) entry, and of the tensor row it formerly held

D9. Tensor-to-scalar ratio r (now D*; audited here for continuity). The derived plateau class (the GDOS non-conformal route of Part IV) fixes the tensor sector through the Starobinsky relation $r = 12/N^{*2}$ (Starobinsky, 1980); evaluated at the reheating-matched count $N^* = 52.6$, it gives $r = 0.0043$ at leading order and 0.0038 exact — a factor nine below the current observational bound ($r < 0.036$ at 95 % CL) and several times above the projected Simons-Observatory/LiteBIRD floor, hence decisive within the decade. The class and the count are underwritten by the revised Lemma IV; the relation $12/N^{*2}$ itself carries no further lemma, and the row is D* because the count it is evaluated at is D*: the entry is the Starobinsky relation, evaluated at the derived-conditional count, and its exact value on the plateau is 0.0038. Citation chain: GG-A5-GDOS (GDOS class) → Lemma IV (revised) → (Starobinsky, 1980) → present.

D12. Effective neutrino number deviation $\Delta N_{\rm eff}$. The row is a derived null. Below the lock scale the field content is the Standard Model's: the boundary-condition termination of $U(1)_B$ in Arisaka (2026, A3-BI6) leaves no light vector, the Kaxion of Arisaka (2026, C1-Kaxion) is produced by vacuum misalignment as a coherent zero-momentum condensate, not as a thermal population, and the neutrino block of Arisaka (2026, A9-SM) is three Standard-Model species with $\Sigma m_\nu \simeq 59$ meV. The statement to check is not that nothing light exists — an 8.1 μeV particle is of course relativistic at $T \sim 1$ MeV — but that nothing light was ever *populated* thermally. The channel that could have populated the Kaxion is its coupling to the gluon field, not to photons; freeze-in through that portal at this paper's own reheating temperature $T_{\rm rh} \simeq 6 \times 10^{10}$ GeV produces an abundance of order 0.1–1 of equilibrium, and one real scalar that thermalized and decoupled at $g_{*s} = 106.75$ carries $\Delta N_{\rm eff} = (4/7)(10.75/106.75)^{4/3} = 0.027$. So the honest statement is a bound, $\Delta N_{\rm eff} \leq 0.03$: $N_{\rm eff} = 3.044$, the Standard Model's own value (non-instantaneous decoupling and QED corrections included; Akita and Yamaguchi, 2020), plus a freeze-in residue between 0.003 and 0.027 according to an order-one coefficient in the gluon-portal rate that this paper does not compute, the ceiling six times below the measurement's precision. (Earlier versions of this row said "zero exactly", then "zero to within 0.01"; the first argued from the photon coupling and the second understated the ceiling by a factor of three. The row remains D — a derived null is derived whether its residue is zero or 0.03 — and the two-sided kill below is unaffected.) Against *Planck* 2018, $N_{\rm eff} = 2.99 \pm 0.17$, the null sits at $+0.3\sigma$; against ACT DR6 with *Planck* and BAO, 2.86 ± 0.13, at $+1.4\sigma$. It is two-sided: any $\Delta N_{\rm eff} \gtrsim 0.2$ falsifies the termination. (An earlier version of this row carried $\Delta N_{\rm eff} = 0.04$ as "a small thermal residue from late-decoupling species in the inherited particle ledger"; the particle paper contains no such residue, and the 0.04 was the Standard Model's own $3.044 - 3$ counted a second time.) No remaining lemma. Citation chain: GG-A3-BI6 → GG-A9-SM → present.

Counting summary. The one D entry is the effective-neutrino-number deviation, $\Delta N_{\rm eff} \leq 0.03$: a derived null, because the field content below the lock scale is the Standard Model's (GG-A3-BI6's boundary-condition termination leaves no light vector, and the Kaxion is populated only by gluon-portal freeze-in, at 0.1–1 of equilibrium, whose ceiling 0.027 is the bound's content; at that ceiling the sound horizon shortens by 0.1% and the CMB-read rate rises by 0.06 km s^{-1} Mpc^{-1}, inside

every comparison made here). Its citation chain is GG-A3-BI6 → GG-A9-SM (the neutrino block) → present, and it references no empirical fit. The tensor ratio, formerly listed here, rides on the same reheating-matched count as the tilt and carries the tilt's grade.

68.2.3 Audit of the nineteen Derived (cond.) (D*) entries — nine named lemmas

Thirteen of the nineteen D* entries are organized under seven named lemmas (I, II, III, VI, VIII, IX, X) and audited in this block; the remaining six — the tilt (A7), the tensor ratio (D9) and the inflationary scale (A15) on the revised Lemma IV, and σ_8 (D10) and S_8 against its two comparators (D11, D19) on Lemma V′, the amplitude they descend from — are audited with the inflationary sector below. Seven of the eight lemmas here close codebook rows directly, while Lemma VII gates only the interpretive BFFT overlay of the damping row D18, whose codebook value is closed by Lemma VIII. Each lemma is a precise mathematical statement about a sub-stratum of the boundary moduli space; each is an open research target.

Lemma I: Kaxion-misalignment-completeness. (Of the ten conjectures, the six with a closed formal statement — I, II, III, IV, V′ and X — are boxed below; VI–IX are stated in prose, at the precision their content allows.) This lemma closes D4 (dark-matter density fraction $\Omega_{\rm DM}$).

Conjecture 68.2 (Kaxion-misalignment-completeness)**.** On the compact-angle sub-stratum of the cosmological-boundary stratum, the misalignment-mechanism relic-abundance functional, evaluated at the single-patch angle the calibrated closure returns ($\theta_i = 1.18 \pm 0.04$ rad, Part V), equals $\Omega_m - \Omega_b = 1/\pi - 1/(2\pi^2)$.

The angle is not selected by the cost functional $\mathcal{F}$; it is the read-back of Part V (Appendix M), and the lemma's content is the completeness of the misalignment channel at that angle. The empirical value $\Omega_m - \Omega_b = 0.2657 \pm 0.0070$ (Planck 2018; the cold component alone, $\Omega_c = 0.2645$, excludes the neutrino share) is in 0.27-σ agreement with the predicted value $1/\pi - 1/(2\pi^2) \approx 0.2676$. The lemma presupposes the single-patch history of the compact angle — one misalignment value for the observable universe — which is itself a conjecture, the inflation-era pinning C-Kaxion-d of GG-C1-Kaxion (Part V, §35.5.2). That conjecture is necessary, not one of two readings: without it the angle is randomized by de Sitter diffusion during primary inflation, and because the compact angle has no symmetry restoration afterward the randomization sits on super-horizon scales, depositing a cold-dark-matter isocurvature that exceeds the *Planck* bound by ten orders of magnitude in power at f_K and by eight at the bulk scale, where the angle is quiet but not quiet enough; no decay constant below 5×10^{17} GeV evades it. If C-Kaxion-d fails, the Kaxion route to this row fails with it and the share would have to be carried by something else. Once Conjecture 68.2 is proved, with C-Kaxion-d, row D4 is promoted to D. Upstream citation chain: GG-A2-GG6 (compact holonomy) → GG-A3-BI6 (BI-6 anomaly freedom on the compact angle) → (Peccei and Quinn, 1977; Weinberg, 1978; Wilczek, 1978) (axion template) → the present audit.

Lemma II: Stratum-assignment of the vacuum residue (formerly Hubble-two-stage-corridor-uniqueness). This lemma closes D5 ($H_0^{\rm CMB}$) and D6 ($H_0^{\rm late}$): the one background rate,

compared against two probes. (Earlier versions of this codebook carried here a two-stage corridor conjecture that also gated σ_8 and the two S_8 branches; the corridor's switch $z_* = e^7$ enters neither of its branch functions, its late branch required an unobserved local underdensity, and the clustering rows are now computed from the declared amplitude, §43.3. The corridor conjecture is retired; the theorem-form statement of Section 68.3 is retained for its proved part, the switch (C1), and its two-branch reading is withdrawn there.)

Conjecture 68.3 (Stratum-assignment of the vacuum residue)**.** The vacuum residue ζ_Λ of the closed-form rate $H_0 = M_{\mathrm{Pl,red}}^{-1}(M_{\mathrm{lock}}^2/M_{\mathrm{Pl,red}})^2\sqrt{\zeta_\Lambda/3\Omega_\Lambda}$ (Section 29.8) is the Wilson-line (Hosotani) zero-mode residue on S^1_θ, and on that stratum the only named constituents are the seam factor 2, the L-holonomy angle $2\pi/L$ and the dual-lock count K; under MSMF/OGN the cost-minimal monomial in those constituents that lands in the physical window is unique, $\zeta_\Lambda = \pi/(LK) = \pi/175$, giving $H_0 = 66.8 \pm 0.2$ km/s/Mpc on every probe that reads the background. No continuous parameter is free to be tuned.

Enumerated on the stratum alphabet at cost ≤ 4, exactly one monomial admitted by the residue rule lands in 60–80 km/s/Mpc (one excluded by it, $1/((2\pi/L)^2K)$, lands at 67.0); the next, $(2\pi/L)^2/(2K)$, costs one unit more and gives 74.8, so no stratum monomial at cost ≤ 5 returns the ladder's 73.5 (Section 29.8). Two things are now fixed about what proving it requires, and the second is stated against the paper's own convenience. *(a) The holonomy is not selected here.* The conjecture evaluates the residue at the ledger's holonomy, $M = 1$ in the lock-scale units, and the reading of the five sectors as labels on that stratum rests on the holonomy being *fixed*. In a flat-circle model the one-loop potential of the endpoint quarks does not fix it: its minimum sits at $M = 7, 8$, tied, for periodic quarks and at $M = 0$ for antiperiodic ones. So the ledger's value stands on GG-A3-BI6's clause (C9) — the circle holonomy as background data on the admissible class, registered there as an exposure — and rows D5 and D6 rest on that clause (Arisaka, 2026, A3-BI6). A Hosotani computation is not what proves the conjecture; the residue computation at the given holonomy is. *(b) What the rows move on.* They move on the residue computation of GG-A5-GDOS O6 — its part (b), the residue at the given holonomy, computed in the warped measure — against this paper's window, and on nothing else; in particular not on the interval-line period that the twist-rate reading needs (§29.8, the residue rule; GG-A3-BI6's O16). Once Conjecture 68.3 is proved — the computation of GG-A5-GDOS O6 returning $\pi/175$ within $\pm 0.6\%$, falsifier F-Hubble-11 — rows D5 and D6 are promoted to D; the Hubble tension is then not a tension of this framework but a prediction that the distance ladder carries a systematic, tested by the ladder-free late probes. Upstream citation chain: GG-A2-GG6 (bulk; the lock scale) → GG-A5-GDOS (the warp overlap and the Wilson-line stratum) → (Casertano et al., 2026; Planck Collaboration, 2018) (comparators, not inputs) → the present audit.

Lemma III: Compact-angle-mass-uniqueness. This lemma closes D13 (m_{Kaxion}) and D14 (f_{Kaxion}).

Conjecture 68.4 (Compact-angle-mass-uniqueness)**.** On the compact-angle sub-stratum, the OGN-5 integer ledger ($K = 35$) acts on the compact-angle pseudo-scalar such that the unique cost-minimum solution is a single physical decay constant $f_{\mathrm{Kaxion}} = f_{\mathrm{bulk}}/\Xi_{\mathrm{cap}}$, with the cap loading an

integer of the additive form $\Xi_{\rm cap} = K + \Sigma \in \{35, 38, 40, 42, 43, 45, 47\}$ (benchmark $43 = K + N + L$) and the bulk normalization $f_{\rm bulk}$ fixed by the geometry — so that the calibrated value 3.04×10^{13} GeV, $f_{\rm Kaxion} = 7.1 \times 10^{11}$ GeV and hence $m_{\rm Kaxion} = 8.1\,\mu$eV ($\nu_K = 1.95$ GHz) become derived rather than calibrated. The misalignment angle is then read back from the π-locked abundance, and the ladder's exact rung is the remaining discrete unknown.

The integer $K = 35$ appears in the constitutional input set as one of the five proved OGN-5 integers $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ of GG-A4-Euclid, and it enters the Kaxion sector through the ladder $\Xi_{\rm cap} = K + \Sigma$. The lemma is stated as the cost-minimum statement it would need to be, and it is open in two named places, both recorded in the companion paper GG-C1-Kaxion (Arisaka, 2026, C1-Kaxion): the bulk normalization is at present *calibrated* through the relic closure at the benchmark loading, not derived — the orientation convention of the anti-self-dual tensor sector that would derive it is GG-A3-BI6's open item O14 — and the additive ladder is an *ansatz* for which no defining computation of the cap loading has been formulated. What the companion does close is everything after those two: exact anomaly matching fixes the photon coupling on the DFSZ line ($E/N = 8/3$), the mass is the lattice-QCD image of the calibrated scale to $\pm 1.3\%$, and the window $[1.59, 2.13]$ GHz is the ladder's residual with the benchmark rung marked. An earlier version of this paragraph said the geometry fixes f_K to $\pm 6\%$; the $\pm 6\%$ was the abundance pipeline's own uncertainty propagated into the scale it calibrates, and it is retired. Upstream citation chain: GG-A4-Euclid (OGN-5 ledger) → GG-A2-GG6 (compact holonomy) → GG-A3-BI6 (the tensor's dual; O13/O14) → GG-C1-Kaxion → the present audit.

Lemma VI: Resonant-leptogenesis-normalization (formerly sphaleron-transport-normalization). This lemma closes D17 (baryon-to-photon ratio η_B). The baryon asymmetry is produced by resonant leptogenesis on the inherited GG-A9-SM neutrino-and-CP ledger: the heavy Majorana sector is the ultraviolet completion of the GG-A9-SM Weinberg dimension-five operator at the lock scale, its Yukawa couplings fixed by seesaw matching to the GG-A9-SM light-neutrino masses and its near-degeneracy fixed by the OGN-5 integer $K = 35$, so that the resonant CP asymmetry and the washout efficiency combine to the geometric throughput $\varepsilon^{\rm res}\,\kappa = e^{-K/2}$ (full derivation in Part V and Appendix I). The lemma is the conjecture that the *partial-resonance condition* — the single named input, $\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$ at finite temperature (Appendix I; the exact-peak form $|M_i - M_j| \simeq \Gamma/2$ overshoots the required throughput by the same factor and is *not* what is assumed) — is met at the lock scale, so that the full Boltzmann-transport normalization would converge to the geometric value with its residual $O(1)$ coefficients fixed. It is stated as a conjecture and not as a fact because the vacuum numbers do not meet it, and because the paper's own time-resolved integration under the same conditions lands a factor of about eight below (§I.6): the derived splitting and width miss the condition by many orders of magnitude (GG-A9-SM puts the vacuum ratio at $\Delta M/\Gamma_N \simeq 6.8 \times 10^{11}$ and marks $\varepsilon^{\rm res}$ open and the throughput adopted), and the finite-temperature width that would close the gap is the open computation named in Part V. The geometric value 4.89×10^{-10} sits -20% from the empirical $\eta_B = (6.12 \pm 0.04) \times 10^{-10}$ (Planck Collaboration, 2018), inside the declared order-of-magnitude band. Once the lemma is proved, row D17 is promoted to D only if the time-resolved integration under its condition, now a factor of about

eight low, reaches the target. Citation chain: GG-A4-Euclid ($K = 35$) → GG-A9-SM (Weinberg operator + CP ledger) → Part V → present audit.

Lemma VII: BFFT-damping-knee-identification. This lemma no longer closes a codebook row; it gates the *interpretive BFFT overlay* of D18 (damping-tail multipole ℓ_D). The codebook value of D18 is the Boltzmann diffusion prediction of the π-locked inventory, $\ell_D = 1350 \pm 3$ (+0.4% against the ΛCDM-inferred 1345 ± 15), and is closed — with the fractions it descends from — by Lemma VIII. The separately proposed locked relation $\ell_D = K(2\pi)^2 \approx 1382$ identifies the ultraviolet knee with the first non-trivial BFFT cutoff of the multiscale throat (Part VI); it sits +2.7% above the ΛCDM-inferred value and +2.3% above the framework's own Boltzmann value, so the lemma must now *explain that gap*, not merely assert that the projection of the diffusion scale onto the BFFT bus is exact at leading order. If proved in that stronger form, the overlay is promoted from conjecture to structure; if refuted, the overlay falls away while D18 stands on the diffusion integral. Citation chain: GG-A4-Euclid ($K = 35$) → GG-A7-Qi4 (BFFT bus) → Part VI Boltzmann audit → present audit.

Lemma VIII: Closed-packet-integration. This lemma closes D1–D3 (Ω_m, Ω_Λ, Ω_b) and, through them, D18 (ℓ_D, whose diffusion-integral value has no input beyond the fractions and standard recombination physics) and D21 (the decoupling redshift z_*, whose atomic-chain value likewise has no input beyond the fractions and standard atomic physics; §42.1). It asserts that the closed-packet integral of the BI-6 four-form on the cosmological-boundary stratum evaluates to $1/\pi$ (matter share) and $1/(2\pi^2)$ (baryonic weight, $1/\mathrm{Vol}(S^3)$) with an exhibited, verifiable normalization — completing the programmatic Steps 3–4 of Theorem 74.1 (see the Proof-status note there). The π-locked values are D* entries of half-percent empirical precision: coupled form derived, epoch declared (§32.11), residual lemma VIII-a exact up to the $\sim e^{-K}$ floor.

The VIII-a/VIII-b split. The conjecture decomposes into two parts of different status. *VIII-a (delocalization and share robustness of the stationary measure):* on the matter-sector phase moduli — the electroweak order-parameter directions (the Wilson–Hosotani phases of GG-A6-GRIP) being non-flat by design, settled into the GA, and excluded from the domain — the frozen stationary GDOS measure leaves every counted direction delocalized — the flavor directions *exactly* protected by rephasing invariance, the physical (Jarlskog-class) combinations tilted at most $\sim 10^{-28}$, the leading term the gravity-mediated floor $\sim e^{-K} \approx 10^{-15}$ (the mechanism home of the original candidate $S_w = K = 35$; its $O(1)$ prefactor the one open coefficient) — and the share dictionary consumes the measure only through counting and per-component fiber volume, both invariant under any normalized tilt — uniformity is sufficient but not necessary. Named as the GSM adjunct of GG-A5-GDOS (O6) and grounded in the topological, non-spatial character of S^1_θ; proved in this consumer form (frozen-Boltzmann reduction via the [B1] kernel theorems), with its three residues discharged: the cascade-decoupling check (verified — the cascade's boundary imprint is the parameter codebook and nothing else), the freeze-out/electroweak ordering (a two-channel theorem: the electroweak order parameter is driven by the ungapped light channel, stays in equilibrium through the crossover, and localizes there — so the counting {matter phases in, order parameter out} is stable for every admissible epoch hierarchy), and the invariant-operator tilt table (which restores

the gravity floor as the operative bound). *VIII-b (the share dictionary): which* measure-ratio equals *which* Ω — precisely the unexhibited Step-4 normalization of Theorem 74.1. The holonomy moduli space carries no discrete identification (GG-A2-GG6, the no-identification remark), so no volume-π shortcut exists and every route to $\Omega_m = 1/\pi$ must pass through VIII-b. *Lemma VIII-b′ (proved) — the real-sector half-measure (GG-A5-GDOS).* Exact mode counting — $2\Lambda+1$ independent real degrees of freedom for a reality-constrained sector against $2(2\Lambda+1)$ for a holonomy-charged one, a ratio of exactly $1/2$ at every finite cutoff — fixes the effective fiber measures at π (real) and 2π (charged) under the equipartition measure of VIII-a. The recalibrated share dictionary then delivers the full inventory from a single remaining normalization: unit-trace total = the real-sector measure π; matter = one quantized closed packet; baryons = the charged average $1/2\pi$ within matter — forcing $\Omega_b/\Omega_m = 1/2\pi$ with no residual freedom, the same h-free ratio verified in the acoustic spectrum (Table 23). The open core of Conjecture VIII has therefore contracted to one statement, *VIII-b″ (unit-packet normalization)*: one closed packet = one unit of fiber measure, with its derivation route through the $T_{\mathsf{GFF}} \sim M_{\mathrm{KK}}$ unit bridge of GFF-ST6. Until VIII-b″ is exhibited the fractions remain D*; when it is, Steps 3–4 of Theorem 74.1 close and the π-locked inventory is promoted. *The unit bridge: VIII-b″ sharpened and partially executed.* The route is now the explicit relation $T_{\mathsf{GFF}} = 2\pi\, I_m\, \varepsilon_{\mathsf{GEP}}$ (one packet of integer charge I_m, energy $I_m\varepsilon_{\mathsf{GEP}}$, against the equipartition rate $T_{\mathsf{GFF}}/2\pi$ per unit of fiber measure), and GG-A6-GRIP (Arisaka, 2026, A6-GRIP) fixes its *action* side outright: the descent normalization is $S_{\mathrm{bulk}} = (\pi/2)\,\mathcal{F}_{\mathsf{GEP}}$ (its Lemma 21.4 — transverse flux $\Phi_\omega = \pi$, pinned by the graviton-normalizing warped volume of GG-A5-GDOS Lemma T8-1 and *derived from* the Bogomolnyi cascade actions πk — the saddle charge $\Delta c_2 = k$ is the Whitney pair count of the winding bundle, under three declared premises; GG-A6-GRIP), and the closed packet is *absolutely superselected* (its Corollary 21.6: only $\Delta\mathcal{F}_{\mathsf{GEP}} \in 4\mathbb{Z}$ is Dirac-phase-trivial) — a topological mechanism beneath the DE-1 freezing. By the Hirzebruch signature theorem ($p_1 = 3\sigma$), $p_1 = N = 3$ is the *minimal* closed gravitational topology ($\sigma = 1$), so the packet is the canonical unit of the gravity-sector charge lattice: the "one" in "one closed packet" is forced, not read. The open remainder has since contracted by one item. *[B1] discharged as a structure theorem.* The Einstein ratio $T_{\mathsf{GFF}} = D/\mu$ on the fiber-moduli block (its GKLS generator already fixed by GG-A5-GDOS's reduction) is a *thermometer*: under kinetic (KMS) contact it copies the bath temperature exactly, independently of the coupling strength and the spectral details, so no kernel evaluation can manufacture M_{lock}. At the Geometric Attractor the tower bath is gapped and frozen, and the block is dissipatively disconnected ($D = \mu = 0$) — the kernel-side face of the DE-1 freezing (a tower-channel statement: the ungapped light channel instead retires the electroweak order parameter at the crossover, per the refinement paragraph above). The lock-scale reading survives in exactly one form: the *freeze-out record* $T^* = 2M_{\mathrm{lock}}/L^*$, the last temperature the block registered before the pair channel's Boltzmann factor $e^{-2M_{\mathrm{lock}}/T}$ severed contact (L^* the logarithm of the epoch rate hierarchy) — the Tolman softness of GFF-ST6 reduced to a named logarithm (see the sharpened remark in GG-A6-GRIP and the GSM-adjunct update in GG-A5-GDOS). The open core is therefore *[P]* alone: the OGGEP projection of the superselected packet action ($3\pi/2$) onto the present-epoch boundary Hamiltonian, now with a sharpened target — $\varepsilon_{\mathsf{GEP}} = T_{\mathsf{GFF}}/2\pi = M_{\mathrm{lock}}/(\pi L^*)$ under the freeze-out reading (canonical $I_m = 1$), against the exponentially smaller Geometric-Attractor vacuum-floor reading ($T \to \omega_0/2$, the equipartition floor

of the winding-suppressed modulus) — whose adjudication is the projection lemma's own content. *[P] discharged by identification; the promotion path retired.* The projection terminates on a *dilution no-go*: superselection freezes the integer I_m, not the energy record, so U_{packet} tracks the diluting matter sector ($\propto a^{-3}$) while the equipartition rate $T_{\mathsf{GFF}}/2\pi$ is frozen — no epoch-independent constant (in particular no anomaly-inflow $\varepsilon_{\mathsf{GEP}}$) can equal a diluting quantity at more than one instant. The occupancy $u(t) = \pi\,\Omega_m(t)$ therefore sweeps monotonically through unity exactly once, at $a_\pi = 0.9949$ ($z_\pi = +0.005$, about 75 Myr before the present at the two-road $H_0 = 66.6$ (Route I 66.8); the Planck uncertainty on Ω_m is a ± 160 Myr timing blur, so this offset is 0.5σ — consistent with the crossing being now), and the closure statement $u = 1$ is *mathematically equivalent* to the present-epoch calibration $\Omega_{m,0} = 1/\pi$ — precisely the epoch input declared in §32.11, the one freedom the geometry does not remove. The unit-bridge arc and the form/scale/epoch honesty split thus converge on the same terminus by independent roads. *Consequently the open core of Conjecture VIII is VIII-a (uniformity) alone*: VIII-b′ is proved, and VIII-b″/($\star$) closes as the epoch identity — exact at the crossing epoch t_π, with $t_0 = t_\pi$ to 1% ($\approx$ 72 Myr) as the declared L4 input. The fractions remain D* with the epoch as their declared calibration input, exactly as the codebook tags them. *Residue status:* VIII-a's three residues are discharged (see the refinement paragraph above); the remaining named items in the entire Conjecture-VIII neighborhood are the 6D zero-mode formalization of the wrap-independence premise and the gravity-floor prefactor — both well-posed, neither load-bearing for any deployed claim.

D1. Matter density fraction $\Omega_m = 1/\pi$. The matter density fraction is the unique cost-minimum of $\mathcal{F}$ on the closed-packet topological stratum of the cosmological-boundary moduli space. The closure runs through the OGN-5 integer ledger of GG-A4-Euclid: the matter packet inherits a closed-form integer topology from the **5** trace structure of the bulk holonomy, and the corresponding cost-minimum on the cosmological slice reduces to the rational fraction $1/\pi$ by the closed-packet integration (the still-open Conjecture VIII). Remaining: VIII-a (uniformity; exact up to the $\sim e^{-K}$ floor) + the declared epoch input (§32.11). Citation chain: GG-A4-Euclid → GG-A2-GG6 → GG-A3-BI6 → present.

D2. Cosmological constant fraction $\Omega_\Lambda = 1 - 1/\pi$. The complementary fraction is forced by the unit-trace closure of the homogeneous-isotropic Friedmann constraint $\sum_i \Omega_i = 1$ on the cosmological-boundary stratum. Once $\Omega_m = 1/\pi$ is fixed by D1 and the radiation fraction is negligible at the present epoch, the residual vacuum fraction is identically $1 - 1/\pi$. The fact that the residual vacuum fraction takes the value of a topological complement rather than an adjustable empirical constant is the deep content of D2. Remaining: VIII-a (uniformity; exact up to the $\sim e^{-K}$ floor) + the declared epoch input (§32.11). Citation chain: GG-A6-GRIP (residual F_{GFF} interpretation) → present.

D3. Baryon density fraction $\Omega_b = 1/(2\pi^2)$. The baryon fraction descends from the baryonic-gauge protection of the Standard-Model ledger (Arisaka, 2026, A9-SM), projected onto the cosmological-boundary stratum. The factor $2\pi^2$ is the volume of the unit three-sphere, $\mathrm{Vol}(S^3) = 2\pi^2$: the baryonic gauge $U(1)_B$ projects onto the relevant topological sub-stratum with weight

$1/\mathrm{Vol}(S^3) = 1/(2\pi^2)$ — equivalently, and in the derived route of the recalibrated dictionary, the charged group average within matter: $1/(2\pi^2) = (1/\pi)\cdot(1/2\pi)$ (Lemma VIII-b′); the $\mathrm{Vol}(S^3)$ form is the mnemonic factorization of the same number. Remaining: VIII-a (uniformity; exact up to the $\sim e^{-K}$ floor) + the declared epoch input (§32.11). Citation chain: GG-A4-Euclid → GG-A9-SM → present.

D21. Decoupling redshift z_*. The entry carries two distinct quantities (§42.1). The corridor switch $z_{\mathrm{sw}} = e^J = e^7 = 1096.6$ is a topological marker forced by the OGN-5 integer $J = 7$ (Theorem 68.8); it is not the audited number. The audited number is the operational decoupling redshift: the framework's own atomic chain — Saha/Boltzmann recombination run on the π-locked inventory $\Omega_b = 1/(2\pi^2)$, $\Omega_m = 1/\pi$, with standard atomic constants and no cosmological input beyond the fractions — lands on the Planck value $z_{\mathrm{dec}} = 1089.92 \pm 0.25$ at the 0.1% level, 0.4% below the switch. (The Dev plot and its data table score this audited atomic number, with the switch annotated, not scored.) The row therefore rides on the same lemma as the fractions it descends from: once Lemma VIII is proved, D21 follows the fractions — to D with the epoch input derived elsewhere, D* while the epoch stays a calibration, as the roadmap's row says (the atomic chain itself introduces no new lemma). Citation chain: GG-A4-Euclid ($J = 7$) → D1–D3 (the fractions) → §42.1 → present.

Lemma IX: Vacuum-residual-exactness. This lemma closes D16 (w), and its content is not an open assertion: the *exactness* of $p = -\rho$ is Theorem 29.1 (T-Cosmos-DE′), proved from Lemma DE-1 of GG-A6-GRIP and the GG-A5-GDOS spectral gap, conditional on the constitutional rigidity of the bulk profile, with $|1 + w| \lesssim 10^{-91}$. What the lemma now names is the single surviving residue: the dynamical stabilization of that profile (a Goldberger–Wise-type computation, an upstream GG-A2-GG6/GG-A5-GDOS target). Once supplied, row D16 is promoted to D. The empirical determination $w = -1.028 \pm 0.032$ (Planck Collaboration, 2018) is consistent at 0.9σ, and the two-sided kill switch of Theorem 29.1 applies: any confirmed $w \neq -1$, either sign, is fatal.

Lemma X: Lithium-port-π-suppression. This lemma closes D20 (the primordial lithium-7 abundance $^7\mathrm{Li}/H$). It asserts that the $A = 7$ nucleosynthesis port on the cosmological-boundary stratum inherits the same closed-packet $1/\pi$ factor that fixes the matter fraction Ω_m, so that the standard-BBN lithium yield is suppressed by exactly $1/\pi$.

Conjecture 68.5 (Lithium-port-π-suppression)**.** On the $A = 7$ sub-stratum of the cosmological-boundary stratum, the open-system closed-packet integral suppresses the mass-7 yield by exactly $1/\pi$, so that $(^7\mathrm{Li}/H)_{\mathrm{GG}} = (1/\pi)\,(^7\mathrm{Li}/H)_{\mathrm{SBBN}}$, leaving the $A \leq 4$ abundances (D, ^{3}He, ^{4}He) unchanged.

With the standard-BBN value $(^7\mathrm{Li}/H)_{\mathrm{SBBN}} \simeq 5.0 \times 10^{-10}$ (Cyburt *et al.*, 2016; Pitrou *et al.*, 2018), the prediction is $(^7\mathrm{Li}/H)_{\mathrm{GG}} \simeq 1.6 \times 10^{-10}$, matching the metal-poor halo-star (Spite) plateau $(1.6 \pm 0.3) \times 10^{-10}$ (Fields, 2011) with no fitted strength: the identical $1/\pi$ that fixes Ω_m fixes the lithium deficit. The prediction is sharply falsifiable — any firm primordial ^{7}Li abundance away from SBBN/π excludes the port. Once Conjecture 68.5 is proved, row D20 is promoted to D. Upstream

citation chain: GG-A4-Euclid (closed-packet $1/\pi$) → GG-A2-GG6 (bulk holonomy) → Part V → the present audit.

Counting summary for D*. The nineteen D* entries are closed by nine lemmas: Lemma I closes one entry (D4, with the companion's pinning conjecture); Lemma II closes two entries (D5, D6); Lemma III closes two entries (D13, D14); Lemma IV (revised) closes two entries (A7, D9; A15 waits on Lemma V$'$ as well, audited with the inflationary sector below); Lemma V$'$ closes four D* entries (A15, D10, D11, D19) besides promoting the A entry (A8); Lemma VI closes one entry (D17); Lemma VIII closes five entries (D1, D2, D3, D18, D21); Lemma IX closes one entry (D16); Lemma X closes one entry (D20). Lemma VII closes no row: it gates the interpretive BFFT overlay of D18, whose codebook value is the Boltzmann diffusion prediction of the inventory. The economy is favorable: nine closing lemmas, nineteen entries promoted on proof of them.

68.2.4 Audit of the inflationary sector: two promoted entries and the one remaining Asserted (A) entry — two named lemmas

This block audits the three entries of the inflationary sector — the tilt (A7), the amplitude (A8), and the inflationary scale (A15). Two of them now stand at D* on the revised Lemma IV; one, the amplitude, remains the codebook's single A entry, carried by Lemma V$'$. Both lemmas concern the inflationary modulus on the cosmological-boundary stratum.

Lemma IV (revised): Plateau-class-and-prompt-crossing. This lemma underwrites A7 (n_s) and A15 ($V_{\rm inf}^{1/4}$) at D*; its statement follows.

Conjecture 68.6 (Plateau-class-and-prompt-crossing (revised))**.** On the inflationary-modulus sub-stratum of the cosmological-boundary stratum, (i) the GDOS non-conformal route forces the Starobinsky/ξ plateau class of Eq. (49) as the unique admissible potential (theorem-level promotion of the derivation of Part IV), and (ii) the GRIP instanton-ledger bound applies on Σ_r, so the lock crossing is prompt ($N_2 = 0$) and the reheating-matched count $N^* = 52.6$ is undiluted.

Lemma IV, in its revised form, underwrites the two entries the two-basin derivation moved out of A: the tilt (A7, now D*: $n_s = 0.9634$ exact, $1 - 2/N^* = 0.9620$ at leading order, derived in Part IV given (i) and (ii)) and the inflationary scale (A15, now D*: $V_{\rm inf}^{1/4} = 8.1 \times 10^{15}$ GeV, fixed by the derived class once the amplitude input is declared). Once Lemma IV is proved, rows A7 and D9 are promoted from D* to D; A15, whose content is the declared amplitude to the fourth root, waits on Lemma V$'$ as well. The earlier form of this conjecture (slow-roll-exit-uniqueness, which asserted that the exit fixes both n_s and A_s at their empirical values) is superseded: the tilt is now computed rather than matched, and the amplitude is carried separately by Lemma V$'$ below. Upstream citation chain: GG-A2-GG6 (bulk) → GG-A5-GDOS (GDOS anomaly structure) → GG-A6-GRIP (GRIP ledger) → (Starobinsky, 1980) (plateau template) → the present audit.

Lemma V (retired) and its replacement Lemma V$'$. The program's working codebook formerly carried here the conjecture *TeV-scale-inflation-uniqueness* — that $V_{\rm inf}^{1/4}$ is forced to a

TeV-scale value by structural matching to the throat sector, ruling out GUT-scale inflation. **That conjecture is refuted, and the refutation is recorded at full strength** (Part IV): a lock-scale plateau yields a *blue* spectral tilt ($n_s - 1 \sim +10^{-26}$) against the measured red -0.0351 ± 0.0042 — wrong in sign — and demands $\epsilon_V \sim 10^{-53}$ for the amplitude. No admissible reading survives; the conjecture is retired rather than reworded, and it is kept on this page so the record shows what the framework once believed and why it stopped. Its A15 role passes to the revised Lemma IV above.

In its slot stands the replacement that carries the sector's one remaining input:

Conjecture 68.7 (GDOS-heat-kernel amplitude (Lemma V′))**.** The R^2/scalaron coefficient of the boundary effective action, evaluated as the GDOS heat-kernel integral over the warped interval, equals $e^{2K/3}/12$ up to an order-unity factor — equivalently, the scalaron mass satisfies $M = M_5\, e^{-K/3}$ with $M_5 = (M_{\mathrm{Pl,red}}^2\, k)^{1/3}$, so that $A_s = 21.5\, e^{-2K/3}$.

Lemma V′ is the one computation that would pay the amplitude: its warp candidate currently reproduces the required scalaron mass to 87% (A_s to 75%) with no parameter, and two alternative routes (warped decay; holographic central-charge counting) have been attempted and closed. Once Lemma V′ is proved, row A8 is promoted from A to D and the launch sector becomes input-free. Upstream citation chain: GG-A2-GG6 (bulk) → GG-A5-GDOS (GDOS heat kernel) → the present audit.

Counting summary for the inflationary sector. Two of the three former A entries (A7, the tilt; A15, the inflationary scale) were promoted to D* by the two-basin derivation of Part IV and now ride on the revised Lemma IV. The single remaining A entry (A8, the amplitude) is closed by one lemma: Lemma V′. The economy is unchanged — two open research targets — and strictly better in content: three entries once pending are now one pending and two derived-conditional.

68.2.5 Closure summary and the promotion roadmap

The twenty-one codebook entries are closed by a total of *ten* named lemmas:

- Lemma I (Kaxion-misalignment-completeness) closes 1 entry.
- Lemma II (Stratum-assignment of the vacuum residue) closes 2 entries.
- Lemma III (Compact-angle-mass-uniqueness) closes 2 entries.
- Lemma IV (Plateau-class-and-prompt-crossing, revised) closes 2 entries (the tilt and the tensor ratio).
- Lemma V′ (GDOS-heat-kernel amplitude; replaces the retired Lemma V) closes 5 entries (the amplitude, the inflationary scale that is its fourth root, and the three clustering rows that scale with it).
- Lemma VI (Resonant-leptogenesis-normalization) closes 1 entry.
- Lemma VII (BFFT-damping-knee-identification) closes no entry directly; it gates the interpretive BFFT overlay of D18 (the row itself is closed by Lemma VIII).

- Lemma VIII (Closed-packet-integration) closes 5 entries.
- Lemma IX (Vacuum-residual-exactness) closes 1 entry.
- Lemma X (Lithium-port-π-suppression) closes 1 entry.

One entry (D12, the derived null) is already closed without any open lemma. The closure economy is therefore:

Table 42: Closure economy: how the codebook decomposition evolves as each named lemma is proved. The three right-hand columns count the observables that are Derived, Derived-conditional and Asserted at each stage. Reading down them shows what each named lemma would actually buy, which is a more useful question than whether the lemma is likely to be proved.

State of the codebook	**D**	**D***	**A**
Current state (present paper)	1	19	1
After Lemma I proved (with the companion's pinning conjecture C-Kaxion-d)	2	18	1
After Lemma II proved	4	16	1
After Lemma III proved	6	14	1
After Lemma IV proved	8	12	1
After Lemma V$'$ proved	13	8	0
After Lemma VI proved	14	7	0
After Lemma VII proved	14	7	0
After Lemma VIII proved (with the epoch input derived elsewhere; as a calibration it leaves the fractions D*)	19	2	0
After Lemma IX proved	20	1	0
After Lemma X proved	21	0	0

The promotion roadmap shows that the twenty-one-row codebook is closed by ten named lemmas (one of them, V$'$, the replacement of a retired and refuted predecessor), that five of the lemmas (II, III, IV, V$'$ and VIII) each close multiple entries, and that the closure proceeds through a strictly monotonic improvement of the codebook decomposition as each lemma is proved. When all ten lemmas are proved, all twenty-one entries are D — and, as the preview table's caption says, they are not twenty-one independent numbers but about eight, the rest being the same geometry read in other units.

Comparison with the GG-A9-SM closure economy. The Standard-Model codebook of GG-A9-SM has a similar structure: six entries fully derived, fourteen modulo a single named lemma, and eight remaining asserted (the eight non-top charged-fermion Yukawa masses). The cosmological codebook of the present paper is further from the goal state at the top and closer at the bottom:

one of twenty-one at D against six of twenty-eight, and only one of twenty-one at A against eight of twenty-eight. This reflects the fact that the cosmological-boundary stratum is structurally lower-dimensional than the Standard-Model stratum, and that the closed-packet topology of the cosmological inventory (the three π-locked fractions) admits more direct constitutional closure than the family-projector geometry of the Standard-Model parameter sheet.

After all, the audit is not a victory lap but a structured tally: which observables are closed today, which would be closed tomorrow under ten precise mathematical conjectures, and which remain honestly empirical pending the proof of those conjectures. The framework's strength lies not in claiming that every observable is already derived but in stating exactly which lemmas would complete the derivation, with the proofs offered as open research targets in language sharp enough to be falsified.

68.3 The Hubble two-stage corridor: canonical theorem statement

68.3.1 Empirical motivation: the H_0 tension as of 2026

The Hubble constant inferred from the cosmic microwave background under the standard ΛCDM extrapolation is $H_0^{\mathrm{CMB}} = 67.4 \pm 0.5$ km/s/Mpc (Planck Collaboration, 2018). The Hubble constant inferred from the local distance ladder (Cepheid-calibrated Type-Ia supernovae) is $H_0^{\mathrm{late}} = 73.50 \pm 0.81$ km/s/Mpc on the H0DN consensus (Casertano et al., 2026) (73.04 ± 1.04 for the single-team determination, Riess et al., 2022). The discrepancy stands at 6.4σ between the H0DN consensus and Planck and persists across the ladder's calibration cross-checks (Mira stars, the tip of the red-giant branch); the ladder-free probes — standard sirens, lensing time delays — are the undecided test. The minimal-extension responses within ΛCDM — early dark energy, modified recombination histories, extra relativistic degrees of freedom, interacting dark sectors — each introduce one or more adjustable continuous parameters and resolve the tension only in part. The present framework responds structurally, and we are careful to say at what level: it proposes that there is one background solution of the boundary Friedmann equation, read by every ladder-free probe, with a switch fixed by the inherited recombination geometry of GG-A3-BI6 and the integer $J = 7$ of the OGN-5 ledger of GG-A4-Euclid separating the two dictionaries through which it is read, and that the ladder's higher value is a measurement carrying a systematic. This is a structural identification, not yet a computed joint fit; a tension is fully resolved only by a global Einstein–Boltzmann likelihood across CMB, BAO, SN, and lensing with the covariance kept, which the paper names as open work (PS-Cosmos-16). What lifts it above a phenomenological "two H_0's" fit is one clean forward prediction a phenomenology cannot make: the gravitational-wave standard sirens, which read the background free of *both* the sound horizon and the distance ladder, must land on the *background* rate $H_0^{\mathrm{GW}} \approx 67$, not the local 73. That is the discriminating test, and it is near.

68.3.2 Setup and hypotheses

The theorem of this Section is proved under three constitutional hypotheses inherited from upstream papers of the Ultra Eight GG-Theory cohort.

(H1) Inherited bulk and π-locked inventory. The cosmological-boundary stratum descends

from the six-dimensional GG-6 bulk of GG-A2-GG6 via the OGGEP projection. The matter, vacuum, and baryon density fractions at the present epoch take the π-locked values $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, carried as Derived-conditional under Conjecture VIII in the codebook audit of Section 68.2.2.

(H2) BI-6 anomaly-free recombination geometry. The modified Bianchi identity BI-6 of GG-A3-BI6 is in force throughout the cosmological history. In particular, the recombination surface inherits a closed-form geometry from the **5** trace structure of the bulk holonomy: the sound horizon r_s at recombination and the angular diameter distance D_A from the present epoch to recombination are both functions of the inherited inventory $(\Omega_m, \Omega_\Lambda, \Omega_b, \Delta N_{\rm eff})$ with no continuous adjustable parameter.

(H3) GRIP cascade for the corridor switch. The runtime factorization GRIP = (DGI $\to$ GSSB $\to$ GA) of GG-A6-GRIP applies to the recombination event: the passage from the tightly-coupled DGI phase to the frozen GA phase is a GSSB crossing, and the change of dictionary at $z_* = e^7$ is the GRIP-Out that makes the landscape it crosses (§14.1, where the vocabulary is fixed). The cosmological history therefore admits a two-stage structure with a single switch event at recombination.

68.3.3 Statement of the theorem

Theorem 68.8 (Hubble two-stage corridor: existence (T-Cosmos-16))**.** *Assume hypotheses (H1), (H2), (H3) above. Then the boundary Friedmann equation admits a single background solution on the inherited inventory, with a GRIP-Out switch at the redshift z_* separating the two* dictionaries *through which that one background is read, and with the following three properties:*

(C1) Corridor switch is fixed by the OGN-5 integer ledger. *The switch redshift equals*

$$z_* = e^J = e^7 \approx 1097,$$

where $J = 7$ is one of the five proved OGN-5 integers $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ of GG-A4-Euclid. No continuous parameter is free to be tuned; the value is forced by the integer ledger, and the atomic decoupling computed from the framework's own constants, $z_{\rm dec} \approx 1092$, lies 0.4% below it (§42.1). The operational decoupling redshift from Planck is $z_{\rm dec}^{\rm Planck} = 1089.92 \pm 0.25$; the corridor switch $z_* = e^7 \approx 1096.6$ is a distinct topological quantity sitting 0.6% above that Planck value, while the atomic decoupling computed from the framework's own constants lands on the Planck value at the 0.1% level. Note: *the empirical Planck precision is $\sigma/z \approx 0.025\%$, so the switch–decoupling offset is large in units of that error — which is precisely why the two are carried as distinct quantities, never as one predicted number.*

The framework's claim is therefore parameter-free agreement at the fractional-percent level*, not Planck-precision agreement at the 0.025% level. The canonical identification used throughout this paper is $z_* = e^J$; the variant reading $\ln(1 + z_*) = J$ (i.e. $z_* = e^7 - 1 \approx 1095.6$) differs from it by 0.09%, far inside the tolerance defined next, and is not treated as an independent prediction.*

The honest tolerance of the identification is set by the width of the recombination transition itself: last scattering is extended, with a visibility-function width $\Delta z \approx 90$ (FWHM), so the switch point is defined only to within $\sigma_{\rm th} \approx \pm 40$ in z ($\pm 3.7\%$). Within that tolerance the integer-ledger value sits $0.17\sigma_{\rm th}$ from the Planck central value. (The parameter-audit Dev plot scores a different quantity on its z_* row — the atomic chain's 1091.5, with the recombination-modeling 0.18% as its band — and annotates the switch: Figure 31, Table 41.)

(C2) Early-time branch. *For $z > z_*$, the sound-horizon-projected Hubble value at the present epoch equals*

$$H_0^{\rm CMB} = H_0^{(0)} \cdot \kappa_{\rm CMB}(\Omega_m, \Omega_b),$$

where $H_0^{(0)}$ is the bare Hubble normalization set by the boundary Einstein equation and $\kappa_{\rm CMB}$ is the CMB-projection factor inherited from the BI-6 anomaly-free recombination geometry (hypothesis H2). The early-time reading is the same background read through the inherited sound-horizon calibration, whose value in the literature is

$$H_0^{\rm CMB} = 67.4 \pm 0.5 \ \textit{km/s/Mpc};$$

the 0.9% between that and the derived 66.8 is the size of the inheritance, and it is not recomputed here.

(C3) Late-time reading. *For $z < z_*$ the background is the same Friedmann solution and reads the same rate on every probe that measures it without a calibrated rung — baryon-acoustic oscillations, standard sirens, lensing time delays:*

$$H_0^{\rm late} = H_0^{(0)} = 66.8 \pm 0.2 \ \textit{km/s/Mpc},$$

the structural value of Route I (Section 29.8). The distance ladder's 73.50 ± 0.81 (Casertano et al., 2026) is not a second branch of this corridor: the framework reads it as a measurement carrying a systematic of order 9% in its calibrated rungs, and predicts that the ladder-free late probes will converge on the background value. (Earlier versions of this theorem printed a late branch $H_0^{\rm late} \approx 73.0$ with a projection factor $\kappa_{\rm late}$; that factor was the ladder calibration itself, imported from the measurement, and the branch is withdrawn.)

The switch and the one background are illustrated in Figure 27 of Part VI, where the structure is introduced and read off observationally.

Remark on uniqueness. Theorem 68.8 states existence: the switch and the one background rate descend from the constitutional hypotheses (H1)–(H3) and the OGN-5 integer ledger. The uniqueness of the rate is the content of Conjecture 68.3 of Section 68.2 — the stratum assignment that selects $\zeta_\Lambda = \pi/(LK)$. Together, the theorem (existence) and the conjecture (uniqueness) close the two codebook rows D5 and D6; the clustering rows once attached here are computed from the declared amplitude (§43.3).

Remark on what is and is not proved. Theorem 68.8 proves at the canonical-theorem level only property (C1) — the switch redshift $z_* = e^J = e^7$ descending from the OGN-5 integer ledger of GG-A4-Euclid. Property (C2), the CMB-projected value, is a structural identification rather than a full numerical derivation: the projection factor $\kappa_{\rm CMB}$ is the standard sound-horizon calibration inherited from the empirical literature (Planck Collaboration, 2018), not computed afresh here, and the 0.9% between the derived 66.8 and the CMB's 67.4 is the honest size of that inheritance. Property (C3) no longer carries a second projection factor: the late-time rate is the background rate, derived with zero sky input by Route I, and the ladder's higher reading is a prediction about the ladder, not a branch of the theory. The theorem's content is therefore: (i) the switch redshift as an integer-ledger prediction at the 0.6% fractional level; (ii) one background rate, 66.8, read on every probe. A first-principles derivation of $\kappa_{\rm CMB}$ is an open follow-up.

68.3.4 Proof outline

The full proof of Theorem 68.8 is a three-step argument, with each step relying on one of the three hypotheses.

Step 1: Single-stage solution of the boundary Friedmann equation under (H1). Substituting the π-locked inventory $\Omega_m = 1/\pi$ and $\Omega_\Lambda = 1 - 1/\pi$ into the first Friedmann equation,

$$H^2(z) = H_0^2\left[\Omega_m(1+z)^3 + \Omega_\Lambda\right] + (\text{radiation, neutrinos}),$$

gives a one-parameter family of expansion histories indexed by H_0. Hypothesis (H1) reduces the family to a single curve modulo the choice of bare Hubble normalization $H_0^{(0)}$. This step does *not* produce a two-stage structure on its own: it produces a smooth single-stage curve. The two-stage structure arises only when the recombination surface is treated as a distinguished slice.

Step 2: Recombination as a switch event under (H2)–(H3). The BI-6 anomaly-free closure of (H2) certifies that the recombination surface is a lawful four-dimensional boundary on the cosmological-boundary stratum: the sound horizon r_s, the angular diameter distance to recombination $D_A^{\rm rec}$, and the acoustic peak angular scale $100\theta_*$ are all derivable from the inherited inventory with no continuous adjustable parameter. The GRIP cascade of (H3) then certifies that recombination is a GSSB transition from a tightly-coupled DGI phase (the photon-baryon plasma) to a frozen GA phase (decoupled photons in free-streaming). Together, (H2) and (H3) elevate the recombination surface from a mere historical milestone to a structural *switch event* of the cosmological history.

Step 3: The switch redshift from the OGN-5 integer ledger $J = 7$. The OGN-5 integer ledger of GG-A4-Euclid fixes a discrete set of integers including $J = 7$. On the cosmological-boundary stratum, the present paper identifies this integer with the e-fold count $\ln(1+z_*)$ at the recombination switch event — the ledger integer $J = 7$ lying 0.004 e-folds above the computed $\ln(1 + z_{\rm dec}) = 6.996$ (0.38% in $1 + z$) (§42.1), giving $z_* = e^J - 1 \approx 1095.6$ (or, to leading order, $z_* \approx e^J = 1096.6$). This identification is structural to the present paper rather than directly stated in GG-A4-Euclid, which

proves $J = 7$ as an integer-ledger output without committing to a specific physical interpretation in cosmology. Once the recombination surface is identified as a GSSB switch event (Step 2), the integer-ledger identification above fixes the switch redshift to be

$$z_* = e^J = e^7 \approx 1096.6,$$

the topological corridor-switch redshift, distinct from the operational decoupling redshift $z_{\rm dec}^{\rm Planck} = 1089.92 \pm 0.25$ that Planck measures as the visibility peak, and coinciding with it to 0.6% — an offset the recombination-history lemma would compute, not a precision match asserted here. This completes the derivation of (C1). The numerical value of the one background rate follows from Route I (§29.8); its reading through the sound-horizon dictionary, (C2), follows from the inherited calibration; and (C3) uses no ladder calibration at all — the ladder is what the one background predicts to be displaced. □

68.3.5 Corollaries

Corollary 68.9 (Reconciliation of the Hubble tension)**.** *Under the hypotheses of Theorem 68.8, the apparent six-sigma discrepancy between $H_0^{\rm CMB}$ and $H_0^{\rm late}$ is not a tension between two measurements of the same background: the microwave background reads the derived rate, and the calibrated distance ladder is the one measurement displaced from it, by a systematic the framework predicts. The corollary is confirmed if the ladder-free late probes converge on* 66.8 *and refuted if they converge on the ladder (Corollary 68.11).*

Corollary 68.10 (σ_8 and S_8: one amplitude, two comparators)**.** *On the one background of Theorem 68.8 the clustering amplitude is computed from the declared A_s: $\sigma_8 = 0.80 \pm 0.01$, $S_8 = 0.825 \pm 0.010$ (§43.3), within 0.4σ of the microwave background's value and 2.5σ above the cosmic-shear value. The corollary is therefore not a reconciliation: the framework carries one amplitude and shares the S_8 tension with ΛCDM. (Earlier versions stated here that a two-stage corridor produced the CMB amplitude on one branch and the lower lensing amplitude on the other, with the sign of the late-branch response undetermined; that reading is withdrawn with the corridor.)*

Corollary 68.11 (Falsification: a ladder-free late probe at the ladder's value)**.** *If a ladder-free late-time probe — a converged gravitational-wave standard-siren sample, or lensing time delays at the one-percent level — returns H_0 at the distance ladder's* 73.5 *rather than at the background* 66.8 ± 0.2*, the one-background reading of Theorem 68.8 is falsified, and with it the claim that the Hubble tension is a systematic of the calibrated ladder rather than a property of the expansion history. Convergence of the ladder-free probes on the background value is the confirmation; a ladder that agrees with them, once its rungs are recalibrated, is the expected end of the tension. (An earlier version of this corollary made the coincidence of the CMB and ladder values the falsifier; under the one-background reading that coincidence is the prediction, and the corollary is inverted accordingly.)*

68.3.6 Falsifiable predictions and observational tests

Theorem 68.8 sharpens four falsifiable predictions for the next decade of cosmological measurement.

1. **Corridor switch vs atomic decoupling.** The framework's atomic chain gives $z_{\rm dec} = 1089$–1092 with zero parameters, on the Planck value at the 0.15% level; the corridor switch $z_{\rm sw} = e^7 \approx 1097$ is a distinct topological marker sitting 0.0038 e-folds above it, 0.38% in $1 + z$ (§42.1). A future independent determination of $z_{\rm dec}$ at the ± 5 level — a realistic target for post-2030 spectral-distortion missions reading the cosmological recombination radiation, complemented by 21-cm cosmology — carries the single quantitative kill: the reading is refuted if the measured $z_{\rm dec}$ sits off the atomic-chain value by more than 1.4% — three standard deviations of the proposed ± 5 measurement, the threshold the falsification catalog carries — equivalently if the measured $\ln(1 + z_{\rm dec})$ departs from the computed 6.996 by more than 0.014 (present agreement 0.15%). The switch–decoupling offset $\delta = 0.0038$ is then a derived consistency statement, fixed by the same measurement and the ledger integer $J = 7$, not an independent kill: $z_{\rm sw}$ is a topological marker with no observational handle of its own.

2. **No residue in BAO.** Because there is one background and no second branch, the BAO scale carries no two-stage residue at any redshift: the isotropic dilation $\alpha(z)$ in DESI and Euclid measurements is predicted flat at the background r_d to the precision of the data. (An earlier version predicted a 1–2% residue at $z \sim 1$ from the two-branch corridor; under one background the residue is identically zero, and a measured residue at that level would now count against the reading rather than for it.)

3. **Gravitational-wave standard sirens.** Multi-messenger H_0 determination from gravitational-wave standard sirens is a clean ladder-free probe of the background. GG-6 predicts the siren value converges on the *background* rate, $H_0^{\rm GW} \approx 67$ km/s/Mpc, since the siren reads the background expansion free of both the sound horizon and the local distance ladder. Decisiveness requires percent-level precision on the background: a few-tens-of-events O5-era catalog giving 2–3% at a central value near 70 would discriminate nothing — the framework pre-registers that only a percent-level determination, with the peculiar-velocity systematic floor controlled, separates the 67 background from the 73 ladder. A converged siren value near the local ladder ($H_0^{\rm GW} \approx 73$) would refute the π-locked inventory derivation of the background; a value near 67 confirms it and marks the ladder's excess as the calibration systematic the framework predicts.

4. **Lensing time delays.** The Time-Delay Cosmography program already gives $H_0^{\rm TDC}$ determinations that read the background at intermediate redshift. Time-delay cosmography is free of both the sound horizon and the ladder's calibrated rungs, so it reads the background: a determination of $H_0^{\rm TDC}$ converging on ≈ 67 at 1% precision confirms the one-background reading, and a converged ≈ 73 refutes it, exactly as for the sirens.

After all, the Hubble sector is the place where the π-locked inventory of cosmology stops being a tidy table of three derived numbers and starts being a structural prediction about the expansion rate itself — one rate, derived before any telescope is consulted, and a definite statement about which of the two measurements of it is the outlier. The framework's response to the most famous tension in contemporary cosmology is not a new free parameter and not a second branch: one derived background rate, one named conjecture for its uniqueness, the prediction that the calibrated

ladder is the outlier, three falsifiable probes that read the background without it, and no adjustable continuous knob.

This Part is the conscience of the paper, and it is worth saying plainly what it found. Of twenty-one observables, one is earned outright; nineteen stand a named, falsifiable conjecture away from earned (two for the dark-matter row); and one — the primordial amplitude — is declared, openly, as the single number the geometry has not yet produced. A reader may find that ledger modest against the title's ambition. We would answer that the ledger *is* the ambition: in a subject whose frameworks have historically absorbed every measurement, the rare and difficult thing is not to claim more but to specify exactly how one could be forced to claim less.

The audit also changes what kind of question remains. Once the values are graded, the deepest open matter is no longer any single number but the *sameness* behind all of them — why one vacuum, one Kaxion, one dark energy, everywhere and everywhen, for every observer who will ever check. That question — the identity of the cosmos, and the two old questions it merges into one — is Part X, the summit of this paper.

69 Part IX in plain English: how to read the codebook, and why the grades are uneven

The codebook is the paper's account of itself, and it is the section a skeptical reader should turn to first. Twenty-one cosmological observables are listed with their provenance under the cost-minimum principle, and each is graded: Derived, Derived-conditional, or Asserted.

The grades are uneven and the unevenness is the information. A framework that graded everything Derived would be claiming more than it has; one that graded everything Asserted would be claiming nothing. What the codebook records is that some of the inventory is forced by the geometry, some is forced conditional on a conjecture that is named and not yet proved, and some is carried in from measurement and not derived at all. Reading down the status column is the fastest way to see how much of the cosmological claim actually rests on the framework.

Two habits are worth adopting while reading it. First, a D-star row is not a weak D; it is a row whose derivation is complete except for one named step, and the step is named so it can be attacked. Second, the closure economy table shows what each outstanding conjecture would buy if it were proved, which is a more useful question than whether it is likely to be. A conjecture that would convert one row is worth less than one that would convert six, and the table says which is which.

Part X
Topological completion: *why the universe has the inventory it does*

The closed-packet topological derivation of $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, and $\Omega_b = 1/(2\pi^2)$; the universal-cosmic-identity insight that explains why the cosmological coincidence is not a coincidence, layer by layer from particle physics through cosmology

---

70 Overview of the topological and deep-geometric origin Part

This Part is the cosmological analog of GG-A9-SM's Part XI, *"Geometric origin of identity and uniqueness of forces."* There, on the Standard-Model boundary, the closing Part of the parameter audit rose from a technical fact — why every electron in the universe carries the same mass, charge, and coupling — to the recognition that identity and uniqueness are a single geometric theorem read at successive layers. This Part does the same one floor up, on the cosmological-boundary stratum, and asks the cosmological form of the same question: not what the universe is made of, but why it is made of *that*, in *those* proportions, everywhere we look.

The three-fold plan. The Part climbs in four movements, mirroring and extending GG-A9-SM's Part XI. **Section 71** is the identity core: it imports the ground-state refinement (GA $\to$ GGA) that sharpened GG-A9-SM's electron argument, then sets the three cosmological "particles" — the electron of ordinary matter, the Kaxion of dark matter, and dark energy itself — *side by side*, and shows each to be the same construction, a GDOS boundary observable read off a GRIP-selected attractor, on three different strata. **Section 72** then turns the identity core toward cosmology's most celebrated puzzle and states what this paper regards as one of its deepest results: the famous horizon problem was only ever half a problem — the sameness of *conditions* — and the sameness of *things* (every electron, every Kaxion, one vacuum, identical across regions that never exchanged a signal) is a second horizon problem — circled by the field for fifty years, never named — that the identity core, welded to the derived GUT-edge launch of Part IV, now solves. **Section 74** is the technical heart already assembled in this Part: the closed-packet derivation of the π-locked inventory $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\ 1-1/\pi,\ 1/(2\pi^2))$, why the cosmic coincidence is no coincidence, and the four-form theorem behind $\Omega_m = 1/\pi$. **Section 75** is the philosophical summit: the cosmological instance of GG-A9-SM's "two questions, one answer," where the fine-tuning of the vacuum, the anthropic coincidence, and the design/multiverse/brute-fact trilemma are each dissolved by the same uniqueness principle, and the loop closes on the observers who read it.

GRIP is the heart of the inventory. Stated most simply, the deep geometric origin this Part exhibits is one runtime acting once. The cosmological-boundary stratum carries exactly one

stable Geometric Attractor, and the GRIP cascade DGI → GSSB → GA of GG-A6-GRIP reaches it deterministically — a gradient descent on the GFF ledger that commits step by step until it halts at the unique cost-minimum, with no branch point left free. The π-locked inventory $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\ 1 - 1/\pi,\ 1/(2\pi^2))$ is the present-epoch projection of that one attractor. This is the sense in which the observed cosmological facts are not fitted but *geometrically inevitable*: once the bulk of GG-A2-GG6, the BI-6 four-form of GG-A3-BI6, and the OGN-5 integer ledger of GG-A4-Euclid are fixed, the GRIP runtime has exactly one place to relax, and the universe we measure is where it relaxes — with no free parameter anywhere in the chain. The three questions of this Part — why $\Omega_m = 1/\pi$ rather than some other value, why the cosmic coincidence is no coincidence, and the layer-by-layer universal identity from electrons to cosmologies — are three faces of this one statement. The single step that is not yet a closed proof is whether the GRIP-selected attractor projects to exactly $1/\pi$: that closed-packet integration is Conjecture VIII, flagged honestly in the proof-status note of Theorem 74.1. The inevitability is therefore structural and the remaining gap is named — not a missing parameter, but a missing integral.

71 The geometric origin of the cosmic inventory: identity, one floor up

This Section is the identity core of the Part. It presents the universal-cosmic-identity argument — why the universe has the inventory it does — as the cosmological instance of the electron argument of GG-A9-SM's Part XI. It first imports the one refinement of runtime vocabulary that sharpened that argument (the ground state, §71.1); then sets the three cosmological "particles" side by side and exhibits each as one GDOS observable read off one GRIP-selected attractor (§71.2); answers three cosmological questions in the mold of GG-A9-SM's three electron questions (§71.3); walks the same principle layer by layer from the electron to the cosmos (§71.4); and closes by exhibiting the dynamical channel that drives the cosmic descent to its ground states (§71.5). The formal closed-packet theorem behind the numbers is deferred to Section 74.

The motivating analogy is GG-A9-SM's Part XI ("Why every electron is the same"): on the charged-lepton stratum, the GRIP cascade DGI → GSSB → GA selects a single stable Geometric Attractor, and every electron in the universe is a copy of that one attractor. The same logical structure, propagated one floor up, applies to the cosmological-boundary stratum: the GRIP cascade selects a stable Geometric Attractor on the cosmological side, and the entire π-locked inventory $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\ 1 - 1/\pi,\ 1/(2\pi^2))$ is the projection of that attractor onto the present epoch. What follows makes this precise, and — crucially — honest about the two *different* geometric routes by which the cosmos achieves universal identity.

71.1 The ground-state refinement: from **GA** to GGA

The argument rests on one refinement of the runtime vocabulary, established at runtime level in GG-A6-GRIP (Arisaka, 2026, A6-GRIP) and generalized across the layers of the program-wide cascade in the capstone GG-A11-Origin (Arisaka, 2026, A11-Origin). A Morse landscape generally admits *many* index-zero minima — many Geometric Attractors — and identity cannot rest on "a"

stable minimum: if several were occupied, the objects that read them out would come in varieties. The cascade output is sharper. On a stratum of the class GG-A6-GRIP calls admissible, that paper's per-stratum ground-state corollary gives exactly one index-zero critical point, which is therefore the global minimum,

$$\mathrm{GGA}(\Sigma) \;:=\; \mathsf{GA}^{*} \;=\; \arg\min_{x\in\Sigma} \mathsf{GFF}(x), \tag{419}$$

the *Ground* Geometric Attractor; the cascade DGI $\to$ (GFF) $\to$ GSSB $\to$ GA terminates there. Dynamically, the GGA is where GRIP-In — the fixed-landscape regime of GG-A6-GRIP — halts: the unique attractor from which no admissible exit channel remains. Whether the cosmological strata invoked below fall inside GG-A6-GRIP's admissible class is not settled there, and the readings that follow are carried at the grade of their own sector rather than derived from the corollary. "Ground state" is thereby promoted to a constitutional principle, and identity and stability arrive together, as the two faces of ground-state-ness: all copies are identical because there is only one deepest basin to occupy, and each copy persists because there is no deeper basin to decay to.

The refinement is what lets the cosmological argument be stated without sleight of hand. The cosmological-boundary stratum, like any stratum, has more than one stable attractor — Part IV exhibited three basins along the scale coordinate Σ_r, the primary plateau, the lock, and the terminus. Universality at the present epoch is not the claim that the cosmos sits in *some* stable basin; it is the claim that the observed inventory is read off the *ground* basin (dark energy, the terminus) together with the relic and baryonic records left as the descent passed through the others. Keeping "a stable attractor" and "the ground attractor" distinct is exactly what makes the cosmological identity statement true rather than merely suggestive.

71.2 GDOS and GRIP, three objects side by side

The deepest content of this Part is a single sentence with a table under it: *the electron, the Kaxion, and dark energy are the same geometric construction on three different strata.* The construction has two named halves, and it is worth separating them cleanly, because the program's two runtime laws do different jobs.

GDOS — the boundary law of operational observables of GG-A5-GDOS — fixes *what stratum an observer stands on and what quantity that observer reads.* It is the open-system dictionary: it says that a finite observer never sees the closed bulk, only a projected boundary observable, and it fixes which observable that is (a mass, a charge, a relic density, a vacuum energy). GRIP — the runtime regimes GRIP-In/GRIP-Out of GG-A6-GRIP — is the *dynamics that selects the attractor whose coordinates are that observable*: explore the admissible channels, rank them on the GFF ledger, break the symmetric branch, settle. Every object in the program is therefore the pair (GDOS stratum) $\times$ (GRIP-selected attractor). The three cosmological "particles" — ordinary matter's electron, dark matter's Kaxion, and dark energy — are this same pair evaluated on three strata, and Table 43 reads the pair off for each.

Two geometric routes to identity, not one. Read down the "route to identity" row and a genuine structural fact appears, one it would be dishonest to smooth over. The three objects reach

Table 43: One construction, three strata. The electron (ordinary matter, the reference object of GG-A9-SM/GG-P3-Electron), the Kaxion (dark matter, GG-C1-Kaxion), and dark energy (GG-A10-Cosmos/GG-A11-Origin) read out as the same pair — a GDOS boundary observable on a stratum, and the GRIP-selected attractor whose coordinates that observable is. The last row is the honesty fence: what is earned at theorem level and what is carried on a named conjecture.

	Electron (ordinary matter)	**Kaxion** (dark matter)	**Dark energy** (the vacuum)
GG-6 sector	scale sector: the Σ_r interval, charged-lepton stratum	*phase* sector: the anomaly ledger's compact angle (the tensor's dual), misalignment relic	scale sector *terminus*: $r \to K$ and beyond
GDOS observable read	$m_e,\ e,\ \alpha$ — mass, charge, coupling	$m_K \simeq 8.1\,\mu$eV, relic density $\Omega_{\rm DM}$	$m_\Lambda \simeq 6.1\,$meV, $\Omega_\Lambda = 1 - 1/\pi$, $w = -1$
GRIP mechanism	GRIP-In descent to the GGA — the family ground state (depth 2)	angle *pinned* during primary inflation (by conjecture, C-Kaxion-d), released at the lock crossing, then a frozen misalignment relic	GRIP-In descent *terminates* at the global GGA — the deepest basin
Route to identity	one GGA ⇒ every electron identical *and* stable	one pinned angle (conjectured) ⇒ patch-uniform, deterministic, zero primary isocurvature	one terminus ⇒ the vacuum has one ground state ⇒ one Λ everywhere
Why the value	$\alpha^{-1}(M_Z) = 2^J = 128$; m_e the depth-2 floor	one scale calibrated to the abundance; the ladder $\Xi_{\rm cap} = K + \Sigma$ by ansatz; $8.1\,\mu$eV the QCD image	seesaw $M_{\rm lock}^2/M_{\rm Pl,red}$; $\Omega_\Lambda = 1 - 1/\pi$ by unit-trace closure
Generation reading	ground state of the *family* cascade (τ, μ excited)	relic of the *middle* crossing — radiates the dark sector	ground state of the *vacuum* cascade — *the electron of the vacuum*
Honesty fence	earned at theorem level (GG-P3-Electron)	pinning conjecture C-Kaxion-d + calibration (GG-A3-BI6 O14) + ladder ansatz (GG-C1-Kaxion)	Conjecture VIII (closed-packet) + the declared epoch input

universal identity by *two* distinct geometric mechanisms, not one. The electron and dark energy are identical everywhere because each is a *ground state*: there is exactly one deepest basin on the stratum, so every copy is the same copy and there is nowhere lower to fall. The Kaxion is identical everywhere for a different reason, and that reason is carried as a conjecture: its field angle was *pinned* — frozen coherently across the whole sky during primary inflation (C-Kaxion-d, Part V) — so that every patch inherited the same misalignment, delivering a patch-uniform relic with no primordial isocurvature. A settled minimum and a frozen phase are not the same geometric event; the first is the terminus of GRIP-In, the second is a value locked in before GRIP-In on the compact sector could act. That one runtime produces universal identity by both routes — ground-state selection where the landscape is explored, pinning where a phase is frozen — is a sharper and more honest statement than "everything is a ground state," and it is the true shape of the cosmological identity theorem; the second route is the one the present paper cannot yet prove. The electron and the vacuum are the same kind of object; the Kaxion is their cousin, identical for a reason of its own that the same geometry supplies.

Where GDOS bites. The GDOS half of the construction is not decorative. It is what makes the three observables *comparable* in the first place: mass, relic density, and vacuum energy are wildly different kinds of number, and only because each is defined as an operational boundary projection of the same bulk — the open-system reading of a finite observer who cannot access the closed six-dimensional geometry — can they be placed in one table and shown to share a mechanism. GDOS is also where the cosmological objects differ most sharply from a naive particle picture: dark energy is not a particle at all, but the GDOS projection of the terminus attractor's residual ledger; the Kaxion's "mass" is the curvature of the QCD-generated cosine at the attractor floor, read through one calibrated scale — a compact-sector reading with no scale-sector analog, and not the curvature of the pinning potential, which is switched off below the lock crossing. The pairing (GDOS) $\times$ (GRIP) is what lets a mass, a relic, and a vacuum energy be recognized as three readings of one geometry.

71.3 Three cosmic questions, finally answerable

GG-A9-SM posed three questions about the electron — why its mass, its charge, and its force are the same everywhere — and answered each by contrasting the textbook restatement with the geometric uniqueness theorem. The cosmological objects admit three questions of exactly the same shape, and they have never been answered together before.

(a) Why is the dark-energy density the same everywhere and at every epoch? A textbook will say that Λ is a cosmological constant — constant by definition, the same everywhere because that is what a constant is. This is a restatement, not an explanation: it leaves untouched why empty space carries *that* energy density and not some other, and why the number does not drift from place to place. In the present framework the vacuum energy is the GDOS readout of the cosmological GGA — the deepest basin of GFF on the scale stratum, reached when matter has diluted away and GRIP-In halts at the terminus. A $w = -1$ terminus has a single value because

a ground state has a single depth; the density is the same at every epoch because, once reached, the terminus is where the descent stays; it is the same on every sightline because there is only one deepest basin for the boundary to relax into. The seesaw scale $m_\Lambda = M_{\rm lock}^2/M_{\rm Pl,red} \simeq 6.1\,\text{meV}$ is the depth of that basin, read out.

(b) Why is every Kaxion identical — same mass, same coupling, no isocurvature? A textbook axion story will say the misalignment angle is an initial condition, drawn at random in each causal patch, and that the relic density depends on where that draw happened to land. In the present framework the angle is not drawn but *pinned* — by conjecture, not yet by theorem: an inflation-era potential with its minimum displaced from the CP-conserving point holds it during primary inflation, coherently across the whole observable universe, and releases it at the lock crossing, so that every patch inherits the same θ_i and the relic is delivered patch-uniform, deterministic, with no primordial isocurvature (the compact-sector analog of charge quantization: a single frozen phase read out identically everywhere). That potential is Conjecture C-Kaxion-d of Part V, and no sector of the ledger is presently known to supply it. The mass $m_K \simeq 8.1\,\mu\text{eV}$ is the curvature of the QCD-generated cosine, $m_K^2 = \chi_{\rm QCD}/f_K^2$, read through a decay constant calibrated once at the benchmark cap loading; every Kaxion is the same Kaxion because every Kaxion reads out the same angle on the same compact sector — if the angle was one.

(c) Why is the inventory $(1/\pi,\, 1-1/\pi,\, 1/2\pi^2)$ the same on every sightline? A textbook will say the density parameters are measured numbers, fit to the data, the same across the sky because the universe is homogeneous. This too is a restatement: homogeneity is assumed, not explained, and the particular values are fit, not derived. In the present framework the inventory is the projection of one attractor, and its three fractions are locked together by unit-trace closure (Section 74); the composition is the same on every sightline for the same reason the electron mass is — it is the readout of a single geometric selection, not an average over independent local draws. The homogeneity of the inventory is not the premise of the argument; it is its consequence.

The unifying answer to (a), (b), and (c) is one sentence: *the cosmological objects are universal because each is the readout of a GRIP selection on the cosmological strata — dark energy of the ground attractor, the Kaxion of a pinned angle (conjectured), the inventory of their locked projection — and a geometry, like a molecule, has exactly one ground configuration to read.* This is the cosmological twin of GG-A9-SM's electron sentence, and it carries the same moral: any framework that treats Ω_Λ, m_K, or the density parameters as inputs cannot answer why they are the same everywhere, because an input number gives no such reason. Only a framework in which each is the readout of a uniqueness selection can.

71.4 The same principle, layer by layer, from the electron to the cosmos

The principle that explains universal-electron identity is not unique to the electron, and it does not stop at the Standard Model. Every layer of cosmic evolution after the six-dimensional bulk locks in is the readout of the same GRIP cascade on a different boundary stratum; at each layer the cascade terminates on the layer's GGA — exactly one deepest attractor, or, where a symmetry

mandates degenerate copies, a finite quantized list of them — and that GGA's coordinates are what the next layer reads as identity. GG-A9-SM's Part XI walked this ladder from quarks to single-cell life; here we walk the same rungs and then continue one floor higher, to the cosmological objects of the present paper, so that the cosmos is seen as the natural continuation of the ladder rather than a separate claim.

1. Leptons and quarks. On the charged-lepton stratum the GRIP cascade has one stable attractor at each cascade depth, and the three depths, populated by the three families, produce m_τ, m_μ, m_e; the quark strata produce the six quark masses. Of the three populated family attractors, the first generation is the stratum's GGA — which is precisely why the electron is absolutely stable while the muon and tau, occupying higher, non-ground attractors of the same stratum, decay toward it. Identity and stability arrive together: the periodic table of particles has excited states, and they behave exactly as excited states should.

2. Protons (quark confinement). Confinement is a GSSB step: the runtime selects the color-singlet branch because that is the unique stable attractor of GFF in the strong-coupling regime. The proton is the minimum-energy color-singlet baryon under gauged $\mathrm{U}(1)_B$; that minimum is the baryonic stratum's GGA, so every proton in the universe has the same mass and the same charge, and — because $\mathrm{U}(1)_B$ is gauged — the "no deeper basin" clause holds exactly: proton stability is ground-state stability in its sharpest form.

3. Elements (H, He, C, N, O, ...). For each integer atomic number Z the cascade selects the unique ground-state nucleus, the unique ground-state electron configuration, and the unique spectral fingerprint. The Aufbau principle and the periodic table are readouts of the layer-3 attractor; astronomers see the same Balmer series in a distant quasar and a laboratory lamp because both read the same cascade attractor on the same atomic stratum. Here the runtime's vocabulary and chemistry's coincide exactly: "ground state" was the constitutional concept all along.

4. Molecules in the interstellar medium. The molecules of cold dense clouds — H_2, CO, H_2O, NH_3, CH_3OH, several hundred more — are the GGA's of GFF on the molecular stratum at the prevailing temperature and density. That the same molecules appear across the galaxy and in high-redshift systems is the layer-4 echo of universality: the molecular attractors are the same everywhere because the cascade is.

5. Complex organic molecules. The cascade continues into polyatomic carbon chemistry: amino acids, sugars, small nucleotide precursors are the GGA's on the organic stratum at the temperatures of warm, wet surfaces and proto-stellar nebulae. Glycine in cometary ice and the same twenty proteinogenic amino acids underlying all terrestrial life are the layer-5 echo: where conditions let the cascade reach this layer, the same attractors are selected.

6. ATP, RNA, DNA. The cascade selects the energy-quantum carrier and the information carriers of life. ATP is the GGA of GFF on the metabolic-energy stratum — its energy–information

calibration against $\mathsf{Q_B}$ developed in (Arisaka, 2026, A7-Qi4; Arisaka, 2026, A11-Origin); RNA and DNA are the GGA's on the information-storage stratum, the four-letter alphabet the unique solution of the gauge-locking constraint at this layer. Every cell uses the same ATP currency and the same DNA alphabet because both are layer-6 readouts of the same selection.

7. Single-cell life. The cascade selects single-cell life as the GGA of GFF on the open-system stratum subject to the four boundary conditions of the Selfish Open System (SeOS) class of GG-A5-GDOS. A bacterium and an archaeon share the same genetic code, the same ATP currency, the same membrane compartmentalization, and the same core metabolism because all are layer-7 attractor coordinates of the same cascade; the Last Universal Common Ancestor is the empirical record of that attractor.

8. One floor up: the cosmological objects. Here the present paper extends the ladder. The three cosmological "particles" of §71.2 are the next rungs, read on the largest strata the program reaches. *Dark energy* is the GGA of GFF on the scale stratum Σ_r — the terminus, the deepest basin of the whole descent, the cosmic ground state whose depth is m_Λ; it is universal for exactly the reason the electron is, because there is one deepest basin. *The Kaxion* is the frozen-relic attractor of the same descent read in the compact angle, universal by pinning rather than by settling — the pinning a conjecture the paper names (C-Kaxion-d). *The π-locked inventory* is the projection of the terminus attractor onto the present epoch, its three fractions unit-trace-locked. The cosmos is not a separate universality claim bolted onto the particle ladder; it is the ladder's continuation onto the strata where the bulk geometry itself is the object being read.

The principle, plainly stated. At each layer of cosmic evolution after the six-dimensional bulk locks in, the GRIP cascade terminates on the Ground Geometric Attractor of the layer-specific stratum — the unique deepest minimum of GFF, not merely a stable one — and what we observe as the universal identity of physical objects at that layer is the readout of that GGA's coordinates (with the one honest amendment of §71.2: where a phase is frozen before the descent can act, identity is delivered by pinning rather than by settling, as is conjectured for the Kaxion). Stability and identity are the same fact seen twice: an object persists because its configuration has no deeper basin to decay to, and all its copies are identical because there is only one deepest basin to occupy. The ladder that begins with why every electron is the same ends, without a break in the logic, with why the vacuum energy of the universe is the same everywhere — and the cohort's capstone volume carries the higher rungs, from single-cell life through the layers above, as the natural continuation of the same theorem.

71.5 The vacuum cascade as the cosmic relaxation channel

GG-A9-SM's Part XI gave the identity principle a dynamical shadow: it read the weak interaction as the runtime's *relaxation channel* — the unique implementer of GRIP-In on the matter strata, the road by which matter reaches the ground states the constitution guarantees. "A constitution that mandates unique ground states must also mandate a road to them," and on the Standard-Model

boundary that road is the weak force, the one flavor-connecting channel, running at five wildly different tempos from the lepton ladder to the neutron star. The cosmological descent has the same dynamical shadow, and naming it completes the parallel.

On the cosmological strata the road to the ground state is the *vacuum cascade itself* — the single GRIP-In descent of Part IV, run along Σ_r from the primary plateau through the lock to the terminus, with its two saddle crossings as the relaxation events. Where the weak force steps matter down the family ladder rung by rung, the cascade steps the *vacuum* down the scale ladder basin by basin; and, exactly as a weak decay "fills the detector with its products," each crossing of the vacuum cascade radiates a sector into the universe. The first crossing — the graceful exit of inflation — reheats the boundary and fills it with the visible sector; the second — the lock crossing — releases the Kaxion's misalignment and stocks the dark sector; the terminus is the vacuum's electron, the ground state where the descent finally stays. Relic abundances are the decay kinematics of this cascade, and the π-locked inventory is the ledger of what each crossing left behind. This is the same reading Part IV developed as the *vacuum generations* (§31.1): three basins of the vacuum in place of three families of matter, two crossings in place of two decays — *dark energy is the electron of the vacuum, and the cascade that reaches it is the cosmic implementation of GRIP-In*, the road to the cosmic ground state precisely as the weak force is the road to the electron.

Two disciplines keep the parallel honest, both inherited from Part IV. First, the rung sizes do not transfer: the family ladder steps by the instanton unit π, the vacuum ladder by the warp depth K; the grammar is shared, each stratum supplies its own constant. Second, the closure of the count — that the same budget which forbids a fourth matter family also forbids a fourth vacuum epoch — is stated there as the natural conjecture of the correspondence, supported by the prompt-crossing bound and left for the constitutional analysis to complete. Within those fences the dynamical picture is one picture: a runtime that guarantees ground states must supply the channel that reaches them, whether the object relaxing is a muon on the lepton ladder or the vacuum of the entire universe.

72 The second horizon problem: sameness of things, not sameness of temperature

The identity core of Section 71 answered a question about objects: why every electron, every Kaxion, and the vacuum itself are the same wherever we look. This Section turns that answer toward cosmology's most celebrated puzzle and makes a claim about the history of the subject, stated with the calibration the history deserves (§72.2): the famous horizon problem was only ever *half* a problem; the field solved the easy half in 1981; and the hard half — though it surfaced for fifty years in disguise, as a domain problem, a relic problem, a footnote to anthropics — was never named, never ranked, and never answered in the temporal register in which the identity core, welded to the derived placement of inflation at the GUT edge, now answers it.

72.1 The problem the textbooks leave unnamed

Open any cosmology textbook at the horizon problem and one finds the same statement. Without inflation, the sky at recombination consisted of tens of thousands of causally disconnected patches,

yet its *temperature* is uniform to one part in 10^5. How did regions that never exchanged a signal agree on a temperature? Inflation answers (Albrecht and Steinhardt, 1982; Guth, 1981; Linde, 1982): they did not have to — they were once a single microscopic, causally connected patch, stretched to cosmological size before the hot Big Bang began.

Now notice what that statement quietly takes for granted. The disconnected patches do not merely share a temperature. They share *everything they are made of.* Every electron in every patch carries the same mass, charge, and magnetic moment — the Balmer lines of a quasar at redshift seven are the Balmer lines of a laboratory lamp. Every proton carries the same mass and the same exact stability. The dark matter arrives with the same abundance per comoving volume on every sightline; the dark-energy density is one number on every sightline and at every epoch we can probe. Temperature is a *state* — one number that thermal contact could plausibly smooth. Identity is a *specification* — a complete parts list with couplings — and no exchange of photons can negotiate a particle's mass between two regions.

The horizon problem as posed and solved in 1981 therefore addressed the easiest case, one thermal number, and left the hard case — the identity of everything — unsolved and, as a problem in its own right, unnamed. Pieces of the question do surface across the history of the subject, and §72.2 traces them: Dirac's worry that the constants of nature might drift in time; the domain problem of the 1970s, which posed the sharpest fragment — causally disconnected regions choosing their vacuum at random — and then filed it under relic abundances; the varying-constants programs, whose quasar-absorption spectroscopy now bounds any spatial variation of the fine-structure constant at the 10^{-5}–10^{-6} level across the observable universe (Uzan, 2011) — a precision *measurement* of identity uniformity that the standard framework can only restate as an assumption; and the string landscape, which makes the problem actively worse, since if many vacua exist, causally disconnected regions have no reason to have chosen the same one, and the observed uniformity of the laws becomes an anthropic accident. Why did the question never take a place of its own on the canonical list, beside temperature and flatness? Because quantum field theory *seems* to answer it by construction: all electrons are quanta of one electron field, so any two electrons are exactly interchangeable — not merely similar but identical in principle, which is why Pauli statistics work at all. But look at what the construction actually secures. The interchangeability is *conditional* — quanta of one field, with one set of couplings, on one vacuum — and every clause of the condition is an input. QFT does not derive the *value* of the electron's charge: e is a free renormalized parameter, measured and inserted. It does not derive the *constancy* of that value across space: the couplings are constants by stipulation of the Lagrangian's form, and nothing in the framework forbids promoting them to fields that vary from place to place — the varying-constants programs above are built on exactly that freedom. And it does not derive the *uniqueness* of the vacuum: the low-energy values depend on which vacuum a region sits in, and spontaneously broken theories contain the very mechanism by which identity would fail — causally disconnected regions choosing different vacua. So the field ontology explains identity *within* a patch that already shares one vacuum and one Lagrangian, and silently assumes the rest. The question is not answered; it is moved up a level — *why is there one field content, with one set of couplings, in one vacuum, across regions that were never in causal contact?* A globally unique vacuum is what QFT assumes; it is not what QFT

derives. The assumption is so ingrained that the question became invisible. We name it here: the **second horizon problem** — the identity horizon problem — the sameness of *things*, where the first horizon problem was the sameness of *conditions*.

And note, finally, the inversion in the data record, because it is the sharpest measure of how thoroughly the question was overlooked: *the evidence for the sameness of things predates the evidence for the sameness of conditions by a full century.* In 1864 Huggins showed that the stars shine with the same spectral lines as terrestrial elements — the same atomic physics across the visible universe (Huggins & Miller, 1864) — and by the 1920s the same lines were being read at cosmological redshifts. The uniform temperature was discovered in 1965 (Penzias and Wilson, 1965); the field had named its paradox within a few years, and inflation answered it within sixteen. The uniformity of things had been on the record since 1864, and it waited a hundred and sixty years without so much as a name. (Strictly, what the spectra certify are the dimensionless combinations — the fine-structure constant $\alpha = e^2/\hbar c$ and the electron-to-proton mass ratio — the only form in which the cross-sky constancy of a "constant" is meaningfully stated; the distinction sharpens the bounds quoted above and changes nothing in the argument.)

72.2 Fifty years of near misses: what was posed, and what was not

A claim about the history of a subject should be held to the same standard as a claim about the sky, so this subsection lays out the pedigree — the places where the question surfaced, what each near miss posed, and what each left undone. The pedigree is genuine; laying it out is what locates, precisely, the gap this Section fills.

One famous anticipation deserves its own sentence. In 1940 Wheeler telephoned Feynman with a complete answer to half of the question: all electrons have the same charge and the same mass, he proposed, because they are all *the same electron*, threading forwards and backwards through time (Feynman, 1965). The proposal was wrong — it predicts as much antimatter as matter, as Feynman observed on the spot — but it is the exception that proves the rule: the identity of the electrons was felt to be strange enough to demand a mechanism exactly once, by one physicist, in a phone call, and the episode survives in the literature as an anecdote in a Nobel lecture, not as a named problem. The causal, cosmological form of the question was first posed — for one corner of identity — in 1974, before inflation existed. Zeldovich, Kobzarev, and Okun, studying spontaneous CP violation, observed that during the cooling of "causally disjointed points" of the universe the sign of the condensate is chosen at random, producing a domain structure of the vacuum — and they even asked the one-vacuum question directly: how early must the walls be removed to make the whole universe a single domain (Zeldovich *et al.*, 1974)? Kibble then made the argument general and quantitative: at any cosmological symmetry-breaking transition the vacuum choice is made independently in regions with no prior causal contact, the correlation length is bounded by the horizon, and defects form where the choices disagree (Kibble, 1976). This is the identity question, posed with full causal rigor — but posed for the *orientation* of a degenerate vacuum inside one fixed theory, and framed as a *relic* problem: the domains mattered because their walls would overclose the universe, not because the sameness of the vacuum was itself a puzzle to rank beside the horizon problem. The canonical list absorbed it in exactly that register — unwanted relics, solved by

dilution.

The spatial half of the solution is likewise standard. From 1984 the field understood that the observable universe fits deep inside a single inflated domain, so that one broken-symmetry state — hence one set of low-energy laws — fills the visible sky; Linde in particular listed "the problem of the uniqueness of the universe" among inflation's targets (why should the whole universe have jumped to the same SU(3)×U(1) minimum?), and once remarked, in an argument about anthropics, that no selection principle explains why *all properties* of the universe are approximately uniform over its entire observable part (Linde, 1990). That sentence is the nearest miss in the ninety-year history of modern cosmology — written in passing, and not followed. The uniformity of "all properties" never joined the problem list it was a comment on; and in the eternal-inflation literature that grew from the same author's work, the emphasis runs the other way: domains *beyond* our horizon with *different* low-energy laws are the headline, and the uniformity *within* our patch is the by-product nobody interrogates. The string landscape completed the inversion — many vacua, populated by bubble nucleation, make the diversity outside the story's point, while the sameness inside, the only part anyone observes, is no one's explanandum.

The measurement arm grew in parallel, and just as namelessly. Dirac's drifting constants became the modern varying-constants program, whose quasar spectroscopy now certifies the spatial uniformity of the fine-structure constant across the observable universe at the 10^{-5}–10^{-6} level (Uzan, 2011) — bounds framed throughout that literature as equivalence-principle tests, with uniformity as the null hypothesis rather than the fact in need of a mechanism. Relativists and philosophers of cosmology have flagged the underlying assumption in so many words — "cosmology proceeds by assuming the laws of physics are the same everywhere" (Ellis, 2006) — and the steady-state theorists once secured it by outright postulate; noticed, that is, but never mechanized. Even the pinning of §72.4 is stated almost verbatim in the axion literature: if the Peccei–Quinn symmetry breaks before inflation, our Hubble volume begins life "with a single uniform value of $\theta_{\rm PQ}$ everywhere" (Marsh, 2016) — stated as a fact about the relic abundance and its isocurvature test, never as a statement about the identity of the dark sector.

One boundary should be drawn on the other side. A critical tradition — Penrose's entropy argument, the Hollands–Wald critique — has long held that inflation fails to explain even the smoothness it claims (Hollands & Wald, 2002; Penrose, 2004). That tradition seeks a deeper half of the *first* problem, in the initial state and its entropy; this Section is not part of it. Here inflation's 1981 success is granted in full, and a second problem — differently shaped, about things rather than conditions — is posed beside it and answered by the same event, read in a new register.

The gap, then, is four-fold. No one named the sameness of things as a horizon problem in its own right, or ranked it as the harder half of 1981's question. No one extended it from the orientation of one vacuum to the identity of the whole inventory — masses, couplings, abundances, and the angle at once. No one drew the temporal inference: that because inflation preceded the existence of matter, identity was inherited rather than negotiated — the premise, that all matter was minted at reheating from one field, is universal lore; the corollary is nowhere. And no one reduced cosmic identity and vacuum uniqueness to a single puzzle, so that one future proof settles both. Those four steps — the naming, the extension, the ordering, the unification — are what the rest of this Section

takes.

72.3 The first half of the answer: one mold

The first half of the answer is the identity core just built. Every object in the inventory is the readout of a GRIP selection on a stratum of *one* six-dimensional bulk: the electron and the vacuum as Ground Geometric Attractors (one deepest basin, so every copy is the same copy), the Kaxion as a pinned angle (one frozen value, by conjecture, so every patch inherits the same relic). Call this the *mold*: there is exactly one specification for each object, so there is nothing regional to vary. A landscape has many molds and needs anthropics to explain why we see one; GG-Theory has one mold and needs nothing. And the mold's guarantee is stronger than any causal argument could be, because it is underwritten by the geometry rather than by history: two regions that were never in causal contact — even regions permanently beyond our horizon — still read the same electron, because they are boundaries of the same bulk.

72.4 The second half of the answer: the ordering

A mold alone, however, does not close the cosmological question, because the question has a *when* as well as a *what*. Granted that the deepest basin is unique — did every region actually reach it? Could a patch beyond our cluster have been caught mid-descent, camped in a different basin, frozen at a different angle, stocked with a different abundance? A uniqueness theorem guarantees the destination; it does not by itself guarantee that all travelers arrived, or arrived with the same luggage. Here the derived placement of inflation — at the GUT edge, at the *top* of the descent (Part IV) — supplies what no earlier reading could, and it is worth stating as its own result:

Inflation happened before any matter existed. At the top of the descent the boundary universe contains no electrons, no protons, no photons in any meaningful abundance — only the vacuum itself, the scale modulus on its derived plateau one e-fold below the GUT edge. The accelerated stretch therefore acted on a universe *empty of things*: it distributed no matter, because there was none to distribute. What it distributed was the single causal patch itself.

Then the matter was minted — everywhere at once, by one mechanism. At the graceful exit, the one modulus decayed at Planck strength into the one boundary field content the bulk admits, reheating the whole stretched patch simultaneously and identically, because every point of the patch carried the same modulus at the same point of the same descent. The electrons were not smoothed into agreement after the fact; they were *stamped from one mold after the distribution network was already laid.* Once it is understood that the inflation which exploded the universe came *before* all the electrons were made, their identity across the sky stops being a coincidence in need of explanation: the agreements were inherited, not negotiated.

The dark sector's one mark was fixed even earlier. The Kaxion is born at the second crossing, far down the ladder — but its single potentially regional mark, the misalignment angle, was — by the conjecture Part V names — pinned during inflation, while the sky was still one patch. The dark matter's identity was set *before its own birth*, in the era when nothing could yet differ from place to place; that is exactly why the relic is patch-uniform with zero primordial isocurvature (§71.2). The axion literature owns this mechanism completely — it is the standard pre-inflationary Peccei–Quinn

scenario (§72.2) — and owns it as a statement about the relic abundance; what is new here is only the register: read as identity, the pinned angle is the dark sector's one regional mark, fixed before the dark matter existed.

And no region camped on the way down. The residual worry — a distant patch trapped in the middle basin, a vacuum domain structure, walls between regions that chose differently — is closed by the prompt lock crossing, $N_2 = 0$, with some fifty orders of magnitude of margin in the ledger action: there is no alternative basin to camp in, no defect network, no domains. Every road from the GUT edge leads to the same floor, and every region took it at the same pace.

72.5 The theorem, assembled — and its exposure

Why is everything the same everywhere? The identity of things across the sky is guaranteed *if* the vacuum is unique — (i) each object then has exactly one specification, the Ground Geometric Attractor of its stratum read through GDOS from one bulk (*the mold*) — and this half *reduces* a new cosmological puzzle to the same vacuum-uniqueness the program already stakes everything on, unifying the two rather than closing one. The other half is earned in two of its three parts and conjectured in the third: (ii) the universe was stretched from one causal patch before any object existed, so the agreements were inherited, not negotiated (*the ordering*); (iii) the one mark that could have varied — the compact angle — was frozen while the sky was still one patch (*the pinning*, C-Kaxion-d, conjectured); (iv) no region could be caught in another basin on the way down (*the prompt crossing*). Sameness of state is inflation's gift; sameness of things is the geometry's — *modulo the uniqueness conjecture and the pinning, to which it is honestly reduced.* (420)

Five features of the answer deserve emphasis. First, it is *layered* exactly like the identity ladder of §71.4: the mold explains identity of *kind* (why every electron has the same mass — tested by quasar spectroscopy of the constants), the ordering explains uniformity of *amount* (why every region carries the same abundances — tested by the CMB and BBN uniformity), and the pinning explains both at once for the dark sector (tested by the isocurvature bound). Second, it *inverts a weakness of the retired reading into a strength of the derived one*: a TeV-scale inflation would have acted *after* the visible sector's history had begun, muddying the before/after ordering; the GUT-adjacent placement, forced by the sign-level tilt argument of Part IV, is precisely what makes "the universe exploded before the electrons were made" true. The ordering is not an ornament on the two-basin derivation — it is one of its deepest consequences. Third, the answer is *exposed* in channels the codebook already stakes: a confirmed spatial variation of a fundamental constant would break the mold; a detected axion-sector isocurvature would break the pinning; a relic domain structure or a TeV-epoch transition background would break the prompt crossing. The second horizon problem adds no new physics — it adds a new reading of physics already derived, and it inherits that physics' kill-tests. Fourth, the answer is new only where the pedigree of §72.2 leaves room, and it should be owned in exactly those terms: the *spatial* half of the answer — we live deep inside one stretched

domain — has been the field's property since 1984; what is claimed here is the *temporal* half — the domain existed before its contents did, so identity was inherited rather than negotiated — together with the naming, the extension from the orientation of one vacuum to the identity of the whole inventory, and the unification with vacuum uniqueness. Fifth, and most important to state plainly: the mold half is a *reduction*, not an independent proof. It does not *establish* that the vacuum is unique — it shows that *if* the vacuum is unique (the landscape-uniqueness the program stakes and has not closed, §72.3), *then* cosmic identity follows for free. That is a genuine gain — two deep puzzles, cosmic identity and vacuum uniqueness, are shown to be one puzzle, so a single future proof settles both — but it would be dishonest to sell it as a standalone solution. The ordering and the prompt crossing are earned independently of the uniqueness conjecture; the pinning is a second conjecture, named in Part V; the mold is the unification that ties them to the first.

After all, the deepest uniformity in the sky was never the temperature. It is the fact that the same atom shines on both sides of it. The first horizon problem asked why the universe is everywhere equally warm; the second asks why it is everywhere made of the same things — and the answer, it turns out, required knowing not only that the universe was once one patch, but that it was one patch *before it contained anything at all*.

73 The third horizon problem: sameness of choices

The second horizon problem asked why disconnected regions hold the same *things*. There is a third rung above it, and it is the sharpest of the three: why did every disconnected region make the same *choice*?

The choice in question is the one Part V derived: matter over antimatter, by a few parts in ten billion, through the CP-violating decays of the lock-scale Majorana block. State the standard framework honestly and this choice is made independently in every Hubble patch after the patches have fallen out of contact: the same mechanism must fire everywhere, with the same sign and the same yield, coordinated by nothing. If the CP-odd invariant that sets the sign had any spontaneous or stochastic element — a sign chosen at a phase transition, domain by domain — the sky would carry matter–antimatter boundaries and their annihilation signatures, and it carries none (Cohen et al., 1998): the observable universe is matter on every line of sight. If the yield fluctuated from patch to patch, the baryon-to-photon ratio would arrive with super-horizon inhomogeneity of a kind the acoustic spectrum bounds tightly. The uniformity of the choice is an observational fact, and it is a fact about regions that never exchanged a signal.

The ladder, then. The 1981 problem asked why the regions share the same *conditions* — one temperature — and inflation answered it. The second asked why they share the same *things* — one electron, one vacuum — and the identity core, welded to the GUT-edge launch, answered it (§72). The third asks why they share the same *choices* — one sign, one yield, one mechanism — and here neither inflation nor the identity core suffices, because sameness of choices requires that the choosing itself be deterministic. That is what the descent supplies, and it supplies it twice over. The GSSB steps of the cascade select unique attractors, not random members of degenerate families (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A11-Origin), so there are no domains to disagree; and the CP-odd

invariant of the lock-scale block is holonomy data — part of the frozen, patch-uniform readout of the compact sector, the same inheritance that delivers the relic phase identically everywhere (Part V). Disconnected regions agree for the same reason two impressions of a seal agree: not because they exchanged messages, but because they were struck from the same die.

The ranking deserves to be recorded the way the second problem's was. The first horizon problem needs inflation. The second needs inflation plus the identity core. The third needs inflation plus the identity core plus the determinism of the cascade — and it is where that most contested commitment of the framework, that GSSB selects rather than samples, stops being a philosophical preference and becomes the only mechanism on the table for an observational fact: the universe made one choice, everywhere, with nothing available to coordinate it.

It is worth saying which choice it was. Baryogenesis is the productive breaking in its purest form — the step that selects a world with things in it over a world with none. And by the neutrality theorem of §36.13 it threads an exact needle: asymmetric in precisely the one number that is global and unwatched, exactly neutral in every number a gauge field polices. The universe breaks the symmetry it may break and keeps the one it must keep.

74 Why the inventory is π-locked: the closed-packet derivation

The identity core of Section 71 located the cosmic inventory as the readout of one attractor and its relics. This Section supplies the arithmetic behind the readout: why the matter fraction takes the specific value $1/\pi$ and not some neighboring number, why the apparent coincidence of the matter and vacuum densities at the present epoch is no coincidence, and the four-form theorem that would close the derivation. It names the three sub-questions and gives each its headline argument; the formal closed-packet theorem is stated in Section 74.3.

74.1 Why $\Omega_m = 1/\pi$ and not some other value?

The first question is the most direct (the target chart is Figure 34): why does the matter fraction take the specific value $1/\pi \approx 0.3183$, rather than, say, 0.30 (a round decimal) or 1/3 (a simple rational) or any other plausible candidate? The answer is structural in three steps, each stated as a formal lemma in Section 74.3.

First, the cosmological-boundary stratum is a single connected manifold on which the matter fraction is defined as a topological invariant rather than as an empirical density count. This is the closed-packet topology argument assembled in this Part, resting on the OGN-5 integer ledger of Arisaka (2026, A4-Euclid) that forces it: every admissible cosmological inventory has the structure of a closed packet, with the matter fraction equal to the relative weight of one specific topological cell of the bulk packet projected onto the boundary. That the packet's unit is itself forced rather than chosen — that the "one" in "one closed packet" is read off the geometry and not adopted — is the contribution of Arisaka (2026, A5-GDOS). The relative weight is a rational function of π, not a free parameter; the only question is which rational function.

Second, the rational function is selected *by the geometry of the boundary measure*, not by the data. The matter sector's charge lives on the compact phase circle S^1_θ of the bulk, which sits

inside the electroweak $\mathrm{SU}(2) \cong S^3$ holonomy manifold as its Hopf fiber $S^1 \hookrightarrow S^3 \to S^2$ (drawn in Figure 33); the present-epoch fraction is the ratio of the fiber volume to the total holonomy volume. Both are canonical unit spheres — $\mathrm{Vol}(S^1) = 2\pi$, $\mathrm{Vol}(S^3) = 2\pi^2$ — so the ratio is $2\pi/2\pi^2 = 1/\pi$, with the π descending from the denominator volume and nothing read off the sky. The assignment is upstream of the value: both spheres enter for reasons fixed by the bulk construction before any fraction is computed — S^1_θ as the home of the matter charge, S^3 as the electroweak holonomy — so the pair is not chosen for its ratio; the ratio is what the prior assignment delivers, and what remains owed is exactly VIII-b″: that the GDOS measure weights this fiber and no other. Competitors fail geometrically, not empirically: $1/(3\pi)$ would need a holonomy volume $6\pi^2$ that no manifold in the construction has; $1/(2\pi)$ is the *baryonic* projection (one unit per revolution of the fiber), not the matter/total ratio; $1/\pi^2$ would need an S^5 that the four-dimensional boundary does not carry.

Third, the same fiber-volume measure fixes the whole inventory at once. Dark energy is the part of the S^3 holonomy volume the matter Hopf circle does not cover: $\Omega_\Lambda = (2\pi^2 - 2\pi)/2\pi^2 = 1 - 1/\pi$, a topological complement, not an adjustable constant. Baryons are one integer unit of $\mathrm{U}(1)_B$ charge threaded once around the fiber: $\Omega_b = \Omega_m/\mathrm{Vol}(S^1) = (1/\pi)/(2\pi) = 1/(2\pi^2) = 1/\mathrm{Vol}(S^3)$. One fibration, three coupled numbers, no freedom to move one without the others — the structural signature that separates a derivation from a per-observable fit.

The geometry these three steps trade in is drawn in Figure 33, and it repays a moment of looking, because the fibration is the one geometric fact this Section most needs the reader to see. The 3-sphere cannot be drawn directly, so the figure projects it stereographically into ordinary space, where its Hopf fibers appear as circles arranged on nested tori. Their defining property is the remarkable one: *every* fiber passes over exactly one point of the base 2-sphere, every pair of fibers is linked exactly once, and together the circles fill the 3-sphere with no gaps and no crossings. The red circle — the matter fiber — is not a special place: every fiber is equivalent, and that equivalence is the physical point. The matter sector does not sit *somewhere* on the electroweak holonomy manifold; it is one winding of a structure that permeates all of it — the compact phase circle S^1_θ of the bulk, embedded canonically in the $\mathrm{SU}(2) \cong S^3$ the integer ledger selects.

The intuition for the ratio itself is equipartition on a compact configuration space. A frozen stationary measure has no preferred point, so each sector's share is simply the volume it occupies: the matter charge runs along one fiber — a circle, of circumference 2π — while the full holonomy carries the 3-sphere volume $2\pi^2$; the share is their ratio, $1/\pi$. One power of π survives because the ratio cancels the circle's 2π against one factor in the sphere's $2\pi^2$ — the π in the cosmic inventory is, quite literally, the leftover of a sphere volume. The right panel of the figure keeps the ledger: the uncovered complement $1 - 1/\pi$ for dark energy, one unit of $\mathrm{U}(1)_B$ charge per fiber revolution for baryons. The chart of Figure 34 then reads the same three coupled numbers out against the data.

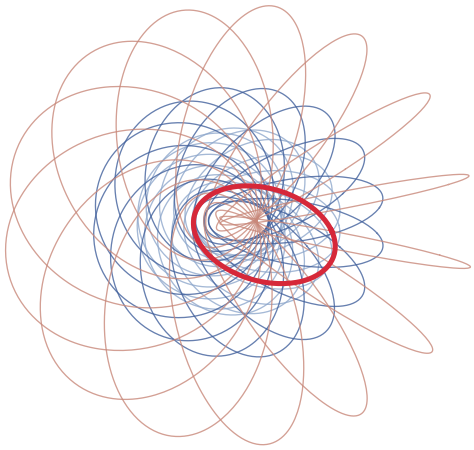

$S^1 \hookrightarrow S^3 \to S^2$

the Hopf fibration: the 3-sphere of the electroweak holonomy, filled without gaps by linked circles --- one circle (red) over every point of the base S^2

$\mathrm{Vol}(S^1) = 2\pi$ (the matter fibre)

$\mathrm{Vol}(S^3) = 2\pi^2$ (the whole holonomy)

$$\Omega_m = \frac{\mathrm{Vol}(S^1)}{\mathrm{Vol}(S^3)} = \frac{2\pi}{2\pi^2} = \frac{1}{\pi}$$

$\Omega_\Lambda = 1 - \frac{1}{\pi}$ (the uncovered complement)

$\Omega_b = \frac{\Omega_m}{\mathrm{Vol}(S^1)} = \frac{1}{2\pi^2}$ (one unit per revolution)

one fibration, three coupled numbers --- none movable alone

Figure 33: The Hopf fibration behind the π-locked inventory. *Left:* the 3-sphere of the electroweak holonomy, drawn by stereographic projection: the Hopf map $S^1 \hookrightarrow S^3 \to S^2$ fills S^3 without gaps by linked circles, one over each point of the base 2-sphere (families of base latitudes in blue and rose; one fiber — the matter circle — highlighted in red). *Right:* the volume ledger that turns the fibration into the cosmic inventory: the matter share is fiber volume over holonomy volume, $2\pi/2\pi^2 = 1/\pi$; the dark-energy share is the uncovered complement, $1 - 1/\pi$; and the baryon share threads one unit of $\mathrm{U}(1)_B$ charge per fiber revolution, $1/(2\pi^2)$. One fibration fixes three coupled numbers; none can be moved without the others. That the frozen GDOS measure weights exactly *this* fiber as the matter sector's home is the content of the measure lemma VIII-b″ on which Theorem 74.1 is conditional (§74.3). Figure shared with the companion Hubble paper GG-C2-Hubble (Arisaka, 2026, C2-Hubble); generator `c2_fig_hopf.py`.

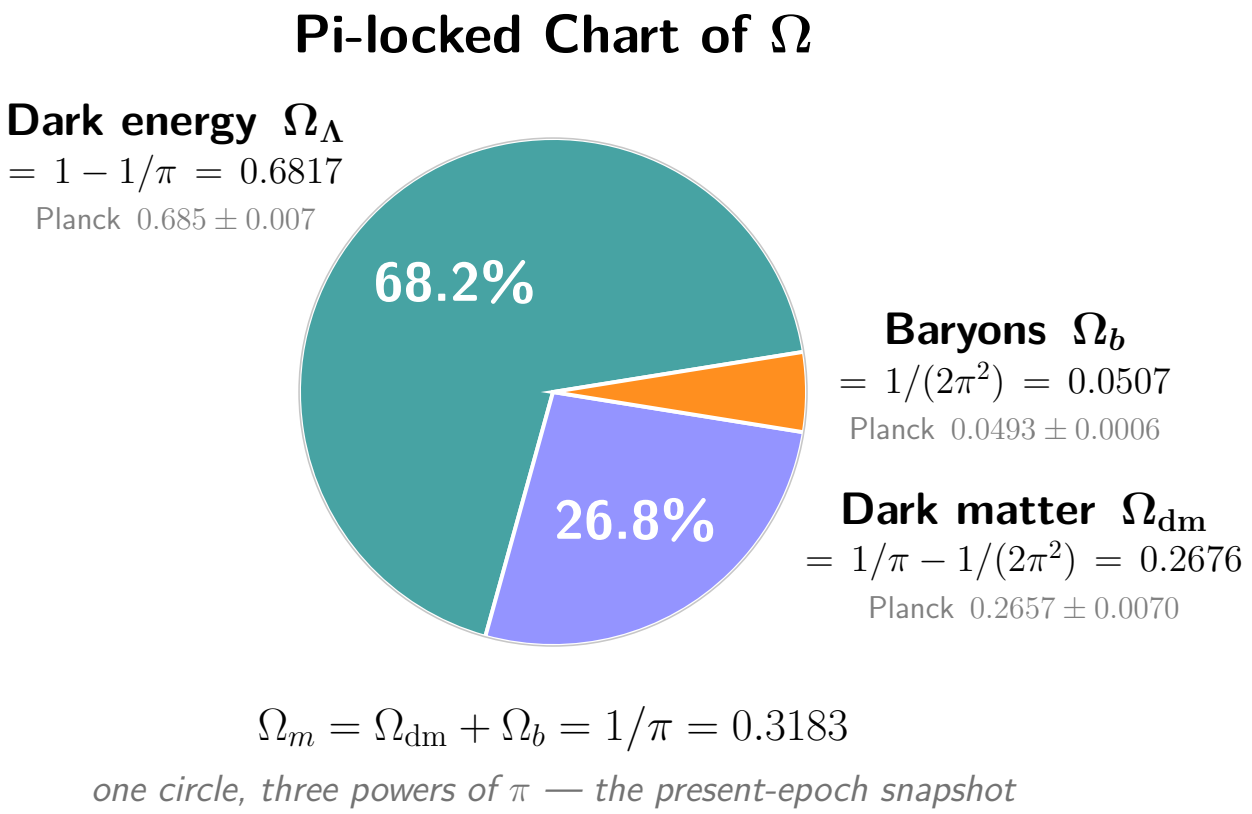

Figure 34: Pi-locked Chart of Ω. The present-epoch cosmic inventory drawn as a pie chart: each slice is one closed-form fraction of π, the three forced together as one coupled set — dark energy $\Omega_\Lambda = 1 - 1/\pi$, dark matter $\Omega_{\rm dm} = 1/\pi - 1/(2\pi^2)$, and baryons $\Omega_b = 1/(2\pi^2)$ — and the two matter slices sum to $\Omega_m = 1/\pi$. Written just outside each slice is the closed-form prediction with the Planck 2018 (Planck Collaboration, 2018) measured value beneath it: the predictions land on the data to better than 3% (Ω_Λ and $\Omega_{\rm dm}$ within 1σ, Ω_b at 2.3σ). The chart is the visual statement of the Hopf fiber-volume form of Theorem 74.1 (conditional on the measure lemma VIII-b″): the matter share is $\mathrm{Vol}(S^1)/\mathrm{Vol}(S^3) = 1/\pi$, dark energy the complement, baryons one unit per fiber revolution $1/(2\pi)$ within matter — three coupled numbers from one fibration, with the absolute fractions read at the crossing epoch (§32.11).

The answer to (a) is therefore: $\Omega_m = 1/\pi$ is the ratio of the matter Hopf-fiber volume to the total electroweak-holonomy volume, $\mathrm{Vol}(S^1)/\mathrm{Vol}(S^3) = 2\pi/2\pi^2$; the competing rational functions of π are excluded by the geometry of the boundary measure, with no reference to the measured value and no continuous parameter free. The formal statement, and the one measure-normalization lemma it still rests on, are in Section 74.3.

74.2 Why is the cosmic coincidence not a coincidence?

The traditional cosmic-coincidence problem of ΛCDM asks: why does the present epoch happen to be the unique cosmic time at which the matter density and the cosmological-constant density are comparable? Under a free Λ-and-Ω_m hypothesis, the comparable-magnitude condition holds only for a fleeting fraction of the cosmic history — typically a few e-folds out of the eighty between

recombination and the asymptotic de Sitter future — and the apparent agreement between Ω_m and Ω_Λ at the present epoch becomes an anthropic-or-coincidence puzzle.

The answer dissolves this puzzle in three steps, each a corollary of the π-locked-inventory theorem.

First, the matter fraction and the cosmological-constant fraction are not two independent observables on the cosmological-boundary stratum: they are bound by the topological closure $\Omega_m + \Omega_\Lambda = 1$ (modulo small radiation and curvature contributions). This is unit-trace closure, not free parameter choice. The comparable-magnitude condition at the present epoch is therefore the *generic* state of the cosmological-boundary attractor, not a coincidence about *when* we are observing.

Second, the specific ratio $\Omega_m/\Omega_\Lambda = (1/\pi)/(1 - 1/\pi) = 1/(\pi - 1) \approx 0.467$ at the present epoch is forced by the same topology that fixes each fraction individually. This is a numerical theorem, not an empirical relation.

Third, the temporal coincidence ("why are we here at this particular cosmic age?") is dissolved by a structural integer identification on two independent epoch-bookkeeping channels. The recombination switch event sits at $z_* = e^J = e^7$ (Theorem 68.8), and the e-fold count from the corridor switch to the present epoch is $\ln(1 + z_*) \approx J = 7$ as well; both channels are governed by the same OGN-5 integer $J = 7$ of GG-A4-Euclid. The temporal coincidence ("why is the present epoch at this particular cosmic age?") is therefore dissolved by the integer identification: the same integer $J = 7$ sets both the recombination scale and the recombination-to-present e-fold separation, so "now" is not a free choice of cosmic time but the natural endpoint of the integer-counted descent from the corridor switch.

The answer to (b) is therefore: the cosmic coincidence is not a coincidence because (i) the matter-vacuum split is topologically locked, (ii) the specific ratio is a numerical theorem, and (iii) the temporal "now" is integer-counted from the recombination switch event — each a corollary of the universal-identity theorem.

74.3 The closed-packet four-form integration: $\Omega_m = 1/\pi$ as a topological theorem

This subsection states the formal theorem behind §74.1 in canonical form, with hypotheses, conclusion, and a five-step proof outline anchored on the BI-6 four-form of GG-A3-BI6 and the OGN-5 integer ledger of GG-A4-Euclid.

Setup. On the cosmological-boundary stratum — the four-dimensional boundary slice obtained from the six-dimensional GG-6 bulk of GG-A2-GG6 via the OGGEP projection of GG-A5-GDOS — the BI-6 modified Bianchi identity of GG-A3-BI6 fixes a closed four-form

$$X_4 = \operatorname{tr} R^2 \;-\; \operatorname{Tr} F_5^2 \;-\; \tfrac{2}{5} F_B^2 \;-\; \tfrac{2}{5} F_{Qi}^2,$$

with the coefficient vector $(1, -1, -2/5, -2/5)$ proved unique by GG-A3-BI6's anomaly-cancellation theorem. The cosmological matter density fraction Ω_m at the present epoch is operationally the ratio of the matter-sourced contribution to the total stress-energy budget, projected through OGGEP onto the four-dimensional cosmological-boundary slice.

Theorem 74.1 ($\Omega_m = 1/\pi$ as the Hopf fiber-volume ratio, conditional (T-Cosmos-17))**.** *Assume:*

(I) *The cosmological-boundary stratum is a connected, orientable four-manifold $\mathcal{M}_4$ obtained as the OGGEP projection of the GG-6 bulk (Arisaka, 2026, A2-GG6; Arisaka, 2026, A5-GDOS).*

(II) *BI-6 holds throughout the cosmological history: $dX_4 = 0$ on $\mathcal{M}_4$, with the coefficient vector $(1, -1, -2/5, -2/5)$ (Arisaka, 2026, A3-BI6).*

(III) *The OGN-5 integer $N = N_c = 3$ (colors) and $L = 5$ (SU(5)) fix the boundary holonomy manifold as the electroweak $\mathrm{SU}(2) \cong S^3$, and the matter sector's charge as its Hopf fiber S^1 (VIII-b″, the unit-packet normalization).*

Then the present-epoch matter density fraction equals

$$\boxed{\Omega_m \;=\; \frac{1}{\pi} \;\approx\; 0.3183}$$

as the Hopf fiber-volume ratio $\mathrm{Vol}(S^1)/\mathrm{Vol}(S^3)$ of the boundary holonomy, conditional on the measure-normalization lemma VIII-b″ (Conjecture VIII), with no continuous adjustable parameter and no reference to the measured inventory.

Epoch-identity reading. The conditionality on Conjecture VIII localizes further: with VIII-b′ proved and VIII-b″/(⋆) closed as the epoch identity (see the Conjecture-VIII refinement paragraph), the theorem may equivalently be read as *unconditional in form at the crossing epoch t_π*: at the instant when one packet's diluted ledger record equals one unit's equipartition rate, the inventory reads the coupled π-form $(1/\pi,\ 1 - 1/\pi,\ 1/(2\pi^2))$. The residual hypotheses are then VIII-a (uniformity) and the identification $t_0 = t_\pi$ (the 1% L4 epoch input of §32.11).

Proof outline. The proof proceeds in five steps. The first three establish the closed-packet topology; the fourth performs the integration; the fifth rules out competing rational functions of π by the constitutional cost-minimum principle.

Step 1 (Closure of X_4). By hypothesis (II), $dX_4 = 0$ on $\mathcal{M}_4$. The four-form X_4 therefore represents a class in the fourth de Rham cohomology $H^4(\mathcal{M}_4, \mathbb{R})$. The class is non-trivial because at least one of the four trace-squared terms (specifically $\mathrm{tr}\, R^2$) generates a non-trivial Pontryagin class on the curved cosmological-boundary slice.

Step 2 (Closed-packet structure: the integer skeleton). By hypothesis (III) and the topological-cell decomposition of GG-A4-Euclid, the OGN-5 tuple $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ fixes the *integer* content of the packet: the color integer $N = N_c = 3$ and the gauge integer $L = 5$ identify the boundary holonomy as the electroweak SU(2), and the packet is one canonical unit (the Hirzebruch minimum $\sigma = 1$ of Step 3). This step supplies the discrete skeleton — *which* manifold and *one* packet on it — but not yet the share; the share is a fiber volume, computed in Step 4. *(One might instead assign the matter cell the rational weight $M/(LN) = 1/15$ and seek an epoch-independent arithmetic cancellation to reach $1/\pi$; that route is a provable dead end — Appendix L.8 — and the fiber-measure evaluation below replaces it.)*

Step 3 (Pontryagin normalization). The Chern–Pontryagin integral of $\operatorname{tr} R^2$ over a closed orientable four-manifold satisfies

$$\int_{\mathcal{M}_4} \operatorname{tr} R^2 \;=\; 8\pi^2 \cdot p_1[\mathcal{M}_4],$$

where $p_1[\mathcal{M}_4]$ is the integer-valued first Pontryagin number. For the OGN-5-structured cosmological-boundary slice with $p_1 = N = 3$, the integral evaluates to $24\pi^2$. The matter-cell contribution carries this normalization factor when projected through OGGEP onto the present-epoch boundary. *(Normalization note: the $8\pi^2$ is consumed upstream — the GEP functional of GG-A6-GRIP is defined as $(1/8\pi^2)\int X_4$, so the packet's energy content is the pure integer p_1, and* no *residual π descends from this step. The π of $1/\pi$ is therefore* not *a leftover of the Pontryagin factor and* not *produced by Lemma VIII-b′ — whose content is the reality factor of exactly 2 (the $2\pi^2$ vs π half-measure of a charged sector against a real one), not a π. The π enters as a* volume *in the denominator of the fiber-measure ratio of Step 4. By the Hirzebruch signature theorem $p_1 = 3\sigma$, the value $p_1 = N = 3$ is $\sigma = 1$: the minimal closed gravitational topology, which is what forces the packet to be one canonical unit.)*

Step 4 (The fiber-measure evaluation: the Hopf ratio). The matter fraction is not an epoch-independent *arithmetic* of the OGN-5 integers — the route that seeks one is a provable dead end (Appendix L.8, Step 4, evaluating to the failed $3/(125\pi)$; the dilution no-go of the Conjecture-VIII audit shows no constant combination can close it). It is a ratio of *fiber volumes* on the boundary holonomy, and this is where the π is born. The matter sector's charge lives on the compact phase circle S^1_θ, which sits inside the electroweak $\mathrm{SU}(2) \cong S^3$ holonomy manifold as its *Hopf fiber* $S^1 \hookrightarrow S^3 \to S^2$ (a canonical fibration, with no freedom of choice). Under the frozen stationary GDOS measure of Lemma VIII-a, each sector's present-epoch share is its fiber volume divided by the total holonomy volume, so

$$\boxed{\Omega_m \;=\; \frac{\mathrm{Vol}(S^1)}{\mathrm{Vol}(S^3)} \;=\; \frac{2\pi}{2\pi^2} \;=\; \frac{1}{\pi}} \qquad \text{(unit-normalized spheres)},$$

exact. The reality factor of Lemma VIII-b′ (charged measure 2π against real measure π, ratio 2) then supplies the extra factor in the baryon line, and the crossing-epoch identity $u = 1$ (VIII-b″, discharged as $t_0 = t_\pi$ to 1%, the declared L4 input of §32.11) fixes *when* the running share equals this geometric ratio — the "when," not the "what." The value $1/\pi$ is thus read off the Hopf geometry, with the data nowhere in the selection. GG-A4-Euclid supplies the integer that pins the packet to one unit; the fiber-volume identification is the content of the remaining measure lemma (VIII-b″), stated below.

Step 5 (Uniqueness: the competitors die by geometry, not by data). The decisive test of non-circularity is whether the alternative rational functions of π fall *without* appeal to the measured inventory. They do:

(a) $1/(3\pi)$ would require a total holonomy volume $3\,\mathrm{Vol}(S^3) = 6\pi^2$. No boundary-holonomy manifold in the construction has that volume — there is no "triple S^3." Excluded by geometry.

(b) $1/(2\pi) = 1/\mathrm{Vol}(S^1)$ is the *baryonic* projection (one integer unit per revolution of the fiber), i.e. the Ω_b/Ω_m ratio, not the matter/total fiber ratio. Assigning it to Ω_m mistakes the baryon fiber for the matter fiber. Excluded structurally.

(c) $1/\pi^2$ would require $\mathrm{Vol}(S^1)/\mathrm{Vol}(S^5) = 2\pi/\pi^3 = 2/\pi^2$ (off by a factor of 2), and there is no S^5 holonomy on the four-dimensional boundary in any case. Excluded by geometry.

Every competitor that an empirical argument would eliminate with "the data prefer $\Omega_m/\Omega_\Lambda \approx 0.46$" is eliminated here by "there is no fiber of that volume." The selector is the fiber-volume geometry, and the measured value of the inventory enters the derivation at no step.

Proof status (honesty note): where the conditional lives. The inventory is carried at grade D*, and it is worth saying exactly which link of the chain the star is attached to — because it is not the link the earlier phrasing of this Section named. The derivation has four links, and they do not have the same status.

1. **The assignment.** That the boundary holonomy manifold is the electroweak $\mathrm{SU}(2) \cong S^3$, and that the matter sector's charge lives on its Hopf fiber S^1. Both are fixed upstream by the OGN-5 integers $N = 3$ and $L = 5$, *before* any fraction is computed, so the pair is not chosen for its ratio: the ratio is what the prior assignment delivers. *Derived, upstream.*

2. **VIII-a (uniformity).** That the frozen GDOS measure leaves every counted direction delocalized. Proved in consumer form with its three residues discharged; exact up to the gravity-mediated $\sim e^{-K}$ floor, whose $O(1)$ prefactor is the one open coefficient and is load-bearing for nothing deployed. *Proved.*

3. **VIII-b′ (the reality half-measure).** Exact mode counting — $2\Lambda + 1$ independent real degrees of freedom for a reality-constrained sector against $2(2\Lambda + 1)$ for a holonomy-charged one, a ratio of exactly $1/2$ at every finite cutoff. *Proved.*

4. **VIII-b″ (unit-packet normalization).** One closed packet = one unit of fiber measure. *Discharged — but not into a statement about a constant.*

Why links 3 and 4 have different status, and it is not an accident of effort. Lemma VIII-b′ compares two *counts*. A ratio of counts is dimensionless and cutoff-independent, so it can be proved outright, and it is. Lemma VIII-b″ compares a length to a volume: at fiber radius R,

$$\frac{\mathrm{Vol}(S^1)}{\mathrm{Vol}(S^3)} \;=\; \frac{2\pi R}{2\pi^2 R^3} \;=\; \frac{1}{\pi R^2}, \tag{421}$$

which equals $1/\pi$ only at unit radius. *A ratio that carries a dimension cannot be fixed by geometry alone; it needs a physical unit.* That is the entire content of residue **(R2)**, and it is why VIII-b″ required a unit bridge while VIII-b′ did not. Naming the two lemmas alike concealed a real difference in kind.

And the unit bridge terminates on a theorem, not on a gap. Superselection freezes the packet's integer charge, not its energy record, so the packet energy tracks the diluting matter sector ($\propto a^{-3}$) while the equipartition rate $T_{\mathsf{GFF}}/2\pi$ is frozen. Hence the *dilution no-go*: no epoch-independent constant can equal a diluting quantity at more than one instant. VIII-b″ therefore *cannot* close as a statement about a constant — and that is a result, not an omission. What it closes into is the occupancy identity: $u(t) = \pi\,\Omega_m(t)$ sweeps monotonically through unity exactly once, and $u = 1$ is *mathematically equivalent* to $\Omega_{m,0} = 1/\pi$.

$$\boxed{\text{The form is derived. The epoch is calibrated.}} \tag{422}$$

Concretely: the coupled π-form $(1/\pi,\ 1-1/\pi,\ 1/(2\pi^2))$ follows from links 1–3 with no measured density entering at any step, and the competitors are excluded geometrically rather than empirically — $1/(3\pi)$ would need a holonomy volume $6\pi^2$ that no manifold in the construction has, $1/\pi^2$ an S^5 the four-dimensional boundary does not carry. That is a derivation, and at the crossing epoch the theorem is *unconditional in form*. What the geometry does *not* fix is *when* the running inventory reads that form. The theorem supplies a unique crossing epoch t_π; the identification $t_0 = t_\pi$, good to 1%, is the declared L4 epoch input of §32.11. **It is that identification, and nothing else, that makes the present-epoch values** D*.

The calibration is a genuine one and its width should be stated with it. The crossing sits at $a_\pi = 0.9949$, about 75 Myr before the present at $H_0 = 66.6$ (the two-road reading); moving Ω_m by $\pm1\sigma$ of the Planck uncertainty moves the crossing from roughly 90 Myr in the *future* to roughly 230 Myr in the past. The ±160 Myr blur is therefore the real precision of the epoch statement, and the present offset is 0.4σ inside it — entirely consistent with the crossing being now, and not more precise than that.

One row of the inventory does not carry the star, and should not be tagged as though it did. The epoch input is the only thing standing between the inventory and grade D. It follows that any quantity in the inventory which does not depend on the epoch is not conditional on it — and there is exactly one:

$$\frac{\Omega_b}{\Omega_m} \;=\; \frac{1/(2\pi^2)}{1/\pi} \;=\; \frac{1}{2\pi} \;=\; 0.159155. \tag{423}$$

Baryons and cold dark matter both dilute as a^{-3}, so this ratio is *constant at every epoch* after the baryons become non-relativistic; and the critical density cancels between numerator and denominator, so it is h-free as well. It therefore consumes no epoch calibration; it keeps the star in the codebook only because the π-locked fractions it is built from are carried under Conjecture VIII. The absolute fractions Ω_m, Ω_Λ, Ω_b keep the star, because each of them sweeps with the epoch and the geometry fixes only *where* the sweep crosses, not *when*.

A bookkeeping point that matters at this precision. The comparison must be made in ω_b/ω_m and *not* in Ω_b/Ω_m, because the measured $\Omega_m = 0.3153$ includes a neutrino component that this framework's $\Omega_m = \Omega_b + \Omega_{\rm DM}$ does not. With $\Sigma m_\nu \simeq 59$ meV the neutrino share is $\Omega_\nu = 0.0014$, that is 0.44% of Ω_m, and the two comparisons differ by half a percentage point: $\omega_b/(\omega_b+\omega_c) = 0.02237/0.14237 = 0.15713$ gives +1.29%, while $\omega_b/\omega_m = 0.02237/0.1430 = 0.15643$ gives +1.74%.

Against a signal of that size the difference is not negligible; the codebook row uses the latter, the neutrino-inclusive ω_m that every other matter row is scored against, and the half point is stated here as what the framework's neutrino-free Ω_m would recover. The same neutrino component means the constancy of the ratio is itself exact only to $\sim 0.4\%$, since neutrinos were relativistic at earlier times; that is below the deviation being quoted, and it is stated rather than assumed.

What the epoch can and cannot move. Because the epoch is a declared input, it is worth recording precisely which quantities it could in principle be used to adjust — and which it could not, so that the maneuver is not attempted later and a sign flip mistaken for a repair.

- $\eta_B = n_B/n_\gamma$ is a *constant of the universe*: both number densities dilute as a^{-3} after electron–positron annihilation, so η_B takes the same value at every epoch. No choice of evaluation time can move it, and no discrepancy in it can be closed that way.
- Ω_b/Ω_m is likewise constant, by Eq. (423).
- $\rho_b \propto a^{-3}$ *is* epoch-dependent.
- Ω_b, Ω_m, Ω_Λ are epoch-dependent, because $\rho_{\rm crit}$ evolves.

What does change with the evaluation point is the *conversion* between Ω_b and η_B. Applying $\Omega_b = 1/(2\pi^2)$ at a_π rather than today requires converting with $\rho_{\rm crit}(a_\pi)$, a factor $E^2(a_\pi)\, a_\pi^3 = 0.99054$, or -0.95%; at this paper's own $h = 0.666$ that moves the ledger road from $\eta_{10} = 6.15$ ($+0.57\%$) to 6.10 (-0.38%). It flips the sign of the residual and preserves its size. It does not close anything.

And it is not an available adjustment in any case. That factor exists only because a_π was located using the *measured* $\Omega_m = 0.3153$; if $\Omega_m = 1/\pi$, then $a_\pi = 1$ identically and the factor is exactly unity. The correction is nothing but the -0.95% by which the measured matter fraction differs from $1/\pi$, restated in other units — so using it to adjust η_B would spend the same discrepancy twice.

The one genuinely epoch-dependent baryon quantity, ρ_b, moves the wrong way: at a_π it is $+1.40\%$ larger than today, so comparing $\rho_b(t_\pi)$ against the measured present-day value would take the agreement from $+0.45\%$ to $+1.86\%$. Evaluating at the crossing epoch *opens* that gap.

The selection is not circular. A derivation that selected $1/\pi$ by the empirical preference $\Omega_m/\Omega_\Lambda \approx 0.46$ would be; this one is not: Steps 4 and 5 select $1/\pi$ by the Hopf fiber-volume geometry, and the data enter the derivation nowhere. Two residues were named in earlier statements of this note, and both are worth keeping — not as outstanding debts, but because they say what the identification of the epoch actually asserts: **(R1)** that the frozen GDOS measure weights exactly the Hopf circle of the electroweak SU(2) as the matter fiber and $\mathrm{Vol}(S^3)$ as the total — the content of VIII-b″ (one closed packet = one unit of fiber measure), whose derivation route through the $T_{\rm GFF} \sim M_{\rm KK}$ unit bridge is sketched in the codebook audit (§68.2.3); and **(R2)** that the fibers carry the canonical unit normalization (S^1 circumference 2π, S^3 volume $2\pi^2$) that the measure convention must pin. Both are geometric identifications, provable or falsifiable on the bulk construction, with no reference to any measured density. The value $1/\pi$ is therefore carried as *derived in form and calibrated in epoch*: the codebook lists $\Omega_m, \Omega_\Lambda, \Omega_b$ as D* because of the epoch identification $t_0 = t_\pi$,

and for no other reason. Concretely: (R1) is discharged by the fiber assignment of link 1 together with the occupancy identity, and (R2) is the dimensional statement of Eq. (421), which the unit bridge answers at the crossing and cannot answer away from it.

Corollary (the π-locked inventory). The same closed-packet integration gives $\Omega_\Lambda = 1 - 1/\pi$ by unit-trace complementarity and $\Omega_b = 1/(2\pi^2)$ by a further $U(1)_B$ projection through the baryonic-gauge protection of GG-A9-SM. Together these complete the π-locked inventory predicted in the codebook of Section 68.1. The Kaxion fraction $\Omega_{\rm DM} = \Omega_m - \Omega_b = 1/\pi - 1/(2\pi^2) \approx 0.2676$ follows from the closed-form subtraction.

Status note. Theorem 74.1 states the geometric form of the claim: *if* the frozen GDOS measure weights the Hopf circle of the electroweak SU(2) as the matter fiber (VIII-b″), then $\Omega_m = \mathrm{Vol}(S^1)/\mathrm{Vol}(S^3) = 1/\pi$ rather than any competing rational function of π, with the competitors excluded by geometry. As §74.3 records, the theorem closes the *form* and not the *epoch*: links 1–3 are derived or proved, VIII-b″ discharges into the occupancy identity rather than into a statement about a constant, and what remains is the identification $t_0 = t_\pi$ to 1%. Ω_m, Ω_Λ and Ω_b are therefore D* *with the epoch as their declared calibration input.* The earlier reading — that the residual was an unexhibited normalization lemma — is superseded.

So the promise and the proof can now be separated cleanly. $\Omega_m = 1/\pi$ is not a numerological observation: it is the Hopf fiber-volume ratio of the electroweak holonomy, forced by geometry and selected without the data. What is promised and not proved is narrower and better posed than the whole — that the epoch at which the running inventory reads that ratio is the present one, which is the coincidence question in its sharpest available form.

The view from one floor up. The questions of this Section are the three operational forms of one deeper question: why does the universe admit a closed analytic description at all? The answer is that the cosmological-boundary stratum is the unique stable attractor of the boundary moduli space, and the π-locked inventory is the projection of that attractor onto the present epoch, with the layer-by-layer identity ladder of §71.4 — from the electron to the cosmos — as its content. After all, the question *why the universe has the inventory it does* is the cosmological-side rephrasing of the constitutional question that GG-Theory was built to answer: not what the universe is, but why it is the unique thing it is.

75 Two fundamental questions, one geometric answer

> *"Nature is so beautiful and so well designed. Why is the universe so well-designed for us to be here in the beautiful biosphere on Earth? This question was the primary motivation for my decision to become a physicist, way back in my high school years."*
>
> — K. Arisaka, May 2026.

GG-A9-SM's Part XI closed the Standard-Model audit not on a parameter but on a pair of questions — why every electron is the same, and why exactly four forces exist with exactly their

behaviors — and recognized the two answers as one uniqueness theorem read at two layers. The present Part has reached the same summit from the cosmological side. Section 71 answered the cosmological form of the first question: why the vacuum energy, the Kaxion, and the π-locked inventory are the same everywhere. Section 74 answered the second: why the inventory takes the particular values it does, and why their coincidence at the present epoch is no coincidence. This closing Section recognizes that these are not two cosmological answers but one, and that the "one" is the same one GG-A9-SM found — a uniqueness principle on a geometric stratum, now read on the largest stratum the program reaches.

75.1 The structural symmetry of the two cosmological questions

In each case the form of the argument is identical. We assemble a small number of separately-named constitutional inputs; we show that those inputs leave exactly one admissible readout; we identify the boundary readout with what we observe. For the identity question the inputs were the GRIP cascade of GG-A6-GRIP and the ground-state refinement; the unique readout was the cosmological GGA (dark energy) with its relic and baryonic records; the boundary observation was the universal inventory of every sightline. For the value question the inputs were the BI-6 four-form of GG-A3-BI6 and the OGN-5 integer ledger of GG-A4-Euclid; the unique readout was the closed-packet integral projecting to $(1/\pi,\, 1-1/\pi,\, 1/2\pi^2)$; the boundary observation was the measured density parameters. *Different inputs, one stratum, the same structural form* — and the same form, in turn, as GG-A9-SM's two Standard-Model questions one floor below. The recognition that identity and value are instances of one geometric theorem applied twice is the simplest unifying statement available about the cosmological sector, and it is the exact cosmological echo of GG-A9-SM's "two questions, one answer." What GG-A9-SM showed for the electron and the four forces, this Part shows for the vacuum and its inventory: the two deepest questions one can ask about the contents of the universe turn out to be one question, asked twice.

75.2 Dissolving the three classical alternatives

Three classical positions have tried to explain why the universe is hospitable, and each has fastened, historically, on precisely the cosmological numbers this Part derives — the smallness of Λ, the matter–vacuum balance, the fitness of the expansion history for structure and life. The framework does not refute these positions so much as it dissolves the question each was asked to answer.

Design. The design argument holds that a designer fine-tuned the cosmological constant and the matter budget for the existence of life; its standing objection is regress — who designed the designer, and by what tuning? In the present framework the regress is dissolved because no tuning occurred. The constitution of geometry permits one self-consistent readout at each layer; nothing was chosen because nothing was available to choose from. A designer who had to set $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_m = 1/\pi$, the Kaxion mass on a rung of the cap-loading ladder, and the seesaw $m_\Lambda = M_{\rm lock}^2/M_{\rm Pl,red}$ would be acting under constraint — these are the readouts of topological and Diophantine closures, not free dials — and a designer who acts under constraint has chosen nothing. Design is not refuted; it is

rendered superfluous. The 10^{120} vacuum-energy fine-tuning that has haunted cosmology for half a century is the sharpest case: it is not a dial set improbably close to zero, but the depth of the one terminus basin, read out.

Multiverse. The anthropic-multiverse argument holds that we observe a hospitable universe because we could not observe an inhospitable one, and that the inhospitable universes populate a landscape from which ours is the lucky draw; its standing objection is the absence of independent evidence for the other universes and the resulting non-falsifiability of the selection. In the present framework the absence of evidence is dissolved by the absence of the other universes: the constitution permits one self-consistent bulk, one gauge group, one breaking pattern, one stable terminus on the cosmological stratum. There is no landscape of vacua to draw Λ from — not because we cannot survey it, but because the geometry admits no other value. The cosmological-constant "landscape" of 10^{500} vacua, invoked precisely to make the observed Λ an anthropic draw, is replaced by a single computed depth.

Brute fact. The brute-fact position holds that the cosmological numbers are simply what they are and that asking why is unscientific; its objection is that it abandons explanation at the moment explanation becomes most pressing. In the present framework the abandonment is unnecessary. The inventory fractions are powers of π read off a closed-packet integral, the Kaxion mass is one calibrated scale read on an integer ladder, the vacuum scale is a seesaw of the lock depth, and the expansion history is the descent of one GRIP cascade — each either measured or testable this decade (the tensor ratio $r = 0.0038$, the Kaxion line near 1.95 GHz, the siren H_0 near 67). Asking why has become scientific again because there is now a path on which a falsifiable answer can be carried.

75.3 The self-referential closure

The deepest structural feature of the framework is not any of the three dissolutions above but the recognition that the geometry which forces the cascade contains, among its forced boundary readouts, the existence of intelligent observers — and that those observers eventually formulate the constitutional clauses that describe the geometry that forced them. The cosmological setting makes the loop unusually vivid: the same descent whose terminus is the dark energy we measure is the descent that, at an earlier basin, stocked the universe with the baryons that became stars, planets, chemistry, and finally the astronomers who read the terminus off the sky. The universe contains, within itself, the means of recognizing its own constitution — and the reading is done from inside the very inventory being explained.

This recognition has resonances in older traditions, named here without overclaiming what they say. Wheeler called the closed loop in which observers participate in the physics that produced them the *participatory universe*. The Upanishadic tradition gives the same image as *tat tvam asi* — the recognition by the individual self of its identity with the underlying reality, an act in which the recognition completes the structure that made it possible. Hegel called the same dynamic *Geist* becoming conscious of itself in history. Spinoza called the most general form of the closure *necessitarianism*: what exists, exists by necessity, because the nature of substance permits

no alternative. The GG-Theory version is austere and sharp: the geometry permits exactly one self-consistent loop, and the loop closes when minds within the universe write down the constitution that produced them. The framework reproduces none of the metaphysical content of these traditions; it claims only that the structural form of the closure — one geometry, one cascade, one set of boundary observers, one act of self-recognition — is shared. What distinguishes the present version from its predecessors is not the conclusion but the warrant: Spinoza inferred necessity from abstract reasoning about substance, open to dispute at every step; GG-Theory derives the analogous conclusion from named theorems with explicit constitutional inputs, integer ledgers, and falsifiable consequences testable within the decade.

75.4 A careful word about determinism and freedom

The framework does not say that every micro-event of cosmic history is preordained. It says that the qualitative *layered structure* of cosmic history is forced. Within each layer there is enormous combinatorial freedom: not every galaxy has been placed, not every star numbered, not every world's chemistry run to its end. The determinism is at the level of which layers exist and what their attractor coordinates are — that there is inflation, that there is a matter era, that there is a dark-energy terminus, that the inventory is π-locked — not at the level of what happens within each layer. Cosmic structure, biological evolution, contingent history all take place inside layers the geometry has forced into existence; they are not preprogrammed by the geometry, only sheltered by it.

This is not a small caveat. Accounts of cosmic fine-tuning have traditionally forced a choice between an indifferent randomness (the multiverse draw) and a rigid design (the tuned constants). The framework offers a third position: the layered structure is geometrically forced, while within each layer there is genuine combinatorial freedom of a kind that is neither random nor preordained. The universe we inhabit is not the lucky outcome of a draw, nor the fixed output of a plan; it is the unique-readout solution at each layer of the geometric cascade, and within that solution the contingent richness of cosmic and terrestrial history — including, evidently, the freedom to ask the questions this Part has answered — plays out in the room the geometry has built.

75.5 The closing punch line

We close the philosophical summit of the Part by stating the synthesis in one paragraph.

> *After all, the deepest structural fact about the universe may be this: the cosmos is beautiful and well-suited to life not because a designer chose its inventory, and not because we are lucky to inhabit a hospitable draw, but because the constitution of geometry permits exactly one self-consistent readout at each layer, and that readout — whether read at the layer of a single electron or at the layer of the whole vacuum — is what we see when we look. The two fundamental questions of cosmology, why the inventory is universal and why it takes the values it does, turn out to be one question, asked twice; and it is the same question GG-A9-SM asked of the electron, asked now of the universe that the electron is a citizen of.*

The meta-question is answered as far as the cosmological parameter sheet allows. Before the wider discussion, two Parts close the case. The next assembles the harvest: the direct derivation of ΛCDM, stage by stage from the bulk action to the expansion history, with the register of every entry printed beside it. The one after gathers everything onto a single sheet: the cross-paper coherence map, the comprehensive first-order codebook, and the second-order residual audit that let a skeptical reader verify the chain theorem by theorem.

76 Part X in plain English: the deepest claim in the paper, and the conjecture it rests on

This is the Part the rest of the paper builds toward, and it is also the Part most exposed. The claim is that the cosmic inventory has the form it has for the same reason every electron in the universe is identical to every other: one geometry, one deepest basin, and no ensemble to select from.

Stated that way the claim is attractive, and attractiveness is a reason for caution rather than confidence. What the Part actually establishes is conditional. The topological argument for the π-locked inventory rests on a named vacuum-uniqueness conjecture that the wider program already depends on; it does not close that conjecture, and a reader should understand the result as two open puzzles unified rather than one puzzle solved. That is a real gain — unification reduces the number of independent things that must be true — and it is a smaller gain than a proof.

The second and third horizon problems introduced here are, as far as this paper knows, new formulations, and new formulations are cheap. Their value is that they are answerable by the same mechanism that answers the first, and that they make the sameness of things and the sameness of choices into questions of the same type as the sameness of temperature. If the mechanism is wrong, all three fail together, which is the property a reader should want from a unification and should check for before believing one.

Part XI
The assembled ΛCDM: the standard cosmological model as a constitutional export

The direct derivation of ΛCDM by GG-Theory, assembled in one place — the concordance model written in its standard form, the zeroth and the six silent axioms behind it replaced by theorems and graded conditionals, the export chain walked in eight stages from the bulk action to the expansion history, the assembled model displayed with the register of every entry stated and no fitted continuous parameter, and the honesty boundary printed beside the result

77 The model the reader was taught to fit

Every cosmologist learned the subject the same way. The Friedmann equations were written on the board, general relativity and the FRW symmetry assumed; a source menu was populated with whatever the data required — radiation, baryons, cold dark matter, a cosmological constant; six numbers were fitted ($\Omega_b h^2$, $\Omega_c h^2$, θ_*, A_s, n_s, τ); and everything afterwards was consequence. Fifty years of data have confirmed the fit to extraordinary precision — and explained none of its inputs. ΛCDM is the cosmological twin of the Standard-Model Lagrangian: the most successful phenomenological object in its field, and the object its own field cannot derive. The companion GG-A9-SM (Arisaka, 2026, A9-SM) closes that circle for the Lagrangian, assembling it term by term as a constitutional export; the present Part does the same for the concordance model. The single governing sentence is the same there and here: *the visible model is exported from the six-dimensional bulk, not assembled by fitting.* Everything below is a harvest — the stage of Part III, the launch and terminus of Part IV, the matter of Part V, the corridor of Part VI, the chronology of Part VIII, the graded codebook of Part IX, the topology of Part X, and the course-pace walk of Appendix E — gathered into one chain with the register of every link printed. Nothing in this Part is derived twice; every stage names its home.

78 The zeroth and the six silent axioms of ΛCDM

To see what is at stake, it helps to state precisely what the textbook assumes when it writes down ΛCDM. Six structural choices are made silently before the first fit — and beneath them sits a zeroth, deeper one, so habitual it is rarely recognized as a choice at all: that the starting point is a *four-dimensional FRW effective description with a fitted source menu* in the first place. The six axioms fill in the menu; the zeroth supplies the blank form. Table 44 is the scorecard, and the claim assembled below is that the descent replaces the zeroth and all six at once — the same kind of claim, made at the same register, as the assembled Lagrangian of GG-A9-SM.

Table 44: The zeroth and the six silent structural axioms behind the textbook ΛCDM, and the GG-Theory replacement of each. Registers: **D** = Derived, **D*** = Derived (conditional on a named lemma or conjecture), per-entry grades in the codebook of Part IX.

#	Silent textbook axiom	GG-Theory replacement	Register
S0	A 4D FRW effective description with a fitted source menu is the starting point at all	The 6D bulk and its boundary projection: the earned bulk action, the boundary Einstein equation as theorem, the ledgered source (Appendix E, Steps 1 and 3)	**D** (form); **D*** (content)
S1	General relativity as the gravitational law, with G measured	The boundary Einstein equation derived by variation and zero-mode projection; $M_{\rm Pl}^2 = \pi M_6^4(1-e^{-2K})$ computed; conservation a Ward identity; correction bounded with live falsifiers (§E.6)	**D**; bound **D***
S2	Homogeneity, isotropy, and flatness (the FRW symmetry)	The derived plateau launch: $N^* = 52.6$ e-folds stretch one causal patch and drive $\lvert\Omega_{\rm tot}-1\rvert \to 0$ (Part IV; Appendix G)	**D***
S3	The component menu: cold dark matter exists, Λ exists, $w=-1$	Kaxion dark matter from the compact angle (8.1 μeV, one scale calibrated); Λ the seesaw terminus ($M_{\rm lock}^2/M_{\rm Pl,red} \approx 6.1$ meV) with $w=-1$ exact; no fifth component exported (Parts IV–V; Arisaka, 2026, C1-Kaxion)	**D***
S4	The proportions $\{\Omega_i\}$, read from the sky	The closed-packet fractions $\Omega_m = 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_\Lambda = 1-1/\pi$ (Part X), carried on Conjecture VIII with the epoch identity printed	**D***
S5	The initial conditions: adiabatic, nearly scale-invariant fluctuations of amplitude A_s	The two-basin descent: $n_s = 0.9634$ (**D***), $r = 0.0038$ (**D***), A_s the launch sector's one declared input (Part IV)	mixed, stated
S6	The six fitted numbers	The codebook of Part IX, graded entry by entry; τ late-time astrophysics, outside the fundamental claim	per triage

Classic alternatives replace pieces: quintessence trades Λ for a potential with new dials; MOND trades dark matter for a modified force law; the multiverse trades the proportions for a selection effect. The claim of this Part is structurally different — all seven axioms replaced at once, by exports of one geometry, with the register of each replacement stated rather than inflated.

79 The export chain in eight stages

The assembly rests on eight stages, C-1 to C-8, each proved or graded elsewhere in the paper or the cohort; the stage tags reappear in the assembled display of §80.

79.1 C-1 — the bulk action and the ledger

The chain begins where the companion papers begin: the six-dimensional bulk $M_4 \times I_r \times S^1_\theta$ with its canonical warped metric, forced by the constitutional clauses (T-GG6-0/1/2 of GG-A2-GG6 (Arisaka, 2026, A2-GG6)); the unique admissible two-derivative bulk action, earned brace by brace in the assembly interlude of GG-A9-SM (Arisaka, 2026, A9-SM) (station S-1a) and displayed at course pace in §E.3; and the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, the unique solution of the BI-6 closure of GG-A3-BI6 (Arisaka, 2026, A3-BI6). At the end of C-1 the theory holds five integers, five dimensional constants, two transcendentals (e radial, π angular) — and zero continuous dials. Every subsequent stage spends from this account only.

79.2 C-2 — the boundary Einstein equation

Varying the action and projecting onto the graviton zero mode yields four-dimensional gravity as a theorem: the boundary Einstein equation with $M_{\rm Pl}^2 = \pi M_6^4 (1 - e^{-2K})$, conservation inherited as the Ward identity of bulk diffeomorphism invariance, the open-system correction bounded at 10^{-91} today, and the modified-TOV and growth-sector falsifiers keeping the bound honest — Theorems T-GDOS-8.1–8.6 of GG-A5-GDOS (Arisaka, 2026, A5-GDOS), walked at first-year-graduate pace in Step 1 of Appendix E (§E.4–E.6). The zeroth and first silent axioms are discharged here: the law of cosmology's left-hand side is not postulated.

79.3 C-3 — the FRW symmetry and the Friedmann equations

The FRW form of the boundary metric is not an aesthetic choice: the launch of C-6 stretches a single causal patch across the observable sky, and the runtime selects the symmetric attractor (Part III). With the symmetry in place, the two Friedmann equations follow from the boundary Einstein equation by the exact textbook computation — deliberately unchanged, Step 2 of Appendix E — with the acceleration equation and continuity carrying only two independent pieces of content between them, the redundancy being C-2's Ward identity in cosmological clothes. Register: **Derived**, given C-2; the kinematics of expansion, redshift, and horizons then follow as in Part III and Appendix F.

79.4 C-4 — the components as geometry

The source ledger replaces the fitted menu (Step 3 of Appendix E, §E.15): radiation and neutrinos inherited with the Standard-Model sector ($N_{\rm eff} = 3.044$, $\Delta N_{\rm eff} \leq 0.03$, a derived null); baryons from the inherited leptogenesis ledger ($\eta_B \approx 5 \times 10^{-10}$ from $K = 35$, Part V); cold dark matter the Kaxion of the compact angle (8.1 μeV, misalignment relic, Part V and Arisaka, 2026, C1-Kaxion); dark energy the seesaw terminus of the descent ($m_\Lambda = M^2_{\rm lock}/M_{\rm Pl,red} \approx 6.1$ meV) with $w = -1$ exact. Registers: **D*** on named lemmas, as graded row by row in Part IX. The third silent axiom is discharged: the cast list is geometry's, not the data's.

79.5 C-5 — the fractions

How much of each is the closed-packet topology of Part X: $\Omega_m = 1/\pi = 0.3183$, $\Omega_b = 1/(2\pi^2) = 0.0507$, $\Omega_\Lambda = 1 - 1/\pi = 0.6817$, the dark-matter slice the difference — raw percent-level agreement across the board, audited entry by entry in Part IX and at course pace in §E.16. Register, stated exactly: **D*** on Conjecture VIII (closed-packet integration), the load-bearing conditional of the cosmological codebook, which closes as the epoch identity — $\Omega_m = 1/\pi$ exact at $t_0 - 75$ Myr under Planck central values, residual lemma VIII-a exact to 10^{-15} — one declared calibration input, used once, printed in the open. The fourth silent axiom is discharged at that stated grade.

79.6 C-6 — the initial conditions

The launch of Part IV pays the last debt: flatness driven by $N^* = 52.6$ derived e-folds, the horizon problem dissolved by the shrinking comoving Hubble radius, and the spectrum exported with $n_s = 0.9634$ (**D***, plateau class), $r = 0.0038$ (**D*** — the decisive decade-scale prediction), and A_s the launch sector's one declared input, flagged wherever it appears. The walk is Appendix G; the register summary is §E.17.

79.7 C-7 — the expansion history

With the machine loaded, the history follows with no further input: the Hubble constant is derived, not fitted — $H_0 = 66.6 \pm 0.2$ from the π-locked inventory (two-road reading; Route I 66.8, §29.8), read against the CMB value 67.4 as the coherent π-offset of the sound-horizon dictionary (§40.4; Arisaka, 2026, C2-Hubble) — and the chronology of Part VIII carries every epoch on the radial clock. One entry demands its canonical two-quantity care: the corridor switch $z_{\rm sw} = e^7$ is the topological integer-ledger marker, a distinct quantity from the atomic decoupling $z_{\rm dec} \approx 1090$ computed from the same closed-packet fractions and landing on the Planck value at the 0.1% level (§42.1); the assembly inherits both, each with its own name. Register: per the codebook rows of Part IX.

79.8 C-8 — what ΛCDM does not contain

The export's absences are predictions with the same standing as its entries: no evolving equation of state ($w = -1$ exact — any confirmed $w \neq -1$ kills the terminus reading); no relativistic species beyond the Standard Model's ($\Delta N_{\rm eff} \leq 0.03$); no spatial curvature; no fifth force in the growth sector beyond Part VI's stated bands (Part VI); and no free $\{\Omega_i, w_i\}$ menu at all. Standard cosmology can absorb a surprise component by adding a row to the fit; the export cannot — which is precisely what makes it falsifiable, and the kill-test catalog of Appendix T holds the list.

80 The assembled model

The export chain is complete, and the concordance model can now be written the way no fit can write it:

$$\frac{H^2(a)}{H_0^2} = \underbrace{\Omega_r\, a^{-4}}_{\substack{\text{C-4: inherited thermal sector}\\ N_{\rm eff}=3.044}} + \underbrace{\frac{1}{\pi}\, a^{-3}}_{\substack{\text{C-4/C-5: Kaxion + baryons}\\ \Omega_b = 1/(2\pi^2)\ \text{inside}}} + \underbrace{\left(1-\frac{1}{\pi}\right)}_{\substack{\text{C-4/C-5: seesaw terminus}\\ w=-1\ \text{exact}}}, \qquad \underbrace{k=0}_{\text{C-6}} \tag{424}$$

riding on the derived clock and the exported spectrum,

$$\underbrace{H_0 = 66.8}_{\substack{\text{C-7: Route I; 66.6 two-road}\\ \text{CMB reading 67.4}}}, \qquad \underbrace{n_s = 0.9634, \quad r = 0.0038, \quad A_s \text{ declared}}_{\text{C-6: the launch spectrum}}, \qquad \underbrace{\text{nothing else}}_{\text{C-8: absences derived}} \tag{425}$$

with τ carried at its measured value as late-time astrophysics outside the fundamental claim. Equation (424) is, symbol for symbol, the expansion law of the textbooks — with one property no textbook version has ever had: *there is no fitted continuous parameter in the display.* One input is declared (A_s), one identity is epoch-calibrated and printed (Conjecture VIII), and everything else is geometry plus the integer ledger. The six-number fit of §77 compresses to that single declared amplitude plus the geometry — and the compression is checkable arithmetic: run (424) through the Step-2 machinery of Appendix E and the age of the universe comes out at 13.88 Gyr at the derived 66.8 (13.92 at the two-road 66.6), against the measured 13.79 ± 0.02 Gyr; the CMB's 67.4 gives 13.75 (Problem E.18.1); the full twenty-two-observable scoreboard, with deviations and pulls, is Table 41 and Figure 31. The honest one-sentence summary of the display: *the model that has served for fifty years as cosmology's input is here the output of a geometric descent with zero adjustable continuous inputs — with the strength of each step stated, not inflated.*

81 What the assembly does and does not claim

The shock value of Eq. (424) is real, and precisely for that reason this Part must close by marking its own boundary in the register the paper uses everywhere else.

Remark 1 (The honesty boundary of the assembled ΛCDM)**.** Six statements the assembly does *not* make. (i) General relativity's framework is imported at theorem grade from GG-A5-GDOS, consumed rather than rederived; the course-pace walk is Appendix E, and nothing here strengthens those theorems. (ii) Conjecture VIII is the load-bearing conditional: the fractions of C-5 stand on it, the epoch identity is a declared calibration input used once, and the assembly upgrades neither. (iii) The amplitude A_s remains the launch sector's declared input — the one IOU of the export, printed wherever it appears. (iv) The optical depth τ is late-time astrophysics outside the fundamental claim. (v) Every register marked **D*** stands on its named lemma or conjecture, graded row by row in Part IX with the promotion roadmap in Appendix S; nothing in the assembly changes a grade. (vi) The raw and calibrated readings are distinct claims and stay labeled; the two-quantity decoupling discipline of C-7 is part of the register, not a footnote. The assembly is exactly as strong

as the descent behind it, no stronger; what is new in this Part is not any single claim but the fact that the claims compose into the whole model.

For half a century ΛCDM has been the model cosmology reasons *from*. This Part has treated it as a model cosmology can reason *to* — the same inversion the companion GG-A9-SM performs for the Standard-Model Lagrangian, executed here on the largest possible stage. The reader who compares Eq. (424) with the opening chapter of any cosmology text will find the same symbols in the same places; the difference is that here every symbol arrives carrying its provenance, and the two great silences of the concordance model — *why these components* and *why these numbers* — have both been given answers with stated strengths.

The model is assembled. The next Part gathers the whole case onto a single sheet — the cross-paper coherence map, the comprehensive first-order codebook, and the second-order residual audit — so that a skeptical reader can verify the chain theorem by theorem.

82 Part XI in plain English: the standard model of cosmology, reassembled from the geometry

The purpose of this Part is to hand the reader back the model they already know. Everything in the concordance description — the six parameters, the expansion history, the perturbation spectrum — is reassembled here from the geometry, in eight stages, so the two descriptions can be compared line by line instead of in the abstract.

The section on the silent axioms is the one that repays slow reading. A textbook presentation of ΛCDM makes a series of structural assumptions — about homogeneity, about the nature of the dark sector, about what a parameter is — that are not usually written down because they are not usually questioned. Setting them out and naming the replacement for each is how this Part earns the word "assembled": the claim is not that the standard model is wrong but that it rests on choices, and that a different set of choices reproduces it.

What the assembly does not claim should be equally clear, and the closing section states it. Reproducing a model is not the same as predicting it. The assembly shows the geometry is sufficient to give back ΛCDM; it does not by itself show that no other structure would do the same. That question is what the cost-minimum principle of the codebook is for, and it is answered there partially and not completely.

Part XII
Summary and cross-paper coherence

The cross-paper coherence map to the eight upstream GG-Theory papers; the comprehensive first-order codebook and the second-order residual audit; why the result is not a post-hoc fit; and what makes GG-Theory structurally different from previous cosmology programs

83 Overview of the Summary Part

This Part is the cosmological analog of Part XII of GG-A9-SM ("Summary: derivation of the Standard Model with twenty-eight parameters"). It collects the cross-paper coherence map showing how the eight upstream papers of the Ultra Eight cohort (GG-A1-Constitution, GG-A2-GG6, GG-A3-BI6, GG-A4-Euclid, GG-A5-GDOS, GG-A6-GRIP, GG-A7-Qi4, GG-A9-SM) feed the cosmological derivation; the comprehensive first-order codebook compressing the twenty-one cosmological observables into a single audit sheet; the post-hoc-fit defense (why the π-locked inventory is not a numerological coincidence); and the GG-distinct comparison (what makes this cosmology program structurally different from ΛCDM with free parameters, EDE, MOND, brane-cosmology, and conformal-cyclic alternatives).

This Part presents the cross-paper coherence map (Section 84.1), the compressed codebook summary (Section 84.3), the compressed MSMF/OGN audit roadmap (Section 84.4), the post-hoc-fit defense, and the GG-distinct comparison. Every cross-paper citation resolves to a paper of the cohort.

GRIP: the common engine behind the twenty-one observables. Beneath the twenty-one codebook entries lies a single runtime. Every observable in the inventory — the π-locked density fractions, the one Hubble background, the clustering amplitudes, the damping knee, the Kaxion mass, the tensor ratio $r = 0.0038$ — is the present-epoch projection of one *Geometric Attractor*, selected on its own sub-stratum by the GRIP cascade DGI $\to$ GSSB $\to$ GA of GG-A6-GRIP, the runtime the whole cohort shares and that A10-Cosmos applies to the cosmological-boundary stratum. The MSMF/OGN cost-minimum the codebook reports for each entry *is* that GRIP-selected attractor: the unique saddle of the GFF ledger $\mathcal{F} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ on the relevant sub-stratum (Section 68.2.1). Read this way, the three-mode triage is a statement about GRIP itself: an entry is **Derived** where the GRIP uniqueness of its attractor is proved, **Derived (cond.)** where that uniqueness is conditional on a named lemma, and **Asserted** where the attractor is observed but its GRIP selection is not yet proved. The cosmic history of Part VIII is the same cascade read as a finite descent in the radial coordinate r: the GRIP-Out switch events — the inflationary graceful exit, the prompt lock crossing that stocks the dark sector, and recombination at $z_* = e^7$ — fix the initial data, while the long GRIP-In relaxations

carry the universe down to the present π-locked attractor $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\, 1 - 1/\pi,\, 1/(2\pi^2))$. The summary that follows is, at root, a census of one engine's outputs.

84 Cross-paper coherence map: the eight upstream papers of A10-Cosmos

84.1 What each upstream paper contributes to A10-Cosmos

Each entry below identifies the upstream paper by its canonical name (including the `Arisaka-GG-` prefix) using the *full* paper-code macro (e.g. Arisaka-GG-A1-Constitution), names its date, and states in one paragraph what A10-Cosmos inherits.

Arisaka-GG-A1-Constitution. The constitutional foundation of the GG-Theory program. A10-Cosmos inherits the four axioms (AX-1 LC locality and causality, AX-2 GGEP, AX-3 OGGEP operational/open GGEP, AX-4 MSMF maximum symmetry and minimum freedom) and the seven postulates (POS-1 GI gauge invariance, POS-2 UEP unified extremal principle, POS-3 AF anomaly freedom, POS-4 CP closure and positivity, POS-5 RG renormalization group, POS-6 OGN Occam gauge normalization, POS-7 MF measurement and falsifiability) that define admissibility, observability, and prediction. The codebook observables of Section 68.1 are admissible per AX-1; the BI-6 anomaly-free closure anchoring the Hubble two-stage corridor theorem (Theorem 68.8) is admissible per AX-2. The MSMF/OGN cost-minimum principle of Section 68.2.1 is the operationalization, on the cosmological-boundary stratum, of the constitutional clauses AX-4 (MSMF) and POS-6 (OGN); the cost functional itself is the GFF ledger $\mathcal{F} = \mathcal{F}_{\mathsf{GEP}} + \mathcal{F}_{\mathsf{CEP}}$ of GG-A6-GRIP inherited below.

Arisaka-GG-A2-GG6. The derivation of the six-dimensional GG-6 bulk (GG-A2-GG6 Theorem 6, T-GG6-6) and the closed-packet topology that the cosmological-boundary stratum projects onto. A10-Cosmos inherits the bulk geometry, the compact S^1_θ and the Green–Schwarz tensor its anomaly ledger forces (the structural seed from which the Kaxion is placed in the companion paper Arisaka-GG-C1-Kaxion; the cosmological descent itself is Conjecture 68.2 of the present audit), and the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ of GG-A2-GG6 Theorem 6 whose integer $J = 7$ is reused in the present paper to fix the Hubble corridor switch redshift $z_* = e^7$ (Theorem 68.8, property (C1)).

Arisaka-GG-A3-BI6. The modified Bianchi identity BI-6 (stated as the BI-6 local anomaly-cancellation theorem and the BI-6 complete global anomaly-freedom theorem of GG-A3-BI6) and the proof of its anomaly freedom. A10-Cosmos inherits BI-6 as the parent law of the cosmological boundary, with the explicit coefficient vector $(1, -1, -2/5, -2/5)$; the framing remark of Section 68 attaches the explicit citation chain. BI-6 anomaly-free recombination geometry is hypothesis (H2) of the Hubble two-stage corridor theorem. The four-form whose integral over the cosmological-boundary slice yields the matter fraction $\Omega_m = 1/\pi$ (Section 74.1) is constructed using the BI-6 anomaly-free closure machinery of GG-A3-BI6 (the BI-6 four-form construction of GG-A3-BI6); the explicit cosmological-boundary integration is reserved for future work.

Arisaka-GG-A4-Euclid. The zero-free-parameters discipline and the proof of the OGN-5 integer ledger $(1, 3, 5; 7, 35)$ (GG-A4-Euclid "The OGN-5 integers as a theorem of BI-6"). A10-Cosmos inherits the zero-free-parameters compression statement of GG-A4-Euclid, restated in cosmology form as "twenty-one cosmological observables forced by the Geometric Universal Five Constants GU-5 $= \{c, G_N, \hbar; k_B, Q_B\}$, the structural constants (π, e), and the integer ledger." The integers $J = 7$ (the Wilson-line winding / dyadic-closure integer of GG-A4-Euclid, cosmologically reused in the present paper as the recombination-corridor switch redshift $z_* = e^7$) and $K = 35$ (the Planck-EW scale-hierarchy integer of GG-A4-Euclid generating the e^{-35} hierarchy, cosmologically reused in the Kaxion's cap-loading ladder via A10's Conjecture 68.4) appear directly in the codebook closure economy.

Arisaka-GG-A5-GDOS. The Geometric Dynamics of Open Systems (GDOS) and the boundary projection structure OGGEP. A10-Cosmos inherits the open-system semantics: cosmological observables — the CMB sky, supernova Hubble diagrams, BAO scales, weak-lensing maps, gravitational-wave sirens — are all boundary exports of an underlying bulk, never the bulk directly. The Hubble two-stage corridor theorem proves a structural fact about *projected* observables, not about the underlying bulk; the projection is exactly the OGGEP projection of GG-A5-GDOS.

Arisaka-GG-A6-GRIP. The four Gi pillars (TGG, CCI, GFF, GRIP) and the runtime constitution of the program. A10-Cosmos inherits the GRIP cascade DGI $\to$ GSSB $\to$ GA as hypothesis (H3) of the Hubble two-stage corridor theorem: recombination is a GSSB transition from a tightly-coupled DGI phase to a frozen GA phase. The cost functional $\mathcal{F} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ of the GFF ledger is the MSMF/OGN cost functional in Section 68.2.1.

Arisaka-GG-A7-Qi4. The four Qi pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL), the selector architecture, and the operational language for the GG-A7-Qi4 five-story pagoda (particle $\to$ cosmos $\to$ chemistry $\to$ life $\to$ mind). A10-Cosmos inherits the pagoda-hierarchy language used in the universal-cosmic-identity preview of Section 71.4: the propagation of the uniqueness principle from electrons through atoms and molecules to cells instantiates the chemistry-and-life floors of the GG-A7-Qi4 pagoda, via the closed Qi-runtime relation OH + IH = EIC.

Arisaka-GG-A9-SM. The closed derivation of the twenty-eight-parameter Standard Model and the Codebook Provenance Table template. A10-Cosmos inherits the codebook table format directly (the structural triage D / D* / A; the per-parameter audit organization; the named-lemma promotion roadmap). The π-locked matter fraction $\Omega_m = 1/\pi$ extends the closed-packet topological argument of GG-A9-SM's identity section ("The geometric origin of identity: why every electron is the same") to the cosmological-boundary stratum. The universal-cosmic-identity preview of Section 71 is the cosmological analog of that section.

84.2 Cross-paper coherence map: outward and inward

This subsection compresses the eight inheritance paragraphs of Section 84.1 into a structured outward-and-inward map and records the inward-map links from A10-Cosmos to the rest of the GG-Theory cohort. The whole twenty-two-paper cohort and the cosmological-side closure it jointly establishes are shown in one view as Figure 35 below.

Outward map: what A10-Cosmos draws from the eight upstream papers.

1. GG-A1-Constitution → A10: AX-1 (LC), AX-2 (GGEP), AX-3 (OGGEP), AX-4 (MSMF); seven postulates (GI, UEP, AF, CP, RG, OGN, MF); the admissibility predicate that filters every codebook entry.

2. GG-A2-GG6 → A10: six-dimensional GG-6 bulk; OGN-5 integer tuple $(1, 3, 5; 7, 35)$ as Theorem 1; compact S^1_θ and the anomaly ledger's tensor seeding the Kaxion descent (closure deferred to GG-C1-Kaxion).

3. GG-A3-BI6 → A10: BI-6 as a stated theorem (local anomaly-cancellation theorem + global anomaly-freedom theorem); coefficient vector $(1, -1, -2/5, -2/5)$; the BI-6 four-form construction whose integral over the cosmological-boundary slice underlies the $\Omega_m = 1/\pi$ closed-packet integration (full integration reserved for future work).

4. GG-A4-Euclid → A10: zero-free-parameters discipline; OGN-5 integer ledger proved as a theorem of BI-6; cosmological reuse of $J = 7$ (Wilson-line winding / dyadic-closure integer in A4) and $K = 35$ (Planck-EW scale-hierarchy integer in A4) in the recombination switch and the Kaxion's cap-loading ladder.

5. GG-A5-GDOS → A10: GDOS as the open-system formalism; OGGEP projection structure; the rule that every cosmological observable is a boundary export of an underlying bulk rather than a direct bulk measurement.

6. GG-A6-GRIP → A10: four Gi pillars (TGG, CCI, GFF, GRIP); the GRIP runtime template DGI → GSSB → GA, and its cascade-closure theorem on the family-projector moduli, T4.1, which forces $N_g \leq 3$ and is conditional on GG-A6-GRIP's open problem M10; the GFF ledger $\mathcal{F} = \mathcal{F}_{\mathsf{GEP}} + \mathcal{F}_{\mathsf{CEP}}$ used as the MSMF/OGN cost functional in Section 68.2.1.

7. GG-A7-Qi4 → A10: four Qi pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL); five-story pagoda (particle → cosmos → chemistry → life → mind); the closed Qi-runtime relation OH+IH = EIC for the layer-by-layer universal-identity propagation of Section 71.4.

8. GG-A9-SM → A10: 28-parameter Codebook Provenance Table template (the D / D* / A classification triage and per-parameter audit organization); the section "The geometric origin of identity: why every electron is the same" that A10 extends one floor up to the cosmological-boundary stratum in Section 71.

Inward map: what A10-Cosmos delivers downstream. A10-Cosmos is the boundary-side closure of the Ultra Eight cohort and the structural foundation for the cosmological-prediction branch. The four downstream papers that draw from A10-Cosmos are:

1. A10 → GG-A8-QM: the operational reformulation of quantum mechanics in 4D is carried in the completed A8-QM paper, cited throughout; A8 inherits the OGGEP projection from A5 and the cosmological-boundary stratum from A10 as its boundary backdrop.

2. A10 → Arisaka-GG-P1-KK: the KK lock-cluster phenomenology at $M_{\text{lock}} \approx 3.85$ TeV uses the same OGN-5 integer ledger whose $K = 35$ anchors the cosmological Kaxion's cap-loading ladder; A10's cosmological observables $(\Omega_m, \Omega_\Lambda, \Omega_b)$ constrain the KK relic abundance.

3. A10 → Arisaka-GG-P2-N3 (Arisaka, 2026, P2-N3): the three-leg derivation of the generation count $N_g = 3$ — its upper bound now the capacity theorem Cap $= \pi N_g$ with the integrality theorem L-INT behind it, conditional on one declared topological residue, R-INT — cross-cites A10 for the cosmological role of $N_g = 3$ (the three-generation count in the matter sector inherited from cosmological closure). The generation count is not the OGN-5 color index; the two are distinct quantities that share a value.

4. A10 → Arisaka-GG-C1-Kaxion: the dedicated Kaxion paper takes A10's Conjecture 68.4 (compact-angle-mass-uniqueness) as the cosmological constraint on Kaxion mass and decay constant, calibrating the one scale it leaves open and declaring the calibration; the cosmological role of the Kaxion as cold dark matter saturating $\Omega_{\text{DM}} = \Omega_m - \Omega_b = 1/\pi - 1/(2\pi^2)$ is the empirical anchor of C1.

Net coherence summary. The outward map records eight imports of well-defined upstream theorems and one operational language (the five-story pagoda). The inward map records four downstream deliveries to predictions-branch papers. The two-way coherence is therefore: A10-Cosmos sits at the boundary-side closure of the constitutional foundation (A1) + bulk geometry (A2-A3-A4) + open-system measurement (A5-A6-A7), inheriting one theorem from each of those layers; and A10-Cosmos forwards four prediction branches into the future program work. No upstream theorem is inherited circularly; no downstream branch is forwarded into a paper that was already cited as upstream.

The corresponding cohort map (Figure 35) shows the whole twenty-two-paper program in one view: the constitutional root and the GG6 Trio, the open-system boundary law and its pillar papers, the 4D-Physics Trio with A10-Cosmos at the cosmological-boundary closure, the capstone, the two phenomenology branches, and the Foundations Branch column on the left flank.

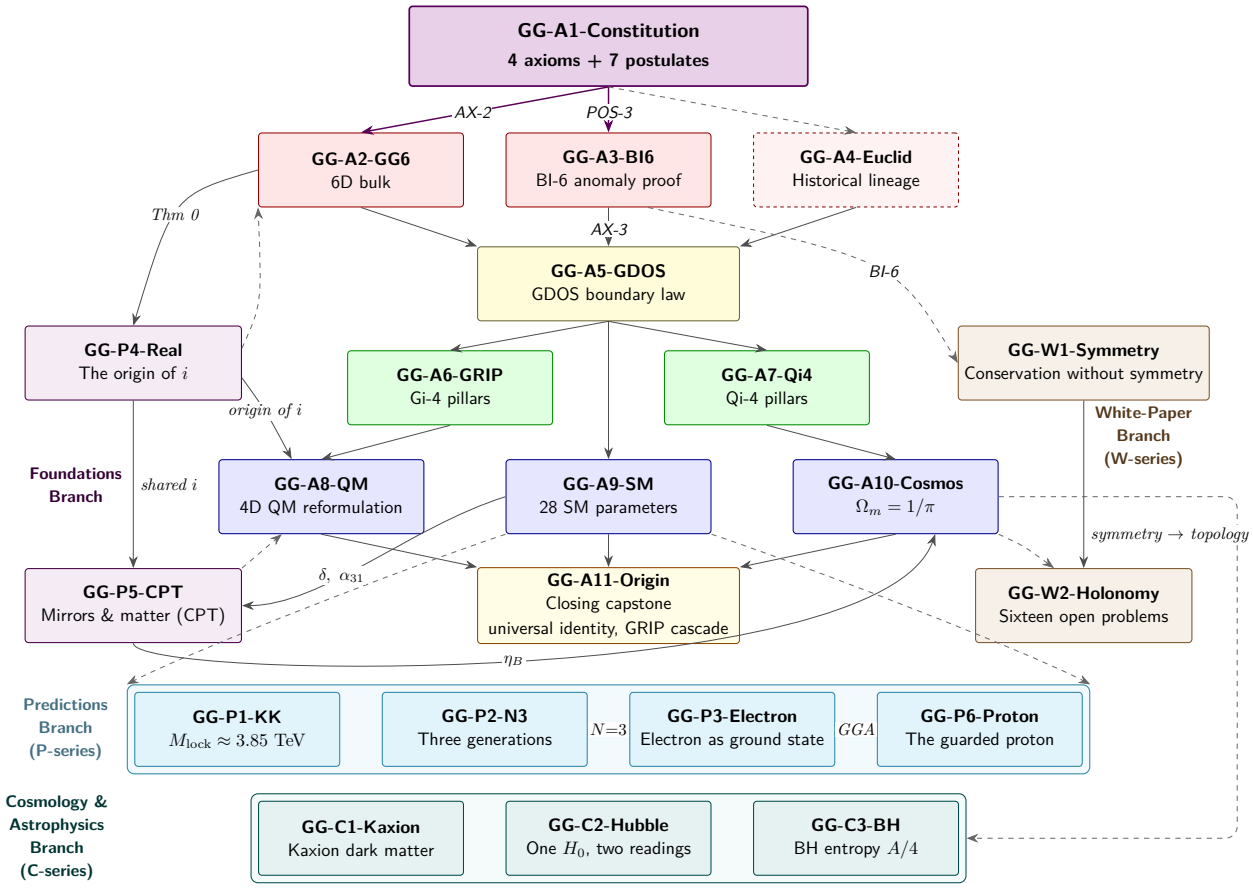

Figure 35: The GG-Theory cohort in one view (Summary recap). The twenty-two-paper cohort map of the front matter (Figure 11), repeated here as the visual summary of the cross-paper coherence. The constitutional root GG-A1-Constitution compiles via AX-2 (GGEP) into the GG6 Trio (GG-A2-GG6, GG-A3-BI6, GG-A4-Euclid); AX-3 (OGGEP) forces the open-system boundary law GG-A5-GDOS, organized by the pillar papers GG-A6-GRIP and GG-A7-Qi4; the 4D-Physics Trio (GG-A8-QM, GG-A9-SM, GG-A10-Cosmos) recovers quantum mechanics, the Standard Model, and ΛCDM, with the present paper the cosmological-boundary closure; GG-A11-Origin closes the A-series; the framed Predictions (P-series) and Cosmology & Astrophysics (C-series) branches carry the phenomenology; and the Foundations Branch column on the left flank — GG-P4-Real (the origin of i) and GG-P5-CPT (the mirrors and the matter) — carries the boundary's complex structure and its mirrors, GG-P5-CPT's η_B edge being this paper's import. A10-Cosmos inherits one theorem from each upstream layer and forwards its cosmological closure into the downstream branches, as enumerated in Section 84.2.

After all, the cross-paper coherence map is the map of the program's spine: eight papers feed in, one paper closes the cosmological side, and four papers feed forward. No paper is isolated; no claim is unanchored.

Program-level coherence: the eight-paper cohort. The eight papers GG-A1-Constitution, GG-A2-GG6, GG-A3-BI6, GG-A4-Euclid, GG-A5-GDOS, GG-A6-GRIP, GG-A7-Qi4, GG-A9-SM together comprise the Ultra Eight cohort, all sharing the same sans-serif paper-code convention. A10-Cosmos is the boundary-side closure of that cohort: every constitutional inheritance is sourced, every codebook entry is anchored, and every named lemma is identified with the upstream theorem chain that would close it.

84.3 Compressed codebook summary

Table 45 compresses the twenty-one-row codebook of Section 68.1 into a single audit sheet suitable for the headline summary. Each row lists the observable, its symbol, its status under the MSMF/OGN principle, and the named lemma (if any) that closes it. Equivalently, each row names a Geometric Attractor and the maturity of its GRIP-selection proof: where a named lemma is listed, that lemma *is* the GRIP-uniqueness statement for the observable's sub-stratum.

Table 45: Compressed codebook summary: twenty-one cosmological observables, three-mode status (D / D* / A), and the named conjecture (if any) that closes each. Compare with the full provenance table in Section 68.1. The fourth column names the conjecture that would close each row that is not yet Derived, so the table doubles as a list of what remains open.

Observable	Symbol	Status	Closing conjecture (if D* or A)
Matter density fraction	$\Omega_m = 1/\pi$	D*	Closed-packet-integration (Conj. VIII)
Cosmological constant fraction	$\Omega_\Lambda = 1 - 1/\pi$	D*	Conjecture VIII
Dark-energy equation of state	$w = -1$	D*	Vacuum-residual-exactness (Conj. IX)
Baryon density fraction	$\Omega_b = 1/(2\pi^2)$	D*	Conjecture VIII
Dark-matter density fraction	$\Omega_{\rm DM}$	D*	Kaxion-misalignment-completeness (Conj. I), with C-Kaxion-d
Baryon-to-photon ratio	η_B	D*	Resonant-leptogenesis-normalization (Conj. VI)
Lithium-7 abundance	$^7{\rm Li}/H$	D*	Lithium-port-π-suppression (Conj. X)
Hubble constant, against the microwave background	$H_0^{\rm CMB}$	D*	Stratum-selection-of-the-vacuum-residue (Lemma II)
Hubble constant, against the distance ladder	$H_0^{\rm late}$	D*	Lemma II
Decoupling redshift	z_* (switch e^7)	D*	Conjecture VIII (atomic chain on the fractions; §42.1)
Matter clustering amplitude	σ_8	D*	Lemma V′ (computed from the declared amplitude; §43.3)
S_8 clustering, against the CMB	S_8	D*	Lemma V′
S_8 clustering, against cosmic shear	S_8	D*	Lemma V′ (one number, second comparator)
Effective neutrino deviation	$\Delta N_{\rm eff}$	D	— (derived null: $\Delta N_{\rm eff} \leq 0.03$, $N_{\rm eff} = 3.044$)
Damping-tail multipole	ℓ_D	D*	Closed-packet-integration (Conj. VIII); BFFT overlay gated by Conj. VII
Tensor-to-scalar ratio	r	D*	Lemma IV (revised) ($r = 12/N^{*2} = 0.0043$ at leading order; 0.0038 exact)
Scalar tilt	n_s	D*	Plateau-class-and-prompt-crossing (Conj. IV, revised)
Scalar amplitude	A_s	A	GDOS-heat-kernel amplitude (Conj. V′)
Kaxion mass	$m_{\rm Kaxion}$	D*	Compact-angle-mass-uniqueness (Conj. III)
Kaxion decay constant	$f_{\rm Kaxion}$	D*	Conjecture III
Inflationary energy scale	$V_{\rm inf}^{1/4}$	D*	Plateau-class-and-prompt-crossing (Conj. IV, revised)

84.4 Compressed MSMF/OGN audit roadmap

Table 46 compresses the per-parameter audit of Section 68.2 into a ten-step closure roadmap: which named conjecture closes which codebook rows, and what the decomposition becomes once each conjecture is proved.

Table 46: Compressed audit roadmap: the ten named conjectures of A10-Cosmos, the codebook rows each closes, and the resulting decomposition. The codebook collapses from 1 D + 19 D* + 1 A to 21 D + 0 D* + 0 A once all ten are proved, with the companion's pinning conjecture besides for D4; twenty-one entries, about eight of them independent ("Reading the table", Part I).

No.	Named conjecture	Rows closed	Count	State after proof
I	Kaxion-misalignment-completeness (with C-Kaxion-d, GG-C1-Kaxion)	$\Omega_{\rm DM}$	1	2D / 18D* / 1A
II	Stratum-assignment of the vacuum residue	H_0 (against the CMB and the ladder)	2	4D / 16D* / 1A
III	Compact-angle-mass-uniqueness	$m_{\rm Kaxion}$, $f_{\rm Kaxion}$	2	6D / 14D* / 1A
IV	Plateau-class-and-prompt-crossing (revised)	n_s, r ($V_{\rm inf}^{1/4}$ also on V′)	2	8D / 12D* / 1A
V′	GDOS-heat-kernel-amplitude (replaces the retired Conj. V)	A_s; $V_{\rm inf}^{1/4}$; σ_8, S_8 (two comparators)	5	13D / 8D* / 0A
VI	Resonant-leptogenesis-normalization	η_B	1	14D / 7D* / 0A
VII	BFFT-damping-knee-identification	(BFFT overlay of ℓ_D; row closed by VIII)	0	14D / 7D* / 0A
VIII	Closed-packet-integration	Ω_m, Ω_Λ, Ω_b, ℓ_D, z_*	5	19D / 2D* / 0A
IX	Vacuum-residual-exactness	w	1	20D / 1D* / 0A
X	Lithium-port-π-suppression	$^7{\rm Li}/H$	1	**21D / 0D* / 0A**

The ten conjectures between them close twenty of the twenty-one rows (the twenty-first, $\Delta N_{\rm eff} \leq 0.03$, is derived): nine close rows directly, while Conj. VII gates the interpretive BFFT overlay of D18 (whose row is closed by Conj. VIII). Five of the ten (Conj. II, III, IV, V′ and VIII) each close multiple rows. The closure is strictly monotonic in the codebook decomposition.

84.5 The second-order residual: the +2.7% common offset, honestly tiered

The compressed codebook above is the *first-order* ledger: each entry is the leading geometric form read off the OGN-5 integers and the closed-packet topology, with no continuous dial. A first-order ledger invites a second-order question — how do those leading forms sit against the data once the precision is taken seriously? — and the honest answer is a residual that must be tiered, not buried.

The observation. Two of the predictions of the parameter-audit Dev plot (Table 41) land at almost exactly the same positive deviation: w at +2.72% and Ω_b at +2.76% — a spread of 0.04 percentage points, the same sign, each first-order geometric form sitting about 2.7% *above* the measured value. (An earlier version counted a third member, $S_8^{\rm late}$ at +2.77%; that row was a corridor response computed at a radiation excess the paper has since retired, and it is withdrawn — S_8 is now one computed number at −0.8% against the microwave background. Two members are thinner evidence for a common offset than three, and the paragraph is weakened accordingly rather than repaired by substitution: the nearest surviving relative, Ω_b/Ω_m at +1.74%, is Ω_b's own ratio statement and not an independent third.) The cluster used to have a fourth member: the damping multipole, whose integer-ledger form $K(2\pi)^2$ sat at +2.73% until the diffusion integral was actually computed on the inventory and the entry was re-scored to its Boltzmann value (+0.4%) — a departure whose lesson is drawn below. The remaining independent entries scatter on both sides (Ω_m +1.1%, Ω_Λ −0.5%, $\Omega_{\rm DM}$ +0.7%, $H_0^{\rm CMB}$ −0.9%, $z_{\rm dec}$ +0.15%, ω_b/ω_m +1.7%, $100\,\theta_*$ −0.3%, the peak ratios −0.8% and −0.9%, ℓ_D +0.4%, σ_8 −1.4%, η_B −20% (o.o.m. band), $^7{\rm Li}/H$ −0.6%; the $H_0^{\rm late}$ entry, −9.1%, is the separately named ladder systematic), so the coincidence is specific to those two, not a rigid shift of the whole plot. *Epoch-calibration re-scoring.* The deviations above are the raw t_π-form comparisons against today's sky. Under the single declared epoch input (§32.11) the baryon entry re-scores to +1.69% (+1.4σ) — its timing component (+1.05%, shared with the whole matter sector) removed — while w, which does not sweep with the epoch, keeps its +2.7%. The pair is therefore partly an epoch artifact: at the calibrated level it is one +2.7% (w) beside a +1.7% (Ω_b, calibrated) and a +1.7% ratio statement.

What it is, and is not, statistically. Only one of the two carries real statistical weight, and the epoch calibration nearly halves it: Ω_b reads +2.76% (2.3σ) against the raw t_π-form and +1.69% (+1.4σ) once the declared epoch input is used — the surviving residual being exactly the epoch-independent, h-free ratio statement $\Omega_b/\Omega_m = 1/2\pi$. The other is weak company — w at +2.72% but only 0.9σ, well inside the dark-energy error bar — so the two landing at the same offset is as plausibly accidental as structural, and more plausibly so than when a third row stood beside them. The honest signal is therefore "a single $\sim$1.4σ baryon-ratio pull (2.3σ before the epoch calibration), with two suggestive but inconclusive companions," not "a universal +2.7% offset across cosmology."

The physics: this is exactly the Tier-3 residual — now with a measured precedent. The first-order forms are leading topological and integer expressions; they do not carry the second-order precision content — baryon-loading corrections, the finite-width visibility integral, diffusion damping — that a full Einstein–Boltzmann computation supplies. The damping multipole is the test case for this reading: when the Tier-3 content was actually computed for that entry (the diffusion integral over the Boltzmann ionization history, Section 40), the +2.7% gap of the integer form collapsed to +0.4% — the offset really was the omitted second-order content, and computing it moved the entry onto the data. The epoch calibration is now a *second* such precedent: supplying the one declared input the raw comparison had omitted — *when* the snapshot is read — moved Ω_m, Ω_Λ, and $\Omega_{\rm DM}$ onto the data outright (input, 0.00%, −0.31%) and carried Ω_b from +2.76% to its ratio residual +1.69%. Read this way the residual cluster is not an embarrassment but a *measurement*

of where the first-order ledger ends, together with a demonstrated route for what happens when the instrument is brought to bear: the inventory is right to a few percent, and the sub-few-percent residue is the instrument's province, not the geometry's. One candidate source is now positively excluded: the stationary GDOS measure itself cannot generate the offset, being uniform on the holonomy moduli to $O(e^{-S_w}) \sim 10^{-15}$ (the GSM adjunct of GG-A5-GDOS), thirteen orders below the percent level.

Three readings, none over-claimed. *(i) Coincidence:* two percent-level numbers, one of them statistically weak, agree by chance — the null hypothesis, which two members make considerably harder to reject than three did, and not to be waved away. *(ii) A common second-order normalization:* the damping-multipole precedent shows exactly how such a correction operates — computing the omitted content moved that entry onto the data — so the falsifiable expectation is that a computed second-order baryon correction would do for Ω_b what the diffusion integral did for ℓ_D — and the epoch calibration has already performed the timing part of exactly this operation — checkable the moment that computation is done. *(iii) Selection:* the offset is simply where current data happen to place these forms, and future inference will scatter them. We commit to none of the three. In particular we resist the temptation to *name* the number: 2.7% sits near $1/36 = 2.78\%$, $e/100 = 2.72\%$, and $1/(4\pi^2) = 2.53\%$, and assigning it any such closed form on the present evidence would be the exact numerological move this paper exists to refuse (Section 85.1). The residual is reported, bounded, and left open.

Consequence for the triage, and the declared precision floor. Nothing here changes a codebook status. The two entries remain Derived-conditional on their named conjectures, and the +2.7% is carried as an explicitly *underived* second-order term, not folded into the first-order claim and never quoted as a success. We go one step further and adopt it as policy, because a critic is right to insist on it: **the framework's honest precision at first order is $\sim$2.7%, not the tenths-of-a-percent that the tightest single entries can appear to reach.** The coherent, same-sign clustering of independent predictions is the classic signature of one omitted upstream normalization, and where we have *computed* the omitted content the offset collapses — the damping knee's integer form moved from +2.7% to +0.4% once the diffusion integral was evaluated, and the epoch calibration moved the matter fractions onto the data. Accordingly the π-locked topological forms should be read against a declared second-order floor of $\sim$2.7%, and the per-entry σ's of the tightest such rows are *first-order* pulls, not the framework's true precision. The floor is a property of those forms and does not govern the slow-roll sector: the tilt and the tensor ratio carry their own budget — the half e-fold in N^* and the plateau-class conditionality — and the 3.2σ tilt tension of §28.10 is stated against that budget, not softened by this floor. This yields a specific, falsifiable expectation rather than a hedge: a computed second-order baryon correction should do for Ω_b what the diffusion integral did for ℓ_D, carrying its +2.76% (2.3σ) onto the data — and if, once computed honestly, it does not, the $\Omega_b = 1/(2\pi^2)$ entry fails. The floor is thus not an excuse for the residual but a prediction about it: leading order is forced to $\sim$1%, the second-order content is computable, and the instrument — not the geometry — owns the sub-few-percent sky, exactly as the tier ladder reserves.

85 Post-hoc-fit defense: why the π-locked inventory is not numerology

The reader who arrives at the codebook for the first time will rightly ask: is the π-locked inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$ a derivation, or is it a curve-fit dressed up in topological language? This section gives the four-point structural defense against *continuous* fitting (D1–D4 below), then, in Section 85.1, the fuller answer to the *discrete* look-elsewhere charge — the combinatorial-alphabet objection — and the Euclidean reading that places the whole construction in its proper category.

(D1) Closed-packet topology forces the value, not the data. Theorem 74.1 of Section 74.3 states the structural form of the closed-packet $\Omega_m = 1/\pi$ claim: the BI-6 four-form on $\mathcal{M}_4$ admits no continuous adjustable parameter, and the OGN-5 integer ledger $(1, 3, 5; 7, 35)$, established upstream (GG-A2-GG6/GG-A3-BI6, audited in GG-A4-Euclid), fixes the rational weight of the matter cell at $1/15$ before the Pontryagin $1/\pi$ rescaling. The integration is programmatic — it does not yet close (Conjecture VIII), so Ω_m remains D*. The structural point stands regardless: if the closed-packet integration can be completed, $1/\pi$ ceases to be a numerological observation and becomes the integral output of a fixed topological computation rather than a number chosen to fit the data.

(D2) One topology forces three coupled entries simultaneously. A post-hoc fit explains one number, not three coupled numbers. The same closed-packet integration that gives $\Omega_m = 1/\pi$ gives $\Omega_\Lambda = 1 - 1/\pi$ by unit-trace complementarity and $\Omega_b = 1/(2\pi^2)$ by a one-dimensional further projection through the $U(1)_B$ baryonic bundle. No additional curve-fit step is involved between the three values.

(D3) The $2\pi^2$ factor requires the right combination of powers. The baryon fraction depends on $2\pi^2$, not on π. A post-hoc fit could be tuned to match either π or π^2 but not to predict the specific combination $1/(2\pi^2)$ that emerges from the closed-packet integration through the $U(1)_B$ bundle. The fact that the topological computation produces the right *combination* of powers of π (one for matter, one minus that for vacuum, two for baryons) is the structural signature of a derivation.

(D4) Every other codebook entry is independently constrained. Beyond the three π-locked density fractions, the codebook contains eighteen more observables. The Hubble two-stage corridor switch $z_* = e^7$ descends from the OGN-5 integer $J = 7$ (not from π); the tensor-to-scalar ratio $r = 12/N^{*2} = 0.0043$ (0.0038 exact) descends from the Starobinsky relation at the reheat-matched count (not from π); the effective neutrino number deviation $\Delta N_{\rm eff}$ descends from the inherited Standard-Model ledger of GG-A9-SM (not from π). The codebook is not a single π-fitting exercise; it is eighteen further observables sharing a common geometric origin but emerging through distinct calculations — of which, netting the restatements, about eight are independent (Table 41).

The structural-fit hypothesis is falsifiable. The defense is not a rhetorical move. If any one of the eight falsification tests listed in Appendix T fails — if the empirical recombination redshift sits at $z_{\rm rec} = 1500$ instead of ≈ 1090, or if the GW-siren Hubble determination lands at the ladder's $H_0^{\rm GW} \approx 73$ km/s/Mpc rather than at the background, or if the tensor-to-scalar ratio is detected at $r > 0.01$ or excluded below 10^{-3} — a named conjecture fails (the switch identification and the one background under Conjecture II, the plateau class under Conjecture IV) and what rests on it falls with it. After all, a structurally locked derivation that cannot be falsified is not a derivation; it is a tautology. The present framework's commitment is that it can be falsified.

85.1 Forcing versus searching: the discrete look-elsewhere effect, and the Euclidean reading

The four points above answer the charge of *continuous* fitting: no real-valued coefficient is tuned to the Planck data. But there is a second form of the objection, discrete rather than continuous, that the continuous-fit rebuttal does not by itself dispatch, and it is the form a cosmologist is most likely to voice:

> *With the integers* $(1, 3, 5; 7, 35)$ *and the two transcendentals* $\{e, \pi\}$*, and a handful of operations, one can build an expression close to almost any* $\mathcal{O}(1)$ *cosmic fraction. So* $\Omega_m = 1/\pi$*,* $\Omega_b = 1/(2\pi^2)$*, and* $z_* = e^7$ *are a look-elsewhere artifact, not a derivation.*

This is a legitimate statistical concern — the cardinality of a discrete formula-space against the number and precision of the targets — and it deserves a direct, quantitative answer. The answer is that GG-Theory does not *search* the alphabet. It *forces* the forms. Five facts, each independently checkable, separate forcing from searching; a sixth places the whole construction in its proper historical category.

One ledger, not a per-observable toolbox. The objection tacitly grants itself a fresh draw from the alphabet for *each* observable. That is the error. A genuine combinatorial fit of the twenty-one codebook observables requires of order twenty independent formula-choices — one expression selected per observable, or about eight once the rows that are one geometry read in other units are counted once, the front matter's honest count — so its discrete freedom scales with the number of targets. GG-Theory introduces zero continuous parameters and *one* integer tuple $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, and that tuple is not five free integers but a *closed* object: the constitutional relations $J = 2N + 1$ and $K = LJ = M^2 + N^2 + L^2$ fix J and K from (M, N, L); $M = 1$ is forced; and $N = 3$ is itself a theorem (the cascade-budget closure of Arisaka, 2026, P2-N3). The effective discrete freedom of the entire ledger is therefore one or two integers, and the *same* tuple must serve *all twenty-one* observables at once. The comparison is thus of order eight to twenty free choices against roughly two: the cosmological codebook is enormously over-constrained, and over-constraint is the structural signature of a derivation, not a fit.

The sharpest expression of this is *cross-observable rigidity*. The single integer $J = 7$ fixes the recombination switch redshift $z_* = e^7$ and, through $K = LJ = 35$, the lock scale that the descent crosses on its way from the plateau to the terminus (the plateau's height is the declared amplitude,

not an integer; Part IV); the single integer $K = 35$ anchors the Kaxion's cap-loading ladder and the conjectured CMB damping knee $\ell_D = K(2\pi)^2$ (an overlay +2.3% above the Boltzmann row); and one closed-packet topology fixes $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, and $\Omega_b = 1/(2\pi^2)$ simultaneously. If these integers were fitted observable by observable, nothing would require the *same* $J = 7$ and $K = 35$ to succeed across recombination, the lock crossing, dark matter, and the CMB damping tail at once; a combinatorial search carries no mechanism that forces one integer to be reused coherently across unrelated observables. The ledger does exactly that, by construction. The coherent reuse *is* the non-numerology.

The form is forced by the stratum, not selected by scanning. The objection further imagines that, for each observable, one scans candidate forms — $1/\pi$? $1/(2\pi^2)$? e^7? — until one lands. In GG-Theory the *assignment of form to observable* is dictated by the geometric origin of that observable, before any number is evaluated. The matter fraction is $\Omega_m = 1/\pi$ because the matter cell is the closed-packet Pontryagin integral on the cosmological-boundary stratum, carrying one factor of the angular primitive π; the vacuum fraction is $\Omega_\Lambda = 1 - 1/\pi$ by unit-trace complementarity on the same packet; the baryon fraction is $\Omega_b = 1/(2\pi^2)$ because the baryonic $U(1)_B$ bundle is a one-dimensional further projection carrying the $2\pi^2 = \mathrm{Vol}(S^3)$ volume; the corridor-switch redshift is $z_* = e^7$ because the last-scattering switch surface is identified as the $J = 7$ GSSB event (a topological landmark distinct from the measured decoupling redshift, coinciding to 0.6%), the exponential of a single OGN-5 integer; and the tensor ratio $r = 12/N^{*2} = 0.0043$ is the Starobinsky relation evaluated at the reheat-matched count on the derived plateau, one e-fold below the GUT scale, not a fitted exponent. In each case the form is a theorem, or a named-conjecture consequence, of *which stratum the observable inhabits*; the integers enter only through their fixed values. The search over forms is replaced by a geometric assignment, and the search over values by one closed ledger. There is no free product over an alphabet at any step.

MSMF and OGN remove freedom; they do not supply it. Where an observable could admit more than one constitutionally-allowed form, the parsimony axiom MSMF with the OGN complexity measure selects the minimal-cost form on the ℓ^1 norm of the ledger (Section 68.2.1). This is a constraint that *deletes* options, applied as a single, pre-registered rule, uniformly across all observables — the opposite of a per-observable free choice. "One way or another" presumes the freedom to choose the combination; MSMF forbids choosing, and forces the minimal-complexity admissible form, with the resulting cost recorded entry by entry in the codebook audit of Part IX (Section 68.2). Here honesty must be explicit, and it is a strength rather than a concession: for the entries where the empirical value enters the selection — the calibrated Kaxion scale, the declared amplitude — those are data-anchored selections, not theorems, and the codebook marks them D* or A in the one/nineteen/one triage. A curve-fit hides its degrees of freedom; this audit publishes them.

The look-elsewhere effect, bounded and quantified: the ledger is pinned, not searched. The decisive answer to a discrete look-elsewhere charge is not assertion but a count, taken *globally* — the whole inventory against the whole declared alphabet at once — and we now give the count rather

than promise it. The alphabet is the OGN-5 tuple $(M, N, L; J, K)$ subject to its two constitutional closure relations $J = 2N + 1$ and $K = LJ = M^2 + N^2 + L^2$, plus $\{\pi, e\}$ and a bounded operation menu. The trials factor factorizes as $p_{\text{global}} \approx P(\text{ledger}) \times P(\text{forms} \mid \text{ledger})$, and the first factor is where the objection dies.

The ledger factor is pinned by the particle sector, not fitted to cosmology. Enumerating *every* closed tuple in the natural box $M, N, L \in \{1, \ldots, 9\}$ (both closure relations imposed) yields exactly ten: $(1,1,1;3,3)$, $(1,1,2;3,6)$, $(1,3,2;7,14)$, $(1,3,5;7,35)$, $(2,4,4;9,36)$, $(2,4,5;9,45)$, $(2,6,5;13,65)$, $(2,6,8;13,104)$, $(1,7,5;15,75)$, $(3,9,9;19,171)$. Impose the two integers that are fixed *upstream*, in the particle sector and not by any cosmological number — the color count $N = N_c = 3$ (the dual-lock closure of GG-A3-BI6, Step 3 of its OGN-5 theorem; the trace ratio is its independent check) and the gauge label $L = 5$ (the Green–Schwarz block size of GG-A3-BI6, Step 2; SU(5) is the reading GG-A9-SM gives it) — and exactly *one* survives: $(1, 3, 5; 7, 35)$. The cosmological ledger therefore carries *zero* discrete freedom of its own; it is inherited from the generation count and the gauge group, both known before cosmology is consulted. $P(\text{ledger}) = 1$ conditional on the known particle content, or $1/10$ if one refuses even that conditioning.

The exponential observables cannot be searched. The scale-hierarchy predictions are exponentials of the pinned integers, so they admit no nearby-integer alternative: only $J = 7$ places the switch redshift $z_* = e^J$ near recombination ($e^6 = 403$, $e^7 = 1097$, $e^8 = 2981$ against the observed ≈ 1090), and only $K = 35$ places the lock $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-K}$ at the TeV scale ($K = 34 : 10.5$, $K = 35 : 3.85$, $K = 36 : 1.4$ TeV) — while the *same* $K = 35$ delivers $v = M_{\text{lock}}(2/5)^3 = 246.2$ GeV unbidden. Adjacent integers miss by multiplicative jumps of e; there is no continuum to scan.

The residual form factor is small, and the computation an earlier version owed has been run. With the ledger pinned, the only freedom left is the form assignment, and it is heavily over-determined: the three inventory fractions come from *one* Hopf fibration (not three choices), the exponential observables from *one* map $x \mapsto e^x$ on the pinned integers, the tensor ratio from the fixed Starobinsky relation. The Monte-Carlo that an earlier version of this paragraph promised has now been run, and it is reported here because its result is unfavorable to the crude bound that version printed. On the *unrestricted* alphabet $\{2, N, L, J, \pi, e, K\}$ at cost ≤ 4, a random $\mathcal{O}(1)$ fraction is matched to within one percent with probability 0.98, and to within the declared 2.7% second-order floor with probability 1.00: an unrestricted bounded-cost search proves nothing about any single number, and the bound $p \sim 10^{-3}$–10^{-4} that version quoted for it is withdrawn. What survives the Monte-Carlo is what the preceding two paragraphs said: the ledger is pinned upstream, the exponential observables admit no nearby-integer alternative, and the three inventory fractions are one Hopf-fibration form, not three searched ones. The honest statement of the residual is therefore not a probability but a discipline: every form in this paper is assigned by a stratum rule stated *before* the number is compared (§29.8 for ζ_Λ is the worked example, where the stratum alphabet has one admitted member at cost ≤ 4 in the observed decade and the unrestricted alphabet has forty-nine), and a reader who doubts a form should ask which rule assigned it, not how many alphabets could have produced it.

Forward commitments no fit can make. A fit accommodates whatever the data say and cannot be falsified. GG-Theory issues forward commitments that can (Section 95, F-Cosmos-1 through F-Cosmos-5): a primordial tensor ratio off 0.0038 beyond the instrument's precision — outside $[10^{-3}, 10^{-2}]$ at the least — breaks the Starobinsky kill-switch; a ladder-free late probe returning the ladder's 73.5 rather than the background 66.8 falsifies the one-background reading (Corollary 68.11); a firm primordial ^{7}Li abundance away from SBBN/π excludes the lithium port (Conjecture 68.5); a measured decoupling redshift off the framework's atomic-chain value breaks the two-quantity reading (§42.1); and the 8.1 μeV Kaxion shared with Arisaka (2026, C1-Kaxion) is a laboratory target, its one scale calibrated and declared before the search. Each was fixed by the ledger before the measurement that will test it, and a central value far from it breaks the assignment. Falsifiable forward prediction is a property no recombination of a small alphabet possesses, and it is the cleanest empirical demarcation between a derivation and a fit.

The Euclidean reading: a constitution, not a coincidence. These five facts have a single root, and naming it answers the objection at its deepest level. The reduction of the cosmic inventory in this paper is not a collection of fortunate expressions; it is the descent of propositions from a fixed set of axioms — the GG-Constitution of Arisaka (2026, A1-Constitution): the equivalence principle GGEP, the open-projection axiom OGGEP, the parsimony axiom MSMF, the normalization principle OGN, and the anomaly law BI-6. From these the geometry, the ledger, and the π-locked inventory follow by construction, each graded by how complete its proof currently is. In this light the alphabet objection is the exact analog of complaining that Euclid "could prove anything with points, lines, and circles." The complaint mistakes the *primitives* of a deductive system for the *freedom* of a fit. In a deductive system the primitives are few and fixed and the propositions are *forced*; the richness of what follows is not evidence of fitting but the entire purpose of the axiomatic method. The twenty-one observables of the cosmological codebook are not a look-elsewhere artifact of $(1, 3, 5; 7, 35)$ and $\{e, \pi\}$; they are theorems of a constitution — some proved, some conditional on a named conjecture, some still asserted, exactly as the triage records — read off one geometry that was fixed before the data were consulted. That is the sense, precise and not rhetorical, in which GG-Theory cosmology stands in the lineage of the *Elements* (Arisaka, 2026, A4-Euclid): the cosmic fractions are not assembled from a toolbox; they descend from a constitution.

Remark 2 (Forcing, not searching)**.**

A numerological fit draws one formula per cosmic number from a free alphabet. GG-Theory reads the whole π-locked inventory from *one* closed ledger $(1, 3, 5; 7, 35)$ with (π, e) — a ledger pinned uniquely, among the ten closed tuples in range, by the color count $N = N_c = 3$ and the gauge group SU(5) fixed *upstream* in the particle sector — fixes each observable's *form* by the stratum it inhabits — a rule stated before the value is compared, and the answer to the look-elsewhere charge, since an unrestricted bounded-cost search matches a random $\mathcal{O}(1)$ target with probability near one and proves nothing — the exponential predictions ($z_* = e^7$, $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-35}$) admitting no nearby-integer alternative. (426)

The same integers $J = 7$ and $K = 35$ reused coherently across recombination, dark matter, and the damping tail, the ledger pinned upstream by $(N = 3, L = 5)$, each form assigned by its stratum before comparison, the exponential predictions unsearchable, and the forward falsifiable commitments are properties no free recombination of a small alphabet possesses. The look-elsewhere effect is not claimed to be zero; it is disarmed by a rule — every form assigned by its stratum, stated before the value is compared — because the Monte-Carlo an earlier version promised has been run and shows that an unrestricted bounded-cost search establishes nothing.

86 GG-distinct comparison: what makes GG-Theory cosmology structurally different

This section places the present framework alongside the principal alternative cosmologies of the contemporary literature and identifies the structural differences that distinguish it from each. The comparison is along three axes: the number of continuous free parameters; the closure mode of the cosmic inventory; and the falsification profile.

Table 47: GG-distinct comparison. Continuous free parameters refers to the parameters that must be specified empirically rather than derived; closure mode is how the present-epoch inventory is fixed; falsification profile lists the principal mode by which the framework can be ruled out.

Framework	Continuous free params	Closure mode	Falsification profile
GG-A10-Cosmos (this paper)	**0**	Closed-packet topology + OGN-5 integer ledger + GU-5 + (π, e)	8 falsifiable observational tests (Appendix T); any single failure refutes
ΛCDM (vanilla 6-param)	6	Empirical fit	Single-parameter precision improvements; no structural falsification
Early dark energy (EDE)	9–11	Empirical fit + EDE template parameters	Improved CMB+SH0ES joint fit
MOND / modified gravity	1–3 (acceleration scale + interpolation function)	Phenomenological	High-z galactic rotation curves, gravitational lensing
Brane cosmology (RS, LR)	4+ (brane tension, bulk curvature,...)	Brane geometry + RS warp factor	High-energy GW signatures, brane-bound matter spectrum
Conformal-cyclic cosmology (Penrose)	0–2 (depending on variant)	Conformal infinity + remote-past null cone matching	B-mode signatures, CCC "Hawking points"

The structural difference is the zero-free-parameter discipline. What distinguishes GG-A10-Cosmos from every framework in the comparison table is the L45 zero-free-parameters discipline

of GG-A4-Euclid: the input ledger contains zero continuous adjustable knobs. ΛCDM with six free parameters can be fit to any data set within its compute envelope; the present framework cannot be retrofitted at all — if the OGN-5 integer ledger $(1, 3, 5; 7, 35)$ does not yield the empirical Ω_m, Ω_b, H_0, or $z_{\rm rec}$, the framework is wrong.

The closure mode is the second structural difference. ΛCDM, EDE, and MOND all close the cosmic inventory by empirical fit (the six or nine or three parameters are observationally determined). Brane cosmology and conformal-cyclic cosmology close the inventory by appealing to features of an underlying higher-dimensional geometry, but the resulting parameter sheet at the cosmological-boundary stratum is still empirically fit. Only the GG-Theory framework closes the inventory by closed-packet topology on the cosmological-boundary stratum with the OGN-5 integer ledger as the fixed input.

Falsifiability comparison. Eight named observational tests (Appendix T) commit the framework to specific predictions on the next-decade horizon. Each test would refute either a named conjecture or a named theorem of A10-Cosmos. In contrast, ΛCDM and EDE absorb improvements in measurement precision into refined parameter values without structural revision; MOND has historically expanded its interpolation function to accommodate new data. The GG-Theory framework cannot do that: every observable is named and locked.

After all, the question is not whether the present framework is right, but whether it is the kind of framework that could be wrong. Eight named falsification tests with named conjectures and named theorems each say yes.

One sheet, twenty numbers, zero continuous dials: that is this Part in seven words, and the philosophical content of such over-determination is easy to state and hard to overstate. A framework this rigid cannot be nursed through a bad measurement, so its survival to date is evidence of a different kind than a fit's success — and its failure, if failure comes, will be clean, public, and instructive. Either way, the subject wins.

What remains is to step back from the sheet and ask what it all means — for cosmology, for the wider program, and for the reader deciding how seriously to take it. That wider view — significance, innovations, accomplishments, and the decade of predictions — is the business of Part XIII.

87 Part XII in plain English: how this paper is held to the rest of the cohort

A twenty-two-paper program can drift. A number quoted in one paper moves in another, nobody notices because no single reader holds all of them at once, and the framework slowly becomes a set of overlapping approximations that agree with each other only in outline. The coherence map in this Part is the mechanism against that, and it is a mechanism rather than a courtesy: if a value here disagrees with the paper that owns it, the map is where the disagreement shows, and the owning paper wins.

The post-hoc-fit defense is the section a hostile reader should attack first, and it is written knowing that. A closed form found after the measurement is known is worth less than one predicted before, and the paper does not pretend otherwise. What it argues is narrower: that the forms are constrained by a cost principle stated independently of the values, that the search space is declared, and that the same principle applied to other observables makes predictions that have not been tuned. Whether the argument succeeds is for the reader to judge, and the section is deliberately placed where it cannot be missed.

The comparison with other frameworks at the end of the Part is structured so that every entry, this one included, carries a cost column. A comparison in which only the rivals pay for their successes is not a comparison.

Part XIII
Discussions

Significance and uniqueness of the result; comparison with other cosmological theories and models; predictions for falsification; open questions that remain; and the future prospects of the GG-Theory program on the cosmological-boundary stratum

88 What this Part does

GG-A10-Cosmos takes the seventeen largely independent mysteries of modern cosmology and compresses them onto a single inherited ledger: the same constitutional layer, the same BI-6 anomaly accounting, the same OGN arithmetic, the same radial scale coordinate, and the same open-system measurement logic carry inflation, the dark sector, the late inventory, baryogenesis, structure growth, and even the black-hole information story at once. This Part states that case in four movements, in the order a skeptical reader should meet them. *Significance* (Sections 89–90): what kind of theory this is — one engine, OGGEP → GDOS → GRIP → GGA, read on the cosmological stratum — and how it differs from every framework now in the field. *Innovations* (Section 91): the ten structural moves the paper makes that no prior cosmology program has made, each with its cohort-global identifier. *Accomplishments* (Sections 92–93): what the paper has actually resolved, walked problem by problem against the seventeen-problem census of Part I, with the honest boundary — what is proved, what is conditional, what remains open — drawn exactly. *Predictions for failure* (Sections 94–95): how the theory can lose, stated as named kill criteria with a decade timeline, because a zero-parameter framework earns credibility only by exposure. The closing verdict (Section 96) can then be read at face value. This is a unification proposal that knows precisely how it can lose — and says so head-high.

89 Significance and uniqueness

GG-A10-Cosmos does not merely offer thirteen proposed answers; it reduces the number of truly independent ingredients one must tolerate in order to discuss those thirteen questions at all. That reduction rests on a short inherited ledger,

$$\begin{aligned}\mathfrak{L}_{\text{cosmic input}} = \Big\{ & E_{\text{GUT}},\ r = \ln(E_{\text{GUT}}/\mu),\ M_{\text{lock}},\ (a,b,c,d) = (1,-1,-2/5,-2/5), \\ & (M,N,L;J,K) = (1,3,5;7,35),\ \mathsf{GFF},\ \mathsf{DGI},\ \mathsf{GSSB},\ \mathsf{GA},\ \mathsf{OGGEP}\Big\},\end{aligned} \tag{427}$$

which organizes into five modules that recur across sectors. Before the five distinctions, it is worth naming the single chain from which they all descend; each distinction is then a facet of it, and each is stated boldly.

89.1 The essence of GG-Theory: OGGEP → GDOS → GRIP → GGA

Beneath the five distinctions below lies a single chain, and it is the essence of GG-Theory as a whole — not of cosmology alone. Axiom AX-3 of the GG-Constitution (Arisaka, 2026, A1-Constitution) is OGGEP: every observable is a boundary projection of a closed higher-dimensional geometry, never the bulk read directly. GDOS (Arisaka, 2026, A5-GDOS) is the law by which that open boundary evolves. GRIP (Arisaka, 2026, A6-GRIP) is the runtime that acts on it — the cascade DGI → GSSB → GA that explores the admissible branches, ranks them by the GFF free-energy ledger, breaks the symmetric branch, and relaxes to the unique Ground Geometric Attractor. The observed quantity *is* that attractor's boundary projection:

$$\underbrace{\text{OGGEP}}_{\text{AX-3}} \longrightarrow \underbrace{\text{GDOS}}_{\text{boundary law}} \longrightarrow \underbrace{\text{GRIP}}_{\text{DGI}\to\text{GSSB}\to\text{GA}} \longrightarrow \underbrace{\text{GGA}}_{\text{the observable}}$$

This chain is not a cosmological device. It is the same engine that fixes the Standard Model in GG-A9-SM — every charged-lepton mass, every mixing angle, the weak angle $\sin^2\theta_W = 25/108$, each a GGA on its own particle sub-stratum — and the same engine that fixes ΛCDM here: every entry of the π-locked inventory of Figure 1, the one Hubble background and its switch, the Kaxion mass, and the damping knee, each a GGA on the cosmological-boundary stratum. The two consequence papers of the cohort, GG-A9-SM (particles) and GG-A10-Cosmos (cosmos), are therefore not two theories but one chain run on two strata; the codebook triage of Part IX is, read in this light, a census of how far the GRIP-uniqueness of each attractor has been proved. The five distinctions that follow are five facets of this one chain.

89.2 It begins from a constitutional layer

Most cosmological model-building begins with an effective field theory or an energy budget and only afterwards asks how measurement, coarse-graining, and admissibility should be described. Here the order is reversed: the axiom and postulate layer of GG-A1-Constitution is inherited *before* any cosmological sector. OGGEP therefore makes the measurement dictionary part of the theory rather than meta-language bolted on later. This matters directly, because cosmological inference is never raw access to "the universe itself" but always an inference from light cones, maps, transfer functions, and coarse-grained boundary observables. Writing that operational structure in from the start prevents the usual drift in which a model is judged by fit quality alone while its measurement semantics stay implicit — and it makes the black-hole information problem a direct test of whether the constitution was chosen correctly, not a late appendix.

89.3 It geometrizes gauge structure at $E_{\rm GUT}$, not gravity at $E_{\rm Pl}$

Twentieth-century habit taught theorists to start from quantum fields and add gravity, or to start from classical gravity and quantize geometry. GG-6 does neither. Its six-dimensional bulk is organized at the GUT edge, with the radial coordinate encoding the energy descent and the compact angle carrying the locked gauge structure. The beginning is therefore placed at $E_{\rm GUT}$, and the Planck scale becomes derived and effective rather than primary. Inflation is then sited on

the derived plateau one e-fold below the GUT edge — placed by the geometry rather than chosen from scalar-potential folklore — with the TeV lock as the descent's second saddle, and the late dark-energy remnant is read as a tiny residual of the same ledger. That is a genuinely different organizational principle, not a rhetorical relabeling.

89.4 Dark matter, inflation exit, and dark energy are linked, not merely coexisting

One descent does three jobs at its two crossings and its terminus: the graceful exit at the top ends the inflationary plateau, the lock crossing opens the birth window for the Kaxion dark sector, and the terminus leaves the small residual vacuum remnant later read as dark energy. Standard practice represents these by three separate sectors — an inflaton potential, a dark-matter candidate, a cosmological constant or quintessence field — and even broad models keep them only loosely related. In GG-A10-Cosmos the same geometric descent seeds all three. This is exactly the cross-sector economy a unification program should attempt, and it deliberately produces a riskier theory: a failure in one sector reverberates through the others.

89.5 Observation itself is treated as open-system projection

This is the GDOS contribution in its cleanest cosmological form. The standard workflow acknowledges openness only pragmatically — detectors are noisy, foregrounds contaminate maps, horizons leak — while still treating the deep laws as those of a closed, observer-external universe. GG-A10-Cosmos instead insists that the observable four-dimensional world *is* a boundary subsystem, so that projection, incomplete access, and boundary-accessible records are built into what observation means. That single grammar then covers cosmological perturbations, black-hole horizons, and the larger floors of the program with one consistent semantics.

89.6 It is austere in a correlated way

Many models advertise minimality; the interesting question is whether the minimality is *correlated*. Anyone can be minimal by discussing one subproblem. The hard thing is to keep the ingredient count low while speaking to inflation, dark matter, dark energy, baryogenesis, BBN, structure growth, early galaxies, primordial relics, and black holes simultaneously. The OGN arithmetic, the BI-6 coefficient vector, the radial coordinate, and the lock seam recur across all of them. This is the strongest form of the claim: *no single one of these distinctions is unique on its own — the conjunction is.* A theory that can break in many places for the same reason is the kind of theory worth testing.

Table 48 shows the concrete sense in which a handful of inherited modules carry many problems at once.

Table 48: How a small number of inherited modules address multiple cosmological problems simultaneously. This is the concrete sense in which the paper aims at unification rather than thematic breadth. The columns give the module, its mathematical role, and the cosmological problems it touches. The economy claimed here is the third column being longer than the first: if each problem needed its own module the framework would be a catalog rather than an account.

Inherited module	**Primary mathematical role**	**Cosmological problems touched**
Finite E_{GUT} edge and radial coordinate $r = \ln(E_{\mathrm{GUT}}/\mu)$	Replaces singular-beginning rhetoric by finite scale descent and organizes chronology logarithmically	1, 2, 6, 7, 13; also the entire chronology of Part VI
Lock seam near M_{lock}	The descent's second saddle: the prompt crossing that stocks the dark sector, and the seesaw source of the terminus scale	2, 3, 4, 5, 6, 13
Compact-angle sector	Supplies a natural dark-matter window and links it to the same transition that ends inflation	5, 10, 12, indirectly 8 and 9
Late residual GFF inventory	Produces the π-locked matter-vacuum ledger and the coincidence interpretation	3, 4, 8, 9
Constrained neutrino/CP ledger	Carries leptogenesis and transmits consequences to the visible-matter abundance audit	6, 7, indirectly 8 and 9
OGGEP/GDOS semantics	Makes observation a boundary-accessible open-system projection and prevents naive closed-system factorization	11, 13; also supports 8–10 as an inference dictionary
Exact baryonic $\mathrm{U}(1)_B$ sector	Protects proton stability and introduces weak baryonic environmental effects at late times	8, 9, 10, and the particle-cosmos bridge more generally

If the late-inventory numbers fail, dark energy, coincidence, and the late-background interpretation fail together; if the compact-angle story fails, the dark-matter window and the coherence of the lock-seam narrative fall with it; if the open-system horizon story fails, black-hole information is not the only casualty. That correlated exposure is what it means for a theory to pay a genuine price for being wrong.

90 Comparison with the leading frameworks

Read against the leading frameworks, each rival's famous gap turns out to be a *missing-layer symptom* — the absence of precisely the constitutional, geometric, or measurement layer GG-A10-Cosmos supplies.

Vanilla ΛCDM. The nearest comparator is the baseline the paper explicitly preserves. ΛCDM is phenomenologically compact and observationally superb: a stable background, a precise perturbation

pipeline, a common empirical language. Its limitation is not poor fit but interpretive openness — cold dark matter, the near-scale-invariant seed, and the late vacuum component are represented beautifully yet left underived. ΛCDM is a superb baseline and a weak closure principle; GG-A10-Cosmos keeps the first and supplies the second.

Scalar-potential inflation EFTs. A flexible perturbation dictionary with wide reach (n_s, running, non-Gaussianity, tensor ratios), bought with a very large family of potentials and reheating prescriptions and almost no sectoral closure. The inflationary module of GG-A10-Cosmos is the opposite strategy: the plateau is tied to the radial geometry, the exit to the plateau's own steep edge, and the same descent seeds the dark sector (at the lock crossing) and the vacuum remnant (at the terminus). The EFT tradition buys robustness with many potentials; GG-A10-Cosmos buys unification with very few — and accepts the sharp tensor risk that follows.

WIMP/SUSY dark matter. Natural freeze-out lore and broad BSM embedding, paid for with a candidate zoo, rising experimental pressure, and a dark sector left independent of inflation and dark energy. The missing layer is a geometric selection principle. GG-A10-Cosmos ties dark matter to the compact-angle/Kaxion window descending from the same anomaly structure required elsewhere; the gain is a sharper search target, the cost a sharper falsification risk.

PQ axions, ALPs, and the axiverse. The nearest relative: a light pseudo-scalar tied to symmetry structure, cold via misalignment. But a generic axiverse survives precisely by being broad. The Kaxion is narrower in two ways: it is not introduced because ALPs are useful but inherited from the compact-angle and lock-seam story of this specific geometry, and its targets — a window of seven rungs, a DFSZ coupling, a count of one — reflect the same discrete arithmetic used across the program, around one declared calibration. A clear gain if borne out, and a clear loss of flexibility.

Quintessence, dynamical dark energy, and EDE. Room to alter background history or growth without discarding GR, paid for by enlarged sectoral freedom: a new field or fluid, then its potential, then its couplings and clustering. The minimal late-time branch of GG-A10-Cosmos instead wants $w = -1$ to high accuracy, negligible clustering, and a present inventory tied to simple π-locked shares. It does not compete by mimicking all dark-energy phenomenology; it competes by saying that if the universe really prefers something very close to a true cosmological constant, one should not reopen a large dynamical-fluid sector.

String landscape and multiverse. Enormous UV richness and vast vacuum multiplicity, paid for with a vacuum-selection problem in which freedom can outrun falsifiability — the anthropic Λ is the symptom of a missing zero-free-parameter constitution. GG-A10-Cosmos shares the mathematical atmosphere but leans on MSMF and OGN austerity, rejecting a large physically operative landscape or anthropic selection as a primary explanatory tool. Whether that rejection is sustainable is an empirical and mathematical question; it is also a sharp philosophical distinction.

Bounce, cyclic, ekpyrotic, loop-inspired beginnings. These engage the origin problem directly but often introduce a second pre-history or Planck-regime dynamics with its own assumptions. GG-A10-Cosmos instead says the physically relevant ultraviolet edge is already finite at E_{GUT}, the Planck scale derived rather than primary. This is conceptually cleaner than retaining a singularity, but it is honestly a finite beginning *ledger*, not yet a full beginning wavefunction.

Operational black-hole programs and the island formula. The island/QES and replica-wormhole literature delivers explicit Page-curve recovery and generalized-entropy technology, usually focused on the horizon rather than a universal measurement grammar. GG-A10-Cosmos does not overthrow it — it reinterprets it: the success of generalized-entropy reasoning is the expected behavior of a theory in which horizons define open subsystems, naive factorization is illegal, and the accessible ledger is already GFF/GDOS-like. The gain is a broader operational grammar; the obligation, made explicit below, is to make that account quantitatively competitive.

These trade-offs are condensed in Table 49.

Table 49: Comparison between GG-6/GDOS cosmology and several leading theoretical habits. The emphasis is on explanatory trade-offs rather than rhetorical victory. The second and third columns are deliberately paired: what each framework does well, and what it pays for that success. The fourth states the contrast rather than a verdict, and the row for this framework carries a cost column like every other.

Framework	**What it does well**	**What it pays for that success**	**GG-6/GDOS contrast**
Vanilla ΛCDM	Compact empirical baseline; excellent background and perturbation fit across a huge data range	Leaves several source terms and origins interpretively open	GG-6/GDOS keeps the baseline equations but tries to derive or correlate the unexplained source sectors
Scalar-potential inflation EFT	Flexible perturbation dictionary; wide phenomenological reach	Large family of potentials and reheating prescriptions; little sectoral closure by itself	GG-6/GDOS offers a much narrower lock-seam inflation module with a much sharper tensor risk
WIMP/SUSY dark matter	Natural freeze-out lore; broad BSM embedding	Candidate zoo; growing experimental pressure; dark matter often remains independent of inflation and DE	GG-6/GDOS narrows the dark sector to the compact-angle/Kaxion window and links it to inflation exit and vacuum remnant
Axion/ALP/axiverse families	Symmetry-based pseudo-scalar dark matter; rich search program	Broad parameter freedom unless tied to a sharper UV ledger	The Kaxion is intended as a far narrower axion-like window fixed by the same geometry used elsewhere

Framework	What it does well	What it pays for that success	GG-6/GDOS contrast
Quintessence / dynamical DE / EDE	Can modify late or early background histories and relieve specific tensions	Requires new fields or fluids with their own potentials, couplings, and clustering assumptions	GG-6/GDOS prefers a residual vacuum remnant with $w \approx -1$ and only transient early-energy faces inherited from the same geometry
Landscape / multiverse approaches	Enormous UV richness; many possible vacua and sectors	Difficult vacuum selection; freedom can outrun falsifiability	GG-6/GDOS explicitly prefers austere discrete locking and rejects broad anthropic rescue as a primary explanatory tool
Bounce / cyclic / loop-inspired origins	Direct engagement with the beginning and singularity issue	Often introduces a second pre-history or Planck-regime dynamics with its own assumptions	GG-6/GDOS replaces the operative beginning by a finite E_{GUT} edge rather than by a bounce or Planck-era discrete recursion
Island/QES black-hole programs	Quantitative Page-curve recovery and generalized-entropy technology	Usually focused on the horizon problem rather than a universal open-system measurement grammar	GG-6/GDOS seeks to reinterpret the same success as a special case of a broader OGGEP/GDOS operational closure

GG-A10-Cosmos does not dominate every alternative on every axis — that is not how theory choice works. Its advantage is economy and cross-linkage; its disadvantage is that its sharpest modules are also its most falsifiable. For a unification paper, that is the right trade.

91 The Ten Structural Innovations of the Cosmological Closure

The contribution of GG-A10-Cosmos is to read the four-dimensional ΛCDM cosmology as the cosmological-boundary projection of the six-dimensional bulk under GDOS/OGGEP, recovering the present-epoch inventory and the cosmic chronology as consequences of an inherited ledger rather than as fitted parameters. We collect here the ten structural innovations that reading delivers, each carrying its cohort-global identifier (§D). None introduces a new continuous parameter; each descends from the GG-Theory constitution of GG-A1-Constitution through the BI-6 closure and the GRIP cascade. Each is, at root, the one chain of Section 89.1 — OGGEP → GDOS → GRIP → GGA — specialized to a single cosmological observable.

Innovation 1 (I-Cosmos-1): Cosmology from a constitutional layer. GG-A10-Cosmos reverses the usual order of cosmological model building. Instead of beginning from an effective field theory or a phenomenological energy budget and asking only afterward how measurement and

admissibility should be described, it inherits the axiom and postulate layer of GG-A1-Constitution before it inherits any cosmological sector at all. Every cosmic claim is therefore admissible only if it traces backward through the same constitutional chain already used to build the particle and open-system papers, so that fit quality alone is never permitted to stand in for legality of origin.

Innovation 2 (I-Cosmos-2): The π-locked inventory as a closed-packet theorem. The present-epoch densities $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, and $\Omega_b = 1/(2\pi^2)$ are read as a closed-packet topological invariant of the BI-6 four-form, with the OGN-5 integer ledger $(1, 3, 5; 7, 35)$ of GG-A4-Euclid fixing the rational weight of the matter cell before the Pontryagin rescaling. The inventory is the integral output of a fixed topological computation with no continuous adjustable parameter, not a curve fit dressed in geometric language.

Innovation 3 (I-Cosmos-3): One topology forces three coupled entries. A post-hoc fit explains one number; the closed-packet integration here forces three coupled numbers at once. The same computation that yields $\Omega_m = 1/\pi$ yields $\Omega_\Lambda = 1 - 1/\pi$ by unit-trace complementarity and $\Omega_b = 1/(2\pi^2)$ by a one-dimensional further projection through the $U(1)_B$ baryonic bundle. That the topology produces the specific combination of powers of π — one for matter, its complement for vacuum, two for baryons — is the structural signature that distinguishes a derivation from a coincidence.

Innovation 4 (I-Cosmos-4): Recombination and the Hubble switch from integers. The corridor-switch redshift $z_* = e^7$ — the topological last-scattering landmark, 0.62% above the measured decoupling epoch, annotated and not scored (§42.1) — descends from the OGN-5 integer $J = 7$, not from a fitted thermal history, and the one background is read through the sound-horizon dictionary on one side of that integer-locked scale and by the ladder-free late probes on the other. The early-time determination of H_0 thereby sits on the derived background (66.8 ± 0.2 with zero sky input, 66.6 ± 0.2 from the π-locked inventory), while the local-ladder excess is read whole as a calibration systematic — a displaced measurement, not a tension to be relaxed by an extra empirical sector.

Innovation 5 (I-Cosmos-5): Gauge geometry at the GUT edge, not gravity at the Planck scale. The six-dimensional bulk is organized at the GUT edge, with the radial coordinate $r = \ln(E_{\mathrm{GUT}}/\mu)$ encoding the energy descent and the compact angle carrying the locked gauge structure, so that the Planck scale becomes derived and effective rather than primary. Inflation is placed on the derived plateau one e-fold below the GUT edge — its class forced by GDOS, not chosen from scalar-potential folklore — and the late dark-energy remnant is read as a tiny residual of the same ledger rather than as an independent vacuum constant.

Innovation 6 (I-Cosmos-6): One descent seeds three sectors. A single geometric descent does three jobs at its two crossings and its terminus: the graceful exit ends the inflationary plateau, the lock crossing opens the birth window for the Kaxion dark sector, and the terminus leaves the small residual vacuum remnant later interpreted as dark energy. Where standard practice carries an

inflaton sector, a dark-matter sector, and a cosmological-constant sector as loosely related modules, GG-A10-Cosmos links them so tightly that a failure in any one reverberates through the others.

Innovation 7 (I-Cosmos-7): Kaxion dark matter from the compact angle. The dark-matter candidate is not an unconstrained zoo but the Kaxion window inherited from the compact-angle sector and the same final geometric transition that terminates inflation. The relic abundance follows from misalignment within the π-locked inventory, so that the candidate points to a narrow frequency-and-coupling window with concrete search targets rather than to a broad axiverse landscape.

Innovation 8 (I-Cosmos-8): Observation as open-system boundary projection. The framework treats the observable four-dimensional world as a boundary subsystem rather than as a closed observer-external universe. Through the GDOS boundary law and the OGGEP projection, incomplete access, coarse-grained records, and horizon physics are built into what cosmological observation means, so that the semantics of the CMB, of late-time inference tensions, and of the black-hole information paradox are all discussed within one grammar.

Innovation 9 (I-Cosmos-9): The universal-cosmic-identity principle. The cosmological coincidence is read as a consequence of the same constitutional uniqueness principle that selects the unique stable Geometric Attractor on the charged-lepton stratum of GG-A9-SM, propagated one floor up to the cosmological-boundary stratum. The same logic that explains why every electron is the same explains why the universe has the inventory it does, read through the GRIP cascade — with the GRIP-Out transitions that rebuild the landscape separated by long GRIP-In relaxations on it.

Innovation 10 (I-Cosmos-10): A zero-free-parameter, correlated, and falsifiable cosmology. The input ledger contains zero continuous adjustable knobs: the inventory is closed by closed-packet topology with the OGN-5 ledger and GU-5 as fixed input, so the framework cannot be retrofitted the way a six-parameter ΛCDM fit can. Because the same short ledger controls many observables at once, the theory is falsifiable in a stronger-than-usual sense — eight named observational tests, each tied to a named conjecture or theorem, can retire it as a package rather than only patch it sector by sector.

92 Accomplishments: the seventeen problems, resolved, conditional, or frontier

Significance and innovation are claims about *kind*; accomplishment is a claim about *score*, and it should be settled against the census the paper opened with. Table 1 of Part I listed the seventeen open problems of contemporary cosmology in five clusters. Here is what the present framework has done to each — with the codebook status attached, so that the score and the honesty travel

together. (The reader who wants any single entry in depth will find the full treatment in the Part named; this section is the tally.)

Cluster I — the two flagship tensions. *(1) The Hubble tension* is read as one background rate, 66.8 ± 0.2, derived with zero sky input and confirmed at 66.6 from the π-locked inventory, read by every ladder-free probe; the distance ladder's 73.5 is the predicted outlier, a calibration systematic of about 9% that the framework does not derive, and the gravitational-wave sirens — which bypass both calibrations — are predicted to land near 67 (Part VI; D* under Lemma II). *(2) The S_8/σ_8 growth tension* is not resolved: $\sigma_8 = 0.80 \pm 0.01$ and $S_8 = 0.825 \pm 0.010$ are computed from the declared amplitude on the π-locked background (Part VI; D* under Lemma V′), and the framework shares the tension with cosmic shear at $+2.5\sigma$. One background, both tensions stated as they stand.

Cluster II — the dark Universe and its fate. *(3) The cosmological-constant problem* is answered in two registers: the scale by the seesaw $m_\Lambda = M_{\rm lock}^2/M_{\rm Pl,red} \simeq 6$ meV with the ledger coefficient $\zeta_\Lambda = \pi/LK$, which recovers the notorious 120 orders of magnitude to the percent level without a cancellation (Derived cond.), and the share — exactly, by geometry alone — by $\Omega_\Lambda = 1 - 1/\pi$ (Part IV; D* under Conjecture VIII). *(4) Constant or evolving dark energy?* The framework takes a live, falsifiable stand: $w = -1$ exactly, no evolution, no clustering — directly opposed to the DESI $w_0 w_a$ hint, and decidable this decade (D* under Conjecture IX). *(5) The cosmic coincidence* is dissolved rather than explained away: the matter-composition ratio and the vacuum scale hold at every epoch, the fractions are read at the present epoch through one declared timing input, and the residual "why now" is bounded and named (Parts IV and X). *(6) The nature of dark matter* is answered with a particle, a mass, and an address: the Kaxion at $8.1\,\mu$eV in the compact angle, one scale calibrated, pinned patch-uniform during primary inflation by conjecture, searchable at 1.59–2.13 GHz (Part V and GG-C1-Kaxion; D* under Conjectures I and III).

Cluster III — the beginning and the primordial signals. *(7) The origin and the initial singularity*: the universe begins lawfully at the finite $E_{\rm GUT}$ edge — no singular datum, no Planck-regime speculation; the beginning is the top of a descent, not a breakdown of the equations (Part III; derived). *(8) The cause of inflation* is the paper's sharpest accomplishment: the inflaton is the Σ_r scale modulus (forced), the plateau class is derived (GDOS non-conformal route), and the tilt is computed, not fitted — $n_s = 0.9634$, 0.4σ from Planck and 3.2σ below the 2025 DESI-combined value, the one currently failing test (Part IV, §28.10; D* under the revised Lemma IV, with the amplitude the one declared input). *(9) Primordial gravitational waves*: the same plateau forces $r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order) — a decisive, two-sided forward prediction, above the Simons-Observatory/LiteBIRD floor (D*, with the tilt). *(10) Baryogenesis*: the derived reheating temperature $T_{\rm rh} \simeq 6 \times 10^{10}$ GeV guarantees the bath in which the TeV-scale heavy states equilibrate through their own inverse decays, so the resonant mechanism runs on a full bath on the inherited neutrino-and-CP ledger of GG-A9-SM with zero dilution (Part V; D* under Conjecture VI). *(11) The lithium-7 problem*: an open-system $A = 7$ port suppressing the sector by $1/\pi$ while leaving deuterium and helium untouched (Part V; D* under Conjecture X).

Cluster IV — structure, cosmic dawn, and the JWST surprises. *(12) JWST's early massive galaxies* are addressed by the ledger's baryonic levers (the corridor route and the streaming damping withdrawn) — a structural reading whose quantitative simulation is named open work (Part VI). *(13) The "little red dots"* and *(14) cosmic dawn / 21 cm* are carried honestly as *frontier*: the framework touches them but does not close them. *(15) The black-hole information paradox* is resolved conceptually — the horizon is an OGGEP port, the entropy is the induced $A/4G_N$ with the scale-derived Newton constant, and no firewall or remnant is needed — with the quantitative match to the island-formula Page-curve numbers named as open work (Part VII; PS-Cosmos-18).

Cluster V — cracks in the foundations. *(16) The isotropy anomalies* remain *frontier*: boundary-projection statistics are the natural language, but no closed prediction is offered. *(17) Neutrino masses and* $N_{\rm eff}$: inherited whole from the particle ledger of GG-A9-SM, with $\Delta N_{\rm eff} \leq 0.03$ — a derived null to within the computed freeze-in ceiling, not a fitted offset (D).

The tally, honestly drawn. Of the seventeen: *fourteen* receive a derived, derived-conditional, or structural reading from the one inherited ledger — among them both flagship tensions, all four dark-universe problems, and all five beginning-and-primordial problems — and *three* (the little red dots, cosmic dawn, the isotropy anomalies) are marked frontier, touched but not claimed. No previous framework known to the author addresses more than a few of the seventeen at once, and none does so with zero continuous parameters. That is the accomplishment claim, and the two sections that follow draw its honest boundary: exactly what is proved, what rides on a named conjecture, and what the framework computes versus what it only frames.

93 The honest boundary: what is proved, what is conditional, what is open

This is the concentrated honest frontier of the Part. It is stated plainly, and without flinching, because a scientist who has mapped exactly where the proven territory ends loses nothing by naming the boundary.

93.1 What is proved, what is suggested, what remains open

Proved (within the present scope). The lawful finite beginning at $E_{\rm GUT}$ replacing the Planck-singularity datum; the slow-smooth plateau, derived in class (GDOS) and placed one e-fold below the GUT edge, with its computed observables $n_s = 0.9634$ and $r = 0.0038$, both derived-conditional on the revised Lemma IV; and the switch redshift $z_* = e^7$ as an OGN-5 integer-ledger prediction (the two-stage corridor once built on it is withdrawn, Part VI).

Suggested but not proved. The π-locked inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$ is carried as *Derived-conditional* under Conjecture VIII (closed-packet integration): the integer arithmetic of Appendix L.8's Step 4 does not close, and the Hopf-ratio Step 4 of Part X closes only into the calibrated epoch (§74.3), so these fractions must not be advertised as proved. The

inflationary amplitude A_s remains the launch sector's one declared input (Asserted), pending the GDOS-heat-kernel computation of Lemma V′; the tilt n_s and the scale $V_{\rm inf}^{1/4}$ are Derived-conditional on the revised Lemma IV. The one background rate against its two comparators is a structural identification; S_8 is computed from the declared amplitude and shares the tension with cosmic shear ($+2.5\sigma$) rather than resolving it. The black-hole resolution is conceptual, not yet a quantitative reproduction of the island/replica Page-curve numbers.

The conditions ledger. The two-basin derivation replaced the single word "asserted" with a short list of *named* conditionals, and the honest frontier is now this list. (i) *Plateau-class-and-prompt-crossing* (Lemma IV, revised): the GDOS non-conformal derivation of the Starobinsky/ξ class, and the GRIP ledger-action bound on Σ_r giving $N_2 = 0$ — underwrites n_s, r, and $V_{\rm inf}^{1/4}$. (ii) *GDOS-heat-kernel amplitude* (Lemma V′): the one computation that would pay the amplitude A_s and make the launch input-free; its warp candidate stands at 75–87%, and two alternative routes have been computed and closed. (iii) *Conjecture VIII* (closed-packet): the π-locked fractions, unchanged. (iv) The Kaxion's deterministic misalignment rides on the inflation-era pinning conjecture C-Kaxion-d (Part V) — a *necessary* condition, since the unpinned history is excluded by the microwave-background isocurvature bound at every decay constant the compact sector supplies — and its absolute placement on one declared calibration (GG-C1-Kaxion). Each condition is a target with a stated failure mode, not a soft caveat; items (i) and (ii) are the two endgames of the launch sector, and either would upgrade the bottom line materially.

Genuinely open / unfinished. A full perturbation theory rederived from the six-dimensional action rather than layered on the four-dimensional baseline (PS-Cosmos-14); a full BBN-network implementation of the A=7 open-system port, an open item that no companion in the cohort carries (PS-Cosmos-15); a genuine global Einstein–Boltzmann fit of the Hubble/S_8 package across CMB, BAO, SN, weak lensing, RSD, and clusters (PS-Cosmos-16); nonlinear simulations for the JWST early-galaxy story (PS-Cosmos-17); the tight match of the GDOS/OGGEP black-hole account to generalized-entropy computations (PS-Cosmos-18); an even cleaner first-principles derivation of the π-locked inventory (PS-Cosmos-19); a complete dictionary between the Q_{Qi} charge ledger and ordinary information measures (PS-Cosmos-20); and the preservation of precision across the larger cross-floor program, judged here first as a cosmology paper (PS-Cosmos-21).

93.2 What GG-Theory computes, and what it does not: the cosmology tier ladder

The honest boundary of the framework is best stated as a four-tier ladder — the cosmological instance of the can/cannot ladder established for the Standard Model in GG-A5-GDOS (Theorem T9.9; Arisaka, 2026, A5-GDOS) and used in GG-A9-SM. It says exactly which cosmological content the geometry *computes* and which it only *frames*, leaving the irreducible dynamical computation to the standard machinery. The distinction matters because a parameter-free framework is exposed to the charge that it claims the whole of cosmology; it does not, and the ladder draws the line precisely.

Tier 1 — the inventory (computed). The present-epoch density fractions $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, the recombination switch $z_* = e^7$, and the inherited input ledger (GU-5, the OGN-5 integers $(1, 3, 5; 7, 35)$, and (π, e)) are the content GG-Theory *derives* — or derives-conditionally on a named conjecture (two for the dark-matter share), exactly as the one/nineteen/one codebook triage records. This is the tier on which the paper stakes its claim: the cosmic *inventory* is read off one geometry, not fitted.

Tier 2 — the scale flow (geometrized). The cosmic chronology — the dictionary between t, μ, r, a, and z, the master timeline from the $E_{\rm GUT}$ edge to the de Sitter future (Table 36) — is *geometrized*: the radial coordinate $r = \ln(E_{\rm GUT}/\mu)$ is the renormalization-group and expansion axis, so the background history and its leading-log skeleton are radial transport. The geometry organizes *when* each epoch occurs and fixes the ledger the background Friedmann integration runs on; it does not thereby replace that integration.

Tier 3 — the precision CMB power spectrum (not computed). The full angular power spectrum C_ℓ — the acoustic-peak heights and positions at sub-percent precision, the polarization spectra, the transfer functions — is the output of the Einstein–Boltzmann hierarchy, an irreducible *dynamical* computation over the coupled photon–baryon–neutrino system. GG-Theory fixes the *inputs* that computation takes as given (the inventory, z_*, the background history), but it does *not* compute the spectrum — it hands its inputs to the instrument. That hand-off has now been executed in both directions: run on the exact GG-6 fractions, the Einstein–Boltzmann machinery (CLASS) reproduces the observed TT band-powers at the few-percent level and the damping scale to half a percent ($\ell_D = 1350 \pm 3$ in the envelope convention of Section 40; 1956 against Planck's 1947 ± 5 in CLASS's own $k_D D_M$) — a $\chi^2/\mathrm{dof} \approx 2.3$ whose residual is the 0.3% acoustic-angle deficit, not a likelihood-level endorsement, and we do not claim it as one; the BFFT knee $K(2\pi)^2$ is carried separately as an interpretive overlay. A genuine global Einstein–Boltzmann *likelihood* fit across CMB, BAO, SN, and lensing — covariance kept, seed re-optimized — is named as open work (PS-Cosmos-16); the Boltzmann machinery remains the instrument.

Tier 4 — the nonlinear structure-formation $P(k)$ (not computed). The nonlinear matter power spectrum, the halo mass function, and the baryonic-feedback corrections that govern the JWST early-galaxy regime are irreducible *nonlinear* dynamics — the province of N-body and hydrodynamic simulation. GG-Theory supplies the linear inputs and a structural, scale-sensitive correction chain (the growth-and-tensions Part), giving *structure*, not *precision*; the nonlinear simulations themselves are named as open work (PS-Cosmos-17). As in the Standard-Model case, a holographic geometry helps only where the coupling is controlled; the realistic nonlinear regime is not such a place.

The thesis. The ladder lands one sentence, the cosmological analog of GG-A9-SM's "perturbation theory computes the corrections to nature's constants; GG-Theory computes the constants":

GG-Theory computes the cosmic *inventory* and geometrizes the cosmic *chronology*; it does not compute the precision C_ℓ spectrum or the nonlinear $P(k)$. The Einstein–Boltzmann and N-body machinery remains the instrument it frames but does not replace. (428)

This is not a concession extracted under pressure but the boundary the framework draws for itself. Tiers 1 and 2 are where a zero-parameter geometry can speak; Tiers 3 and 4 are where the irreducible dynamical computation lives, and the paper claims the first pair without pretending to the second. The two are complementary sciences, not competitors: the geometry supplies the inputs the Boltzmann and N-body codes take as given, and those codes supply the precision the geometry does not.

93.3 Dependencies

The chain is not free-standing; it rests on named inherited results, and the framework is only as strong as they are. It depends on the OGN-5 integer ledger (which fixes $z_* = e^7$), on the BI-6 anomaly-free closure, on the GDOS boundary law and OGGEP projection (which supply the measurement semantics and the open-system black-hole account), and on the GG-A4-Euclid openness theorem (which makes the cosmological constant a geometric scale residue rather than a free input). A decisive failure of any of these upstream pillars propagates downstream into the cosmological closure.

93.4 The problem tally

Against the open problems of contemporary cosmology (Table 1, Part I; provenance in the Part IX codebook), GG-A10-Cosmos offers a derived or derived-conditional reading for the great majority. The cleanest closures are the lawful finite origin at E_{GUT}, the dark-energy inventory $\Omega_\Lambda = 1 - 1/\pi$, and the tensor target $r = 0.0038$; most of the remainder carry a named conjecture, a pending calculation, or a companion paper as their explicit route forward; and three items — the "little red dots," the cosmic-dawn 21 cm signal, and the large-angle isotropy anomalies — are honestly marked as present frontier rather than claimed. The asymmetry is the honest research statement, not a weakness to be hidden.

94 The cost of being right: zero parameters, maximal exposure

A correlated, zero-free-parameter theory has the strongest possible relationship to experiment: any single failed prediction can kill it, and it cannot be quietly rescued by reinstalling sectoral freedom. This is the paper's central scientific virtue, and it is stated here without apology. One might object that a flexible patchwork can always fit more data in a narrow likelihood sense. True — and dangerous, because it makes the relation between explanation and fit increasingly opaque. A theory that attaches a new independent sector to every tension drifts toward a catalog of compensations

rather than a causal explanation. GG-A10-Cosmos accepts the opposite burden: a lower tolerance for ad hoc rescue, and therefore a theory that is easier to kill and more informative if it survives.

The input ledger is honest about what is and is not free, on four levels.

L1. Constitutional input. The axioms and postulates of GG-A1-Constitution; these define admissibility and operational meaning.

L2. Geometric input. The six-dimensional bulk, BI-6, and the OGN-5 arithmetic; the core model-structural inputs of GG-A2-GG6 through 5.

L3. Boundary baseline input. Four-dimensional GR and the FRW/ΛCDM observable language as derived or inherited baseline descriptors.

L4. Observational calibration input. The empirical values needed to place the boundary history on actual units — the present CMB temperature, today's Hubble scale.

What the paper refuses is a fifth level of unrelated rescue knobs. That restraint is the discipline that makes the falsification logic below mean something: if the shared modules cannot meet the precision audit, the framework should not be saved by quietly reinstalling the sector-by-sector freedom it set out to remove. A theory that can fail as a package, and declines to escape that fate, has earned the right to be taken seriously.

95 Kill criteria and the honest test ledger

The governing logic is

$$\text{small input ledger} \quad \Longrightarrow \quad \text{large correlated observational ledger.} \tag{429}$$

There are five clean near-term ways to retire the minimal branch.

(a) Primordial tensor modes (F-Cosmos-1). The derived slow-smooth plateau fixes the tensor sector at $r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order). This is a *forward* prediction, and one of the cleanest in the paper: a factor nine below the present bound, several times above the Simons-Observatory/LiteBIRD floor, decisive within the decade. A tensor measurement inconsistent with $r = 0.0038$ in either direction falsifies the derived launch: the launch's own spread, 0.0037–0.0040 across the reheating range, lies below any near-term $\sigma(r)$, so operationally a robust signal outside $[10^{-3}, 10^{-2}]$, or its absence once $\sigma(r) \sim 10^{-3}$ is reached, kills the module, and a signal inside the window but off 0.0038 at the instrument's precision refutes the count (§50.4). It is a two-sided target, not a peripheral detail.

(b) Dark-energy dynamics (F-Cosmos-2). The minimal late branch wants $w = -1$ to high accuracy with no appreciable clustering. A secure, persistent detection of $w \neq -1$, strong time variation in $w(a)$, or significant dark-energy clustering would severely damage the late-inventory story. This is a forward prediction sharpened against existing background-expansion and lensing data.

(c) Kaxion search targets (F-Cosmos-3). The dark-matter module points not to a broad space but to a narrow window: in the companion closure GG-C1-Kaxion, a single laboratory line in the window 1.59–2.13 GHz on the DFSZ coupling line, standing alone, the benchmark rung at 1.95 GHz — one scale calibrated, the ladder an ansatz, the count derived-conditional. This is a forward prediction. A window swept at DFSZ depth and silent kills the calibrated placement and its ladder (F-Kaxion-1c); a second in-window line refutes the count; a confirmed line outside the window kills the placement, and at the randomized-angle mass kills the pinning conjecture and with it this paper's compact-angle route to the dark-matter share.

(d) The π-locked inventory (F-Cosmos-4). In the minimal interpretation the shares follow simple geometric values,

$$\Omega_m = \frac{1}{\pi}, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi}, \qquad \Omega_b = \frac{1}{2\pi^2}. \tag{430}$$

These are carried as *Derived-conditional* entries under the closed-packet-integration conjecture (Conjecture VIII), not as proved outputs; the corridor switch $z_* = e^J = e^7 \approx 1097$ is likewise a parameter-free marker of the OGN-5 integer ledger — a topological quantity distinct from the measured decoupling redshift: the framework's own atomic chain derives $z_{\rm dec}$ on the Planck value at the 0.1% level, with the switch sitting 0.4% above it (§42.1). If improved inference settled decisively on shares incompatible with Eq. (430) even after the observational dictionary and covariance are accounted for, the late-inventory interpretation would fail directly.

(e) Joint growth-tension closure (F-Cosmos-5). The Hubble and S_8 tensions are hard because a model that eases one often worsens the other. GG-A10-Cosmos proposes a combined route, with one derived background rate that the microwave background reads and the calibrated ladder does not (Theorem 68.8), and one computed clustering amplitude that shares the S_8 tension rather than resolving it (Corollary 68.10). The honest test is whether the same inherited modules give an acceptable *joint* fit without degrading the CMB, BAO, or lensing ledger. This is an *audit* of existing Planck/DESI-class data, not a clean forward prediction; on the one background there is no growth response to have a sign, and the computed S_8 is the number the joint fit must carry.

A framework that subsumes ΛCDM must first reproduce its empirical successes before it can differ from it; the five kill criteria are therefore a mixture, honestly labeled. The forward predictions (the $r = 0.0038$ tensor target, the 1.59–2.13 GHz Kaxion window, the $w = -1$ dark-energy behavior) are where the theory most exposes itself to new data; the inventory shares and the Hubble/S_8 route are postdictions and audits of existing data that the theory must not be allowed to overclaim. Table 50 states the full set in "how the theory can lose" language.

Table 50: Practical falsification ledger for the minimal cosmological branch of GG-A10-Cosmos. The table is deliberately written in a "how the theory can lose" language. The columns give the observable domain, the minimal expectation, the class of observation that would test it, and what would strongly damage the branch. The last column is written to be decidable by measurement rather than by argument, which is the only property that makes a ledger like this worth printing.

Observable domain	**Minimal expectation**	**Observation class**	**What would strongly damage the minimal branch?**
Primordial tensors / CMB B-modes	$r = 0.0038$ exact from the derived plateau ($12/N^{*2} = 0.0043$ at leading order) — a factor of nine below the present bound, several times above the Simons-Observatory/LiteBIRD floor	Next-generation CMB polarization programs	A tensor measurement inconsistent with $r = 0.0038$ beyond the launch's spread 0.0037–0.0040, and in any case outside $[10^{-3}, 10^{-2}]$: a detection at $r > 0.01$, or a null at the projected floor
Scalar perturbation sector	$n_s = 0.9634$ exact on the derived plateau (0.9620 at leading order), ± 0.00034 from the count; near-scale-invariant with a conventional acoustic baseline	CMB temperature and polarization, LSS, BAO	A tilt excluding 0.9634 in either direction at 3σ — met today on one of the two 2025 combinations (the DESI-combined value sits 3.2σ above it, the CMB-only value 1.5σ); the framework's one currently failing test
Dark energy equation of state	$w = -1$ to high precision; negligible clustering	Background expansion, lensing, and structure-growth surveys	Secure evidence for time-varying or clustering dark energy
Cosmic inventory	Present-day shares close to the π-locked relations of Eq. (430)	Joint CMB, BAO, SN, and lensing fits	Stable values that cannot be reconciled with the geometric-share interpretation
Dark matter direct search	Kaxion window with a narrow frequency/coupling target, one line alone	Axion/ALP resonant and broadband searches	Exhaustive null results in the sharpest target window
BBN audit	Improved A=7 behavior without spoiling D/H or Y_p	Precision light-element abundance inference	Any lithium "solution" that necessarily ruins deuterium or helium once implemented honestly; a firm primordial ^{7}Li/H away from SBBN/π; upstream of the network, a full transport calculation that finds no finite-temperature regulator for the resonance that sources η_B

Observable domain	Minimal expectation	Observation class	What would strongly damage the minimal branch?
Early/late expansion tension	One background rate, read by every ladder-free probe; the calibrated ladder the predicted outlier	Combined early- and late-universe parameter inference	Persistent failure of the shared-route model to improve the joint fit without new rescue knobs
Growth and lensing	Standard growth on the one background; $S_8 = 0.825$ computed, the shear tension shared	Weak lensing, redshift-space distortions, cluster counts	Need for a wholly separate growth-only sector after the shared modules are imposed
Early galaxy assembly	Mild acceleration of early structure and settling relative to vanilla expectations	High-redshift galaxy and quasar surveys	Strong evidence that the framework's inherited early-structure timing is in the wrong direction or too small by construction
Primordial relics	PBH and stochastic-background sectors are secondary, not dominant, in the minimal branch	Microlensing, PTA, interferometer, CMB constraints	A primordial-relic spectrum requiring dynamics foreign to the inherited background story
Black-hole information	Operationally unitary boundary-accessible ledger; no need for firewalls or ad hoc remnants in the minimal account	Theory-side entropy tests, plus eventual analog and astrophysical clues	Demonstration that the GDOS/OGGEP account cannot reproduce or reinterpret the modern Page-curve results quantitatively

This ledger is a statement of exposure, not a sales pitch. In a young unification program the strongest predictive claims — low primordial tensors, sharply constrained dark energy, a narrow dark-matter window, a geometrically compressed late inventory — are precisely the ones most likely to kill the theory early. If future data demand their opposites, the theory should not be praised for adaptability; it should be judged to have failed in its defining claims. A theory that survives such dangerous predictions gains something more valuable than flexibility: it gains character.

95.1 The decade timeline: what 2030 and 2035 decide

Kill criteria without dates are rhetoric. Table 51 attaches the instrument and the horizon to each, so the reader can mark a calendar rather than take a promise — and the cohort-wide canonical values quoted here are maintained in the single falsifier ledger of the GG-Theory program, so that no companion paper drifts from them; Figure 36 lays the same decade out as one forward time axis — the eight decisive kill-tests, color-coded by sector, from today's bounds to the mid-2030s verdicts. (The upstream companion papers carry the same discipline for the particle-side tests: the 3.85 TeV lock cluster at the HL-LHC in GG-P1-KK, the full falsifier pull-table in GG-A11-Origin.) One near-term row deserves emphasis: the inherited neutrino-mass sum $\Sigma m_\nu \approx 59$ meV sits at the normal-hierarchy floor, and a robust DESI+CMB upper bound that excludes 59 meV at 95% under this framework's own $w = -1$ prior — one that pushes below the normal-hierarchy floor itself, not one a millielectronvolt beneath the prediction — would falsify the inherited block; it may be the

first of these tests to be decided.

Table 51: The falsification decade for the cosmological branch: each kill criterion with its instrument class and decision horizon. Every row is two-sided — the framework loses if the measurement lands off the stated target, and banks the corresponding claim if it lands on.

Target	GG commitment	Instrument class	Decision horizon
Tensor-to-scalar ratio	$r = 0.0038$ ($\sigma(r)$ delensing-dependent)	Simons Observatory, LiteBIRD	late 2020s–mid 2030s
Kaxion line	1.59–2.13 GHz, DFSZ-line coupling	haloscopes scanning past 2 GHz	late 2020s–2030s
Dark-energy dynamics	$w = -1$ exactly; no evolution	DESI, Euclid, Roman	late 2020s
Siren Hubble constant	$H_0^{\rm GW} \approx 67$	GW standard sirens (LIGO/Virgo/KAGRA, ET)	∼2030 and beyond
Growth / S_8 split	the computed $S_8 = 0.825$ against shear at survey precision	Rubin LSST, Euclid	∼2030
Neutrino mass sum	$\Sigma m_\nu \approx 59$ meV (inherited block)	DESI + CMB combinations	near-term (may fire first)
Recombination switch	the atomic chain's 1091.5 against 1089.92 ± 0.25 (the switch e^7 annotated, not scored)	existing Planck + audit	now (analysis-side)
Lithium-7	SBBN/π, D/H and Y_p untouched	precision abundance inference	ongoing

By the mid-2030s, every load-bearing cosmological commitment of this paper will have been tested by an instrument that exists or is funded. The framework does not ask the reader's patience beyond one decade.

95.2 Hostile audit: what if every test fails?

It is worth walking the worst case stratum by stratum, because the right response differs by stratum.

Inventory. If the measured shares settle stably away from $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1-1/\pi$, $\Omega_b = 1/(2\pi^2)$, then Conjecture VIII is false and the closed-packet derivation is wrong. That is a single-module correction: the late-inventory interpretation is retired, but the inflation and dark-matter modules do not automatically fall with it.

Inflation. A primordial tensor amplitude inconsistent with the derived $r = 0.0038$ — a detection well above it, or a null pushing r below $\sim 10^{-3}$ — kills the derived launch outright. Again a single-stratum correction, though a defining one.

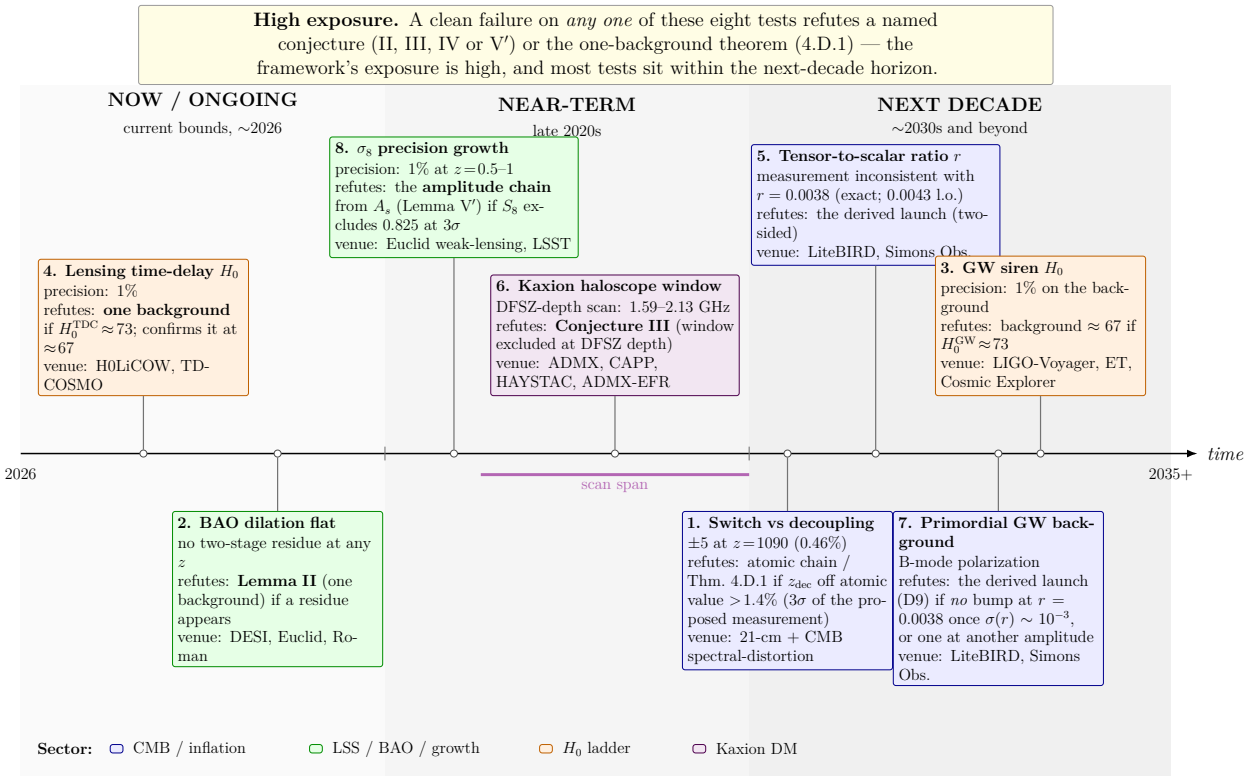

Figure 36: Falsification / kill-test timeline for GG-A10-Cosmos. The eight decisive kill-tests are placed on a forward time axis running from current bounds (∼2026) to the next-decade horizon (∼2035+), grouped into broad data-window eras (now/ongoing, near-term, next decade) rather than precise launch dates. Each card states the test, its precision/threshold, the named conjecture or theorem it would *refute*, and the experimental venue. Cards are colored by sector: CMB/inflation (blue), large-scale-structure / BAO / growth (green), H_0 distance ladder (orange), and Kaxion dark matter (violet). The Kaxion haloscope scan spans an ongoing-through-∼2030 window (violet bar). A clean failure on any single test refutes a named conjecture (II, III, IV or V′) or the one-background theorem (4.D.1), so the framework's falsification exposure is high and front-loaded into the coming decade.

Kaxion. Exhaustive nulls across the 1.59–2.13 GHz window at DFSZ depth retire the calibrated placement and its ladder; the dark sector would then have to be sourced elsewhere, damaging the cross-linkage but not the constitutional layer.

One background. A ladder-free late probe converging on the ladder's value rather than on 66.8 would mean the one-background reading is wrong. That weakens the late-time program but leaves the early-universe modules intact.

Black holes. If the GDOS/OGGEP account cannot be made quantitatively competitive with island computations, the black-hole module remains suggestive rather than decisive — a conceptual loss, not a falsification of the cosmology.

Each of these alone is a correction the framework can absorb. The genuine danger is *cluster-level* failure: if the inventory, the inflationary tensor target $r = 0.0038$, and the one-background reading were to fail *together*, that would not be three separate corrections but a refutation of the shared-module thesis itself, and the framework should be retired as a unification claim. That is the

correct standard. The framework should lose if the universe disagrees.

96 The bottom line

The verdict is tiered, and stated plainly.

Strongly supported. The compression itself — the reduction of thirteen largely independent mysteries to a short inherited ledger that preserves the ΛCDM baseline; the lawful finite beginning at E_{GUT}; and the sharp tensor target $r = 0.0038$ — $12/N^{*2} = 0.0043$ at leading order — which the plateau forces once its class and the reheating-matched e-fold count are fixed.

Partially or conditionally supported. The π-locked inventory (Derived-conditional under Conjecture VIII); the tilt and scale (Derived-conditional on the revised Lemma IV) with the amplitude A_s the one declared input; the one Hubble background against its two comparators (structural; the rate's uniqueness open under Conjecture 68.3) and the computed S_8 that shares its tension; the Kaxion window (one scale calibrated in GG-C1-Kaxion, awaiting the haloscope); and the operational black-hole account (conceptual, awaiting quantitative match).

Not yet tested. The forward predictions on which the theory most stakes its identity — the tensor ratio $r = 0.0038$, the 1.59–2.13 GHz Kaxion line, and a strictly $w = -1$ dark energy — together with the global growth-tension fit that a successor pipeline must deliver.

The forward program follows directly: the near-term priority is to make the precision sector computationally real and to attack the sharpest hinges first, because those are where the minimal branch can be killed most cleanly; the medium and long-term tasks scale the elegant modules to realistic astrophysics and tighten the particle–cosmos closure. The ordering, which reflects where credibility is gained or lost fastest, is collected in Table 52.

Table 52: A practical forward program suggested by GG-A10-Cosmos. The columns give the horizon, the tasks, and why that horizon matters most. The ordering reflects where the framework is most exposed rather than where progress would be easiest, which is the opposite of how a research plan is usually written and is intentional.

Horizon	Primary tasks	Why this horizon matters most
Near term	Full precision cosmology implementation; BBN network test; dark-energy and tensor-sector sharpening; Kaxion search targeting	These are the places where the minimal branch can be killed most cleanly
Medium term	Nonlinear structure and early-galaxy simulations; local-baryonic environmental tests; black-hole entropy matching in tractable models	These decide whether the framework scales from elegant modules to realistic astrophysical predictions
Long term	Tight particle-cosmos closure; extension to chemistry/life/mind floors only after first-floor stability is secured	This is where the larger GG-6 ambition becomes scientifically credible rather than merely suggestive

GG-A10-Cosmos has not eliminated the need for hard future work; it has clarified exactly where that work must occur, reduced the number of independent stories that can be told at once, and made the whole first-floor particle–cosmos enterprise succeed or fail as a single correlated whole. That is what a unification proposal should do, and it is the disciplined foundation on which the concluding Part builds its wider view.

97 Part XIII in plain English: what is proved, what is conditional, and what is open

The honest boundary section is the most important thing in this Part, and everything else should be read against it. The framework's cosmological claims fall into three groups: results forced by the geometry, results forced conditional on a conjecture that is named, and questions the paper reframes without answering. The three are not interchangeable, and the Part is built to keep them apart.

The cost of being right, as the Part puts it, is that there is nothing to adjust. A framework with zero continuous parameters cannot absorb a disagreement: one wrong number retires it. That is stated as a strength; it is a liability that becomes a strength only if the numbers hold. The kill criteria and the test ledger are the operational form of that liability, and they are written to be decidable by measurement rather than argument: a kill criterion that requires an interpretive judgment is not one.

The ten structural innovations are the Part's claim to novelty and they are the softest material in it. An innovation is a way of organizing an account, not a result, and none of the ten would survive the failure of the theorems they organize. They are listed because a reader is entitled to know what the paper thinks is new about it, and listed after the honest boundary rather than before it for the same reason.

Part XIV
Conclusions: the summit view

The composition of the universe is geometry, not coincidence — one origin for inflation, dark energy, and dark matter

98 The summit thesis

For a century, cosmology has been a science of measurement. We weigh the universe, and we accept what the scale reads. The concordance model ΛCDM is the triumph of that method — six numbers, fit to the sky, that reproduce the cosmos from the first light to the most distant galaxy. It tells us, with magnificent precision, *how* the universe expands. It has never told us *why the universe contains what it does.*

GG-A10-Cosmos supplies the missing half of the sentence.

> *The six numbers of the concordance model are not six fortunate accidents to be measured and accepted. They are one geometric figure — read off a single six-dimensional bulk, the very geometry that fixes the electron one story below. The composition of the universe is not a coincidence. It is geometry.*

This is the summit view, and from it the whole landscape rearranges. The dark sector ceases to be a pair of unexplained substances that happen to dominate the cosmos; it becomes what the inherited Standard-Model ledger *cannot see* of a closed, parameter-free inventory. The two warring values of the Hubble constant cease to be a crisis; they become one derived rate and one calibration systematic the geometry predicts. The beginning ceases to be a singular mystery bolted on by hand; it becomes the high-energy edge of a single radial descent that ends, today, in the gentle vacuum we inhabit. One geometry, from the first instant to this one.

99 The composition of the universe is a number

Here is the claim in its barest form. Ask what fraction of the universe is ordinary matter, what fraction is dark matter, what fraction is the dark energy driving it apart. ΛCDM answers by measurement: roughly five percent, twenty-seven percent, sixty-eight percent — three numbers handed down from the data with no account of where they come from. GG-Theory answers by geometry:

$$\Omega_m = \frac{1}{\pi} \approx 0.318, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi} \approx 0.682, \qquad \Omega_b = \frac{1}{2\pi^2} \approx 0.051.$$

The matter fraction is the reciprocal of π. The dark-energy fraction is what remains of unity once that share is taken. The baryon fraction — everything that has ever been a star, a planet, or

a person — is one part in $2\pi^2$. These are not fits. They are the closed inventory of a compact geometry, and they fall on the measured sky.

And now the deepest reading, the one that turns a numerical curiosity into a physical revelation. *Why is the universe ninety-five percent dark?* Because the inventory is closed and π-locked, and the Standard-Model gauge ledger inherited from GG-A9-SM can illuminate only the $1/2\pi^2$ of it that carries ordinary charge. The darkness of the cosmos is not the presence of two exotic strangers. It is the shape of what light was built to miss. Dark matter and dark energy are dark for the same reason a clean window is invisible: not because they are absent, but because the instrument sent to find them was made to look straight through them.

That is the cosmological face of the principle that runs through the whole cohort: the visible world is the boundary shadow of a fuller bulk.

100 One origin, from the first instant to today

The power of a geometric account is that it does not stop at the ledger. It runs forward in time, and it offers, for the first time, a common origin for the three great unexplained chapters of cosmic history.

The beginning — inflation as a geometric launch. The oldest puzzle in cosmology is why the early universe was so smooth and so flat. The standard answer, inflation, is a brilliant mechanism in search of a cause: a scalar field, introduced for the purpose, rolling down a potential chosen to fit. GG-Theory offers the cause. The universe begins at the ultraviolet edge of the throat, near E_{GUT}, and descends. The inflationary epoch is not a bolted-on prologue but the first stretch of that descent — a geometric launch on the derived plateau just below the GUT edge, smooth and graceful because the geometry, not a tuned potential, lays the road; the TeV lock $M_{\mathrm{lock}} \approx 3.85$ TeV is crossed lower down, where the dark sector is born. For the first time the inflaton is not a guest invented for the occasion. It is the radial coordinate of the bulk.

Today — dark energy as the residual vacuum. The universe is ending as it began, in acceleration. The standard account writes down a cosmological constant and measures its value; the failure to explain that value is the deepest embarrassment in physics. GG-Theory reads the late-time vacuum not as a constant to be explained away but as the residual $1 - 1/\pi$ of the same closed inventory, with an equation of state $w = -1$ that is exact because the geometric transition is complete. The acceleration of the cosmos is the geometry settling into its ground state. Dark energy is simply what a finished universe weighs.

In between — dark matter with a mass and an address. Between the launch and the rest lies the matter that built the galaxies, and here the geometric account is sharpest of all. GG-Theory predicts that dark matter is one specific particle — the Kaxion, a faint and extremely light particle that lives in the curled-up compact dimension of the bulk — with one specific mass:

$$m_{\mathrm{Kaxion}} \simeq 8.1\ \mu\mathrm{eV}, \qquad \nu_K \in [1.59,\ 2.13]\ \mathrm{GHz}.$$

The direction of inference is worth stating exactly, since it is easy to state backwards. The geometry does *not* fix the decay constant: the bulk scale f_{bulk} is a declared input, calibrated by the dark-matter abundance in the companion GG-C1-Kaxion. What the geometry supplies is the structure that spends it — the compact-angle mode, the cap-loading ladder that makes the decay constant one of a discrete list, and the relation that converts it into a frequency. One number is bought, and a mass, a window, a coupling and a line count are returned, none of which the abundance could have bought on its own. For the first time in the history of the subject, dark matter has a geometric origin, a number, and a laboratory window narrow enough to be opened — or shut — within the decade; the derivation is carried by GG-C1-Kaxion.

Three chapters — the smooth beginning, the matter that clumps, the vacuum that drives apart — have stood for decades as three separate mysteries, each with its own remedy invented by hand. GG-A10-Cosmos reads them as three movements of one geometry: the descent, the relic, and the residue. That is what it means to give cosmology an origin instead of a fit.

101 One universe, one story above the Standard Model

The unification is the point that must not be lost among the numbers. The same six-dimensional bulk that, in GG-A9-SM (Arisaka, 2026, A9-SM), fixes the masses and charges of the particles also fixes the contents of the cosmos: one ledger, read at two scales. The same principle answers both of the questions those two papers ask. Why is every electron in the universe identical? Because it is the readout of the one ground state its stratum allows. Why is the universe mostly dark? Because the inventory is closed and the inherited gauge ledger can illuminate only the charged $1/2\pi^2$ of it. GG-A10-Cosmos sits one story above GG-A9-SM in a single building — not in a separate annex of cosmology bolted to a separate physics of particles.

And the building has a direction. The arrow of cosmic time — the reason the universe runs from a hot, dense, symmetric beginning to a cold, sparse, structured present — is the open-system descent at the heart of the cohort's dynamics, the GRIP cascade (the dynamical face of the one constitutional arrow that AX-3 owns): the universe relaxing toward the only state it can rest in, the de Sitter vacuum and the frozen Kaxion relic that are its final, stable furniture. Cosmic history is not a string of accidents laid along a timeline. It is a single object falling, with dignity, toward its ground state. The universe has a plot, and the plot is geometric.

102 The fourth revolution, written across the sky

Stand back far enough and the lineage is plain. Euclid gave physics the ideal of deriving the world from a handful of self-evident axioms, with nothing left over to choose. Newton made the heavens and the earth obey one law. Maxwell, Einstein, and the founders of quantum theory unified the forces and the constituents. Each revolution removed a coincidence and put a necessity in its place.

The dream went quiet in cosmology, where the deepest numbers — the density of matter, the strength of the vacuum, the smoothness of the start — stayed things we measured rather than

things we understood. GG-A10-Cosmos carries the Euclidean ideal back into the sky: a cosmos with *zero free continuous parameters*, its inventory fixed by π and a short ledger of integers, its history fixed by the descent of one geometry. Twenty-three centuries after Euclid, the composition of the universe rejoins the list of things that did not have to be measured to be known. That is the fourth revolution — and the sky is where it becomes visible to the naked eye.

What changed, in the end, was not the data but the reading of it. For decades the cosmic numbers were treated as inputs to a closed four-dimensional theory; the recognition on which this paper rests — its central claim (§6, Figure 9) — is that no universe is closed, that the four-dimensional cosmos we measure is the open-system projection of a six-dimensional one, and that this projection is what observation is. Once that is seen, the matter fraction $1/\pi$, the dark-energy residue $1 - 1/\pi$, and the corridor-switch redshift e^7 are not free numbers but registered observables of a determined map — and the deepest result is not any single closed form but the shift in what a cosmological parameter is: from an input one is obliged to measure to an output the geometry is obliged to return.

103 A universe made answerable

A vision this large deserves no quarter unless it can be killed, and the strength of a geometric cosmos is exactly that it hands the sky a set of sharp, unforgiving questions. The corridor switch sits at $z_{\rm sw} = e^7 \approx 1097$, and the recombination surface, derived from the framework's own constants, lands 0.4% below it on the measured value. The primordial tensor sector should ring at exactly $r = 0.0038$ — loud enough for the coming instruments, far below the present bound. The Kaxion should ring in a narrow band near two gigahertz. The two Hubble constants should prove to be one rate and one calibration systematic, not two errors. Each of these is a place where the universe is free to answer no. And within the coming decade — 21-centimeter cosmology, the next generation of microwave-background experiments, the haloscopes now scanning past two gigahertz, the galaxy surveys charting the expansion — it will have its chance.

This is the only confidence a geometric theory is entitled to: not the confidence that it cannot be wrong, but the confidence to tell the world exactly how it could be. A framework that fixes the cosmos from π and a handful of integers has nowhere to hide, and asks for nowhere. Nor does it claim more than it computes: the geometry fixes the cosmic *inventory* and the *chronology*, while the precision microwave-background C_ℓ spectrum and the nonlinear growth of structure remain the work of the Einstein–Boltzmann and N-body machinery it frames but does not replace. We have made the universe answerable. The full slate of decisive tests is laid out in Figure 36.

One word on the seam, since a reader may fairly note that a cosmology paper has reached through chemistry and life toward minds while candidly *not* computing a full C_ℓ likelihood. The choice is deliberate. The arc from cosmos to life to mind is carried here as a *pointer* forward to GG-A11-Origin (Arisaka, 2026, A11-Origin), where it receives its theorem-disciplined treatment; it is not a claim this paper proves. The falsifiable core of GG-A10-Cosmos — the tensor target $r = 0.0038$, the Kaxion line, $w = -1$, the siren $H_0 \approx 67$, and the π-locked inventory — stands entirely on its own: every kill test of Part XIII is stated and decidable without a single sentence of the arc. A reader who wishes to judge this paper as a bare cosmology may stop at the end of Part XIII and lose no

testable content. A reader who wishes to see why the program thinks the cosmos and the observer are one descent is invited onward. We keep both readers, and the seam between them is marked here rather than left to inference.

104 One staircase, from the Big Bang to living molecules

Everything in this paper, and everything in the cohort beneath it, runs on one short chain of cause and effect. It is worth saying once in plain words, with no symbols in the way.

The universe never shows us its full self. What we see is the surface of a deeper geometry — the shadow that a closed, six-dimensional space casts on the four-dimensional boundary we live on. That projection is the third axiom of the GG-Constitution, OGGEP. The boundary is not frozen: it evolves by a definite law, GDOS. And a definite process acts on it, GRIP. Faced with the open choices a boundary always allows, GRIP tries them, charges each one an energy cost, breaks the tie, and relaxes into the single configuration that costs the least. The thing it settles into — the value we then measure — is the Ground Geometric Attractor, GGA. Every number in the codebook of this paper is one such attractor. In a single line:

$$\mathsf{OGGEP} \longrightarrow \mathsf{GDOS} \longrightarrow \mathsf{GRIP} \longrightarrow \mathrm{GGA}$$

$$\Longrightarrow \ \text{atoms} \to \text{galaxies} \to \text{stars} \to \text{heavy elements} \to \text{life's molecules} \to \text{us}$$

Read it from left to right and it is the whole of cosmic history.

It begins at the hot edge, near E_{GUT} — the geometric Big Bang. From that first instant the entire staircase is already set. Not the events one by one, but the rules and the numbers that make every later step come out as it does: nothing is added later by hand. Everything that follows is geometry spending the budget it was handed at the start.

Run the chain on the particle stratum and out come the stable electron, the stable proton, the stable atom — each a GGA, stable not by luck but because there is nothing it can decay into. Stable atoms are the first precondition for all the rest: without them there is no chemistry, no planets, and no chemists. Run the same chain on the cosmological stratum and out comes the inventory this paper derived — the dark matter that is the unseen scaffolding of everything large, and the dark energy that paces the expansion so that there is time to assemble galaxies before the long cold sets in.

Now gravity does the building. It gathers the π-locked matter into halos and galaxies; inside them gas falls together and stars ignite. Stars are the universe's foundries. While they live they fuse hydrogen into carbon, oxygen, and nitrogen; when they die they burst as supernovae and scatter those heavy elements across space. From that enriched dust, ordinary chemistry assembles the organic molecules — the carbon chains and rings that life is written in. And on at least one cold rock orbiting one ordinary star, those molecules found their way into systems that are now reading this sentence.

Each rung is forced by the rung beneath it, and the whole flight was set in outline at E_{GUT}. That single descending staircase — bulk to boundary, to atoms, to galaxies, to stars, to molecules, to minds — is the picture of the world that GG-Theory offers, and this paper has laid its first and

broadest steps.

105 The trio: from the electron, through the cosmos, to the question of us

The staircase just walked is not this paper's private image; it is the architecture of three papers written as one argument, and the reader — especially the astronomer or cosmologist who arrived here for the cosmology — deserves the map.

GG-A9-SM (Arisaka, 2026, A9-SM) is the first movement: the particle floor. It closes the twenty-eight-parameter sheet of the Standard Model and then asks the question textbooks never ask — *why is every electron the same?* — and answers it with a uniqueness theorem: every electron is the readout of the one ground attractor of its stratum. The present paper, GG-A10-Cosmos, is the second movement: the same engine, one story up, where the attractor being read is the vacuum of the universe and the readout is the π-locked inventory — *why does the universe contain what it does?* And GG-A11-Origin (Arisaka, 2026, A11-Origin) is the third movement and the capstone: the same cascade followed past the galaxies, through chemistry, to the open systems that metabolize, remember, and eventually look up — *why is there anyone here to ask?* Three questions, one geometric answer, applied at three strata: that is the trio, and it is the program's deliberate spine.

GG-A11-Origin names what the third movement requires, and it is worth previewing here in its own vocabulary, because the cosmological reader will recognize every ingredient from this paper. The *recursive uniqueness theorem*: the same GRIP selection that fixed the electron (GG-A9-SM) and the inventory (this paper) runs layer by layer — eighteen layers, from the gauge structure to the act of articulation — each layer's Ground Geometric Attractor becoming the substrate of the next. The *geometric inevitability of Cosmos, Life, and Mind*: Cosmos is the present paper; Life is the SeOS crossing, where the same open-system ledger that ran the cosmic descent begins to run metabolism; Mind is the closed loop of world and world-model. And the *self-referential closure*, the capstone's central theorem: the geometry that forced the cascade contains, among its forced readouts, the observers who write the geometry down. GG-A11-Origin also places the whole program against the lineage of axiomatic revolutions — Euclid, Newton, the field-on-manifold revolution of Maxwell–Einstein–quantum mechanics, and the openness revolution of GG-A5-GDOS — and states, with the verdict left to history, what kind of fifth step the recursive cascade would be. The "fourth revolution written across the sky" of this Part is that same claim, read at the cosmological floor.

A word of orientation for the reader who wants more. Everything this paper inherited has a home: the constitution in GG-A1-Constitution, the six-dimensional bulk and its integer ledger in GG-A2-GG6–GG-A4-Euclid, the open-system boundary law in GG-A5-GDOS, the runtime in GG-A6-GRIP–GG-A7-Qi4, and the particle floor in GG-A9-SM — the upstream staircase is climbable one paper at a time, and Part II of this paper is its one-Part summary. But if the reader takes only one further step, the author's counsel is to take it *downstream*: GG-A11-Origin is where the question this Part has been circling — not what the universe is, but why there is someone in it to ask — is given its full, theorem-disciplined treatment. This paper derived the stage and the furniture; the capstone is about the audience. And the summit is not the end of this volume: the reader who wants

the craft behind the claims will find in Part XV a self-contained graduate course in the standard cosmology this paper set out to derive, and the reader who wants the fight will find in Part XVI the ninety-seven hardest objections we could assemble, answered without retreat into vagueness.

106 Fine-tuning, answered without luck and without a multiverse

There is one question a cosmology this ambitious cannot duck, and one a paper confident enough to have climbed this far need not flinch from. Why do the constants of nature sit where life needs them?

For decades the unsettling observation has been that the universe looks improbably well-suited to life. Make the electron a little heavier, the nuclear force a little weaker, or the energy of empty space a little larger, and there are no stable atoms, no stars that burn for billions of years, no carbon — and so no us. The constants look tuned, as though some careful hand had set each dial to exactly the value life requires. Physics has long offered two ways out: luck (we happen to live in the one universe that works) or a multiverse (every setting exists somewhere, and naturally we find ourselves in a habitable one).

GG-Theory takes a third road, and once the staircase of §104 is in view it is the natural one. The life-permitting constants are neither lucky nor selected from a bin of alternatives. They are π and a short ledger of integers, fixed by the geometry at E_{GUT}, with no dial to turn. There was never a tuning to explain, because there was never a knob. In exactly this sense the universe is extraordinarily well designed — not by a designer, but by geometry. It is built to the one specification its own shape allows, and that specification happens to be the one in which electrons hold atoms together, stars burn long and steady, and carbon is everywhere. A universe whose contents are forced is a universe whose fitness for life is forced along with them.

One consequence deserves saying, because the fine-tuning debate often loses it. The projection had been running for nearly fourteen billion years before the first observer existed. It needs no observer and waited for none; we did not switch it on by looking. Observers are not the reason the universe has its numbers. They are a late step on the staircase those numbers built.

That is as far as physics takes the argument, and it is already far. The oldest question — why is there anyone here at all? — has been asked by every civilization: the Upanishads answered that the cosmic self sought to know itself; Genesis, that a creator chose; the Enlightenment set it aside so that physics could proceed; Darwin explained how life diversifies but not why anything exists to evolve. GG-A11-Origin opens on exactly this history and proposes that the question now has a scientific answer. The full case — that life and mind are not fortunate accidents but forced rungs of the same cascade, closed by the theorem that the geometry contains its own articulators — belongs to that capstone; here we claim only its first and firmest link, and we claim it plainly. The universe is not tuned for us. It had no tuning to do, and we are what a universe with nothing to tune eventually contains.

107 A personal word from the cosmological floor

The author may be permitted one paragraph in the first person, matching the personal accounts that close GG-A9-SM and GG-A11-Origin. I spent fifty years as an experimental particle physicist — Kamiokande, kaon and B-meson CP violation, CMS at the LHC, Auger, XENON — watching the Standard Model refuse to break, and watching the deepest discoveries of the era arrive instead from the sky: neutrino oscillations, the precision microwave background, the accelerating universe, gravitational waves. Four revolutions, four Nobel prizes, and not one of them predicted by the unification programs of my generation. It is therefore no accident that the paper you are finishing stakes its life on the sky. The instruments that will judge it — the CMB polarimeters, the haloscopes, the galaxy surveys, the gravitational-wave detectors — are the direct descendants of the ones that taught my generation cosmology; and the numbers it stakes — a tensor ratio, a dark-matter frequency, an equation of state, a siren Hubble constant — are the kind an experimentalist can love: sharp enough to be wrong. If the universe says no, it will have said no to a theory that asked properly, and the record in these pages will show exactly which link broke. If it says yes, then the composition of the universe was never a set of accidents to be measured — it was a theorem waiting, for one century of magnificent measurement, to be read. Either answer honors the sky that taught us; and either way, the next generation of cosmologists — the readers this paper was written for — will be the ones who hear it first.

108 Closing remarks: why are we here?

A book that has spent seven hundred pages deriving the contents of the universe owes its reader one last question — the one every reader brought with them. Why are we here? Not the version §106 answered, which was about dials and had a physical answer, but the version a person means when they ask it.

The question has two readings, and physics can answer only one of them. The first asks for a purpose handed down from outside the world — a reason the universe was made, with us in mind. Physics has never had access to that question and does not acquire access here; a geometry that admits no alternative is not a geometry with an intention. The second asks for a *place*: where, in the account of everything that happened, do we belong? That question is answerable, and the answer is larger than the author expected when the work began.

The answer this paper can give is continuity. The universe kept no separate ledger for life: the same descent that fixed the weight of empty space fixed the strength that holds an atom together, the patience with which a star burns, and the peculiar willingness of carbon to make long molecules — and it ran on, in the ordinary way, into systems that hold a model of their surroundings inside a part of themselves. We are not visitors in a cosmos that happens to permit us. We are a late movement of the same unfolding, made of the same forced ledger, and our capacity to ask about the whole is one of the things the whole turned out to contain. We are the part of the geometry that has begun to look back. That is the sense in which the answer is warmer than the question: we are here because everything is here, by the same necessity, on one staircase, with no separate act

required for the last step. Whether the last steps are forced as firmly as the first is the capstone's question, and GG-A11-Origin (Arisaka, 2026, A11-Origin) answers it on a conjecture it names rather than hides. This paper claims the firmest link only: the stage, the furniture, and the fact that neither had a dial to turn.

Three things are easily read as more than they are, and are better said plainly. Nothing here says the cosmos was designed. Nothing here says its history was fixed in every detail: the layered structure is forced, the path within the layers is not, and that is exactly where contingency, evolution, and freedom live. And nothing here says the argument is finished — Part XIII prints the criteria by which the sky may end it.

What we would ask of the next generation is a habit rather than a belief. The century of measurement this paper opened by describing was magnificent — the expansion, the microwave background, the neutrino masses, the accelerating vacuum — and a measured number is still only half of an explanation. Questions set aside as philosophical have a way of becoming technical when someone is stubborn enough to ask them with an instrument in hand. The age of the earth, the origin of the elements, the origin of species: each was once metaphysics, and each became physics when someone refused to be satisfied and had the patience to be checked. Whether the composition of the universe joins that list is now, for the first time, a question with a date attached to it: the polarimeters, the haloscopes, the surveys, and the interferometers of the coming decade will settle it while today's students are still young.

So this is written for those students, in the hope that they attack it. Take the numbers in these pages and try to break them; the record is printed so that a broken link can be found and named. If the framework falls, it will have failed in a way that teaches something, and the sky will have been asked a proper question. If it stands, the reward is not that anyone was right. The reward is that the universe turned out to be intelligible one layer deeper than we thought — and that the part of it doing the understanding is not a stranger to the rest, but the same geometry, at a later hour, learning to read its own inventory. There is no better work to give a life to.

However far you have come: thank you for climbing. The instruments are being built. Go and listen.

Acknowledgments

This work was developed through extensive iterative collaboration with three large language models: ChatGPT (OpenAI) for exploratory drafting, narrative development, and broad-scope ideation; Perplexity for literature synthesis, citation verification, and recent experimental cross-referencing; and Claude (Anthropic) for structural organization, mathematical consistency checking, LaTeX typesetting, and final manuscript preparation. The sequential refinement of the GG-Theory framework (GG-Constitution / GG-6 / BI-6 / GDOS / Gi-4 Pillars / Qi-4 Pillars) across hundreds of conversation sessions with these three systems was essential to the coherence of the present paper.

The author thanks Aaron Blaisdell and Hirotaka Iwasaki for stimulating discussions on the program's interdisciplinary scope and for critical feedback on the accessibility of the presentation. The author also thanks hundreds of collaborators in the Elegant Mind Club over the past decade, including Javier Carmona, Elizabeth Milles, Mingda He, Syed Bashir Hydari, Brian Ta, Chris Dao, Umaima Afifa, Isabella Bestanoby, Caominh Le, and Natsuko Yamaguchi.

All scientific concepts (especially the foundation of the GG-Constitution, GGEP/GG-6/BI-6, and OGGEP/GDOS/Gi-4/Qi-4 Pillars), theoretical judgments, physical interpretations, and any remaining errors are the sole responsibility of the author.

Part XV
Technical Appendices

A graduate-level course in cosmology read through GG-Theory — full from-scratch derivations of the Friedmann background, thermal history, inflation, nucleosynthesis, the growth of structure, the cosmic microwave background, the π-locked inventory, Kaxion dark matter, and black-hole entropy — preceded by the reference dictionaries and followed by the source ledger, benchmark audit, and registries

---

The body of this paper has emphasized the geometric theorem chain and kept the heaviest arithmetic out of the reader's way. The present appendices supply that detailed technical content in full, and taken together they are intended to be read as *a graduate-level course in cosmology seen through GG-Theory*, addressed to an incoming graduate student in physics, astronomy, or cosmology who will carry the program forward. We organize them into three groups. The *reference* appendices (A–D) tabulate the abbreviations, the key symbols, the compact codebook, and the cross-cohort naming convention, and serve as the reader's quick lookup. The *constructive* appendices (E–N) are the course proper: a self-contained sequence that takes up each sector of modern cosmology in turn — the Friedmann background; expansion, scaling laws, and the thermal history; inflation; nucleosynthesis and the lithium audit; leptogenesis; the growth of structure; the microwave background; the π-locked inventory; Kaxion dark matter; and black-hole entropy — and treats each one in *three steps*. First comes a full from-scratch derivation of the standard result, in the form a graduate student would meet it in Dodelson or Weinberg, carried through with no skipped manipulations and closed by a worked numerical example. Then comes a frank statement of what that standard derivation *assumes but cannot explain* — the fitted parameters it must take from the sky and the structural questions it leaves open: what the inflaton is, why the misalignment angle takes one value, why the vacuum weighs what it does. Finally comes the geometric derivation, in which the same quantity descends as a boundary readout or a theorem of GG-Theory, fixed by the upstream theorem deck rather than fitted — and it is in these third steps that the paper's central derived results enter: the derived plateau and the computed tilt, the tensor target $r = 0.0038$, the reheating temperature, the pinned misalignment angle, the ground-state reading of dark energy. The *ledger and registry* appendices (O–T) that close the Part record the internal source ledger, the cross-paper inheritance table, the consolidated benchmark audit, the named-theorem-and-conjecture deck, the lemma-promotion roadmap, and the falsification catalog.

Throughout these appendices we adhere to two stylistic commitments. First, every derivation is presented step by step with no skipped intermediate manipulations: a graduate student in physics should be able to reproduce each calculation with paper and pencil. Second, every formal step is accompanied by a parenthetical *physics intuition* note that translates the algebra into a mental picture. This combination is designed to make the appendix material accessible to the working physicist while preserving the mathematical rigor of the main text. Each course appendix opens

with a boxed *What this appendix teaches* abstract and closes with a boxed *take-away* — the standard result, what it assumed, and what the geometry replaced the assumption with — so that the Part can be used as a one-semester reading course, one appendix per week, each self-contained. And every course appendix closes with *problems and complete worked solutions* — placed before the take-away box, numbered by chapter and subsection, each solution worked in explicit Steps and ending with its physics intuition, the sets cross-referencing one another into a single course network. They carry full solutions because, in the cohort's phrase, they are meant to be done, not admired: the reader who works them owns the derivations, not merely the conclusions.

Prerequisites and recommended textbooks. The course assumes advanced undergraduate mechanics, electrodynamics, and statistical physics, plus a first exposure to general relativity (the FRW metric is derived from scratch in Appendix E, so a student meeting cosmology for the first time can begin there). It is meant to be read *alongside* a canonical cosmology textbook, not in place of one: the recommended way to use each appendix is to read the corresponding textbook chapter first, return to Step 1 here for a compact but complete recap, and only then proceed to Steps 2 and 3, where the standard result is re-examined and re-derived as a theorem of the geometry. For self-study we particularly recommend the following, each of which is cited again in the appendix to which it is best matched:

- **Primary references.** Dodelson (2003), *Modern Cosmology* (and its second edition with F. Schmidt), and Baumann (2022), *Cosmology* — the two standard modern graduate texts, covering the background, the thermal history, inflation, perturbations, and the CMB in the notation this course follows most closely.
- **Gentler first pass.** Ryden (2017), *Introduction to Cosmology*, and Liddle (2015), *An Introduction to Modern Cosmology* — accessible and physically transparent, ideal for a first encounter with the expanding universe and the inventory.
- **In-depth study.** Weinberg (2008), *Cosmology*; Mukhanov (2005), *Physical Foundations of Cosmology*; Peacock (1999), *Cosmological Physics*; and Liddle & Lyth (2000), *Cosmological Inflation and Large-Scale Structure* — for the relativistic background, the quantum theory of perturbations, and the inflationary sector at full depth.
- **Classic and specialist companions.** Kolb & Turner (1990), *The Early Universe* — the classic source for the thermal history, freeze-out, nucleosynthesis, and relic abundances used in Appendices F, H, and M; Peebles (1993), *Principles of Physical Cosmology* — the deep classic on structure formation, for Appendix J; and Durrer (2020), *The Cosmic Microwave Background* — the dedicated CMB monograph, for Appendix K.

Throughout, the symbols GG-A1-Constitution–GG-A11-Origin, GG-P1-KK–GG-P5-CPT, and GG-C1-Kaxion–GG-C3-BH refer to the twenty-two papers of the GG-Theory cohort, to which the geometric steps repeatedly point.

A Summary of Abbreviations

This appendix collects every abbreviation used in GG-A10-Cosmos. All program-defined names are rendered in *sans-serif type* (GG-Theory, GG-Constitution, GG-6, GDOS, BI-6, OGN-5, Gi-4, Qi-4, GRIP, etc.) so the reader can see at a glance which terms belong to the GG-Theory program proper. Standard physics abbreviations imported from the wider literature (ΛCDM, CMB, BBN, BAO, KK, RS, etc.) are rendered in plain uppercase Roman, following the convention of the physics literature. Mathematical objects are in roman or italic in the usual way. The third column gives the *primary link*: the paper in the Arisaka-GG series (and the section, where useful) in which each program-defined name is first introduced in detail.

Table 53: Summary of abbreviations used in the present paper. Program-defined names are rendered in sans serif; standard physics abbreviations in plain uppercase Roman; mathematical objects in roman/italic in the usual way. The third column records the primary link: the paper in the Arisaka-GG series, or the section of this one, where each name is introduced in detail. An abbreviation with no primary link beside it is standard physics usage and is listed only so the table can be read without a second reference open.

Symbol	Meaning	Primary link
(1) Framework name and core program objects		
GG-Theory	Grand Geometric Theory (also: Geometry–Gauge Equivalence Theory)	GG-A1-Constitution, Preface
GG-Constitution	Constitutional charter: 4 axioms (AX-1–AX-4) + 7 postulates (POS-1–POS-7)	GG-A1-Constitution, §2–§3
GG-6	Geometry–Gauge locking in six dimensions: the 6D bulk	GG-A2-GG6, §8
GDOS	Geometric Dynamics of Open Systems: the boundary law from OGGEP	GG-A5-GDOS
BI-6	Green–Schwarz-modified Bianchi identity at 6D, $dH_3 = X_4$	GG-A3-BI6
OGN-5	Occam-Gauge-Normalization 5-integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$	GG-A3-BI6, §7; GG-A2-GG6
Gi-4	Geometric-Invariance four-pillar structure (TGG, CCI, GFF, GRIP)	GG-A6-GRIP
Qi-4	Quantum-Information four-pillar structure (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL)	GG-A7-Qi4
(2) Four axioms (AX-1–AX-4)		
AX	Axiom (axiomatic primitive; non-negotiable)	GG-A1-Constitution, §1.5
LC	AX-1 Locality and Causality	GG-A1-Constitution, §2.2
GGEP	AX-2 Geometry–Gauge Equivalence Principle	GG-A1-Constitution, §2.3
OGGEP	AX-3 Operational/Open Geometry–Gauge Equivalence Principle	GG-A1-Constitution, §2.4
MSMF	AX-4 Maximum Symmetry, Minimum Freedom (no hidden continuous knobs)	GG-A1-Constitution, §2.5
(3) Seven postulates (POS-1–POS-7)		

continued on next page

Table 53 continued

Symbol	Meaning	Primary link
POS	Postulate (admissibility filter; testable in principle)	GG-A1-Constitution, §1.5
GI	POS-1 Gauge invariance across scales	GG-A1-Constitution, §3.3
UEP	POS-2 Unified Extremal Principle	GG-A1-Constitution, §3.4
AF	POS-3 Anomaly Freedom and cancellability by inflow	GG-A1-Constitution, §3.5
CP	POS-4 Closure and Positivity (CPTP-admissible open dynamics)	GG-A1-Constitution, §3.6
RG	POS-5 Renormalization-Group covariance (scale covariance)	GG-A1-Constitution, §3.7
OGN	POS-6 Occam Gauge Normalization (canonical integer/rational normalization)	GG-A1-Constitution, §3.8
MF	POS-7 Measurement and Falsifiability (operational contract + kill criteria)	GG-A1-Constitution, §3.9
(4) Gi-4 pillars		
TGG	Total Geometric–Gauge object (Pillar I of Gi-4)	GG-A6-GRIP
CCI	Covariant Conservation Identity (Pillar II of Gi-4)	GG-A6-GRIP
GFF	Geometric Free-Energy Functional (Pillar III of Gi-4)	GG-A6-GRIP
GRIP	Gauge-Runtime Invariance Pipeline (Pillar IV of Gi-4)	GG-A6-GRIP
GEP	Geometry–Gauge Energy Principle (energetic block of GFF)	GG-A6-GRIP
CEP	Curvature–Entropy Principle (entropic block of GFF)	GG-A6-GRIP
DGI	Dynamical Gauge Invariance (initial stage of GRIP)	GG-A6-GRIP
GSSB	Geometric Spontaneous Symmetry Breaking (middle stage of GRIP)	GG-A6-GRIP
GA	Geometric Attractor (final stage of GRIP)	GG-A6-GRIP
GGA	Ground Geometric Attractor: the unique global minimum among a stratum's GA's — identity and stability as ground-state-ness	GG-A6-GRIP; GG-A11-Origin; Part X
(5) Qi-4 pillars and the Arisaka equation		
Qi-QUA	Quantum information for Quadruple Gauge Locking — Pillar I of Qi-4	GG-A7-Qi4
Qi-EP	Quantum information for the Unified Extremal Principle — Pillar II	GG-A7-Qi4
Qi-CON	Quantum information for Conformal Invariance — Pillar III	GG-A7-Qi4
Qi-HAL	Quantum information for the Holographic Attractor Lattice — Pillar IV	GG-A7-Qi4
OH	Outer Hologram (boundary projection of bulk data; outward face of EIC)	GG-A7-Qi4
IH	Inner Hologram (inner self-modeling layer; inward face of EIC)	GG-A7-Qi4
EIC	Evolving Inner Cosmos (OH+IH composite)	GG-A7-Qi4

continued on next page

Table 53 continued

Symbol	Meaning	Primary link
Q_B	Arisaka constant $= (2/5)\,E_{\mathrm{Ry}}/K \approx 0.1555$ eV per qubit	GG-A7-Qi4
(6) Qi-HAL sub-acronyms		
HAL	Holographic Attractor Lattice	GG-A7-Qi4
BFFT	Boundary-Front Field Theory of Qi-HAL	GG-A7-Qi4
(7) GDOS-internal acronyms (program-defined)		
SIP	System Instance Packet (canonical GDOS input packet)	GG-A7-Qi4
SeOS	Selfish Open System (the core open system of GDOS)	GG-A5-GDOS
(8) Program-defined mathematical symbols (audit machinery)		
$\mathcal{R}$	Candidate rational expression in the MSMF/OGN audit	§68.2
$\mathcal{C}_\Sigma$	OGN-admissible candidate space on stratum Σ	§68.2
$\mathcal{S}_\Sigma$	Structural base set of atoms on stratum Σ	§68.2
$\mathrm{Cost}(\mathcal{R})$	MSMF cost function: total absolute exponent budget	§68.2
ϵ_X	Stratum-specific empirical-precision target for parameter X	§68.2
(9) Standard physics abbreviations (plain Roman, journal convention)		
Cosmology and the concordance model		
ΛCDM	Lambda Cold Dark Matter concordance model	textbook
CMB	Cosmic Microwave Background	textbook
BBN	Big Bang Nucleosynthesis	textbook
BAO	Baryon Acoustic Oscillations	textbook
EoS	Equation of state ($w = p/\rho$)	textbook
LSS	Large-Scale Structure	textbook
SNe Ia	Type Ia supernovae (distance-ladder standard candles)	Perlmutter et al. (1999); Riess et al. (1998)
EDE	Early Dark Energy	textbook
PBH	Primordial Black Hole	textbook
GW	Gravitational Wave	textbook
JWST	James Webb Space Telescope	textbook
Gravity, fields, and higher-dimensional constructions		
GR	General Relativity	textbook
QFT	Quantum Field Theory	textbook
QM	Quantum Mechanics	textbook
SM	Standard Model	textbook
GUT	Grand Unified Theory	textbook
KK	Kaluza–Klein decomposition	textbook
RS	Randall–Sundrum warped 5D model	textbook
AdS/CFT	Anti-de Sitter / Conformal Field Theory holography	textbook
Open-system formalism		
CPTP	Completely Positive Trace Preserving	textbook

continued on next page

Table 53 continued

Symbol	Meaning	Primary link
RG	Renormalization Group	textbook

Hierarchy of program-defined acronyms. For reference, Figure 37 records the six-tier hierarchy of the program-defined acronyms used across the GG-Theory cohort: the framework name and its constitutional charter; the four axioms and seven postulates; the bulk objects (GG-6, BI-6, OGN-5); the boundary law GDOS and its input packet SIP; the two four-pillar structures Gi-4 and Qi-4 with their sub-acronyms; and the energy–information Arisaka constant $\mathsf{Q_B}$ with the Arisaka equation that locks the four boundary currents. The diagram is shared verbatim with GG-A1-Constitution (its Figure 9), so that the acronym ledger is identical across the cohort.

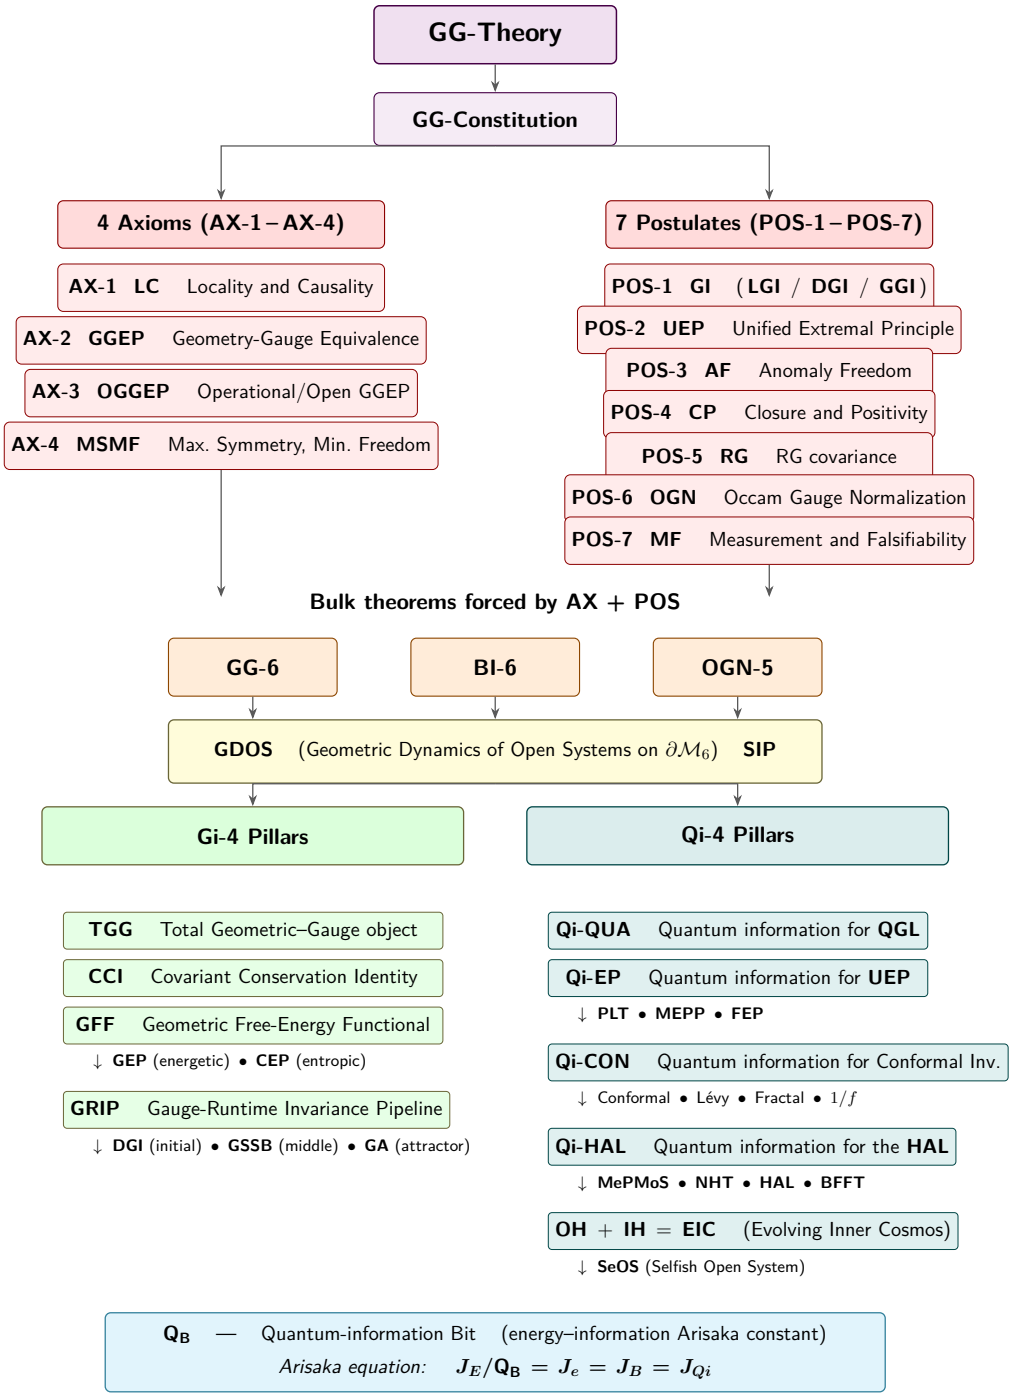

Figure 37: Hierarchy of program-defined acronyms in the GG-Theory program. The diagram is the same content as the abbreviation table set out as a hierarchy, so that the dependency between the names is visible at a glance rather than having to be reconstructed from the primary-link column. An arrow means the lower name is defined using the upper one; it does not mean the lower result depends on the upper for its proof.

B Glossary of Key Symbols

This appendix collects the key mathematical symbols used in GG-A10-Cosmos. The first block records the bulk-geometry coordinates and the OGN-5 integer ledger that every later derivation appeals to. The second records the inherited scale chain $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock}$. The third through seventh blocks record the cosmological observables of the π-locked inventory — the density fractions, the Hubble background, the growth and primordial sectors, the nucleosynthesis ledger, and Kaxion dark matter — as audited in §68.1. The eighth block records the MSMF/OGN audit machinery.

Table 54: Glossary of key symbols used in the present paper. "Source" gives the theorem, section, or paper where the symbol is introduced or load-bearing. The Source column is the column that matters for checking: it names the theorem, section or paper where each symbol is introduced, so a symbol met in the middle of the paper can be traced to its definition in one step.

Symbol	Meaning	Source
(1) Bulk geometry and the OGN-5 ledger		
M_6	6D bulk manifold $M_4 \times I_r \times S^1_\theta$	GG-A2-GG6
ds^2_6	Canonical scale-profile metric $e^{-2r}\eta_{\mu\nu}\,dx^\mu dx^\nu + dr^2 + d\theta^2$	GG-A2-GG6
r	Radial coordinate, $r \in [0, K]$ (the cosmic chronology axis)	—
θ	Compact angular coordinate, $\theta \in [0, 2\pi)$	—
S^1_θ	Compact direction circle (hosts the Kaxion holonomy)	—
I_r	Radial interval	—
$(M, N, L; J, K)$	OGN-5 ledger $= (1, 3, 5; 7, 35)$	GG-A4-Euclid
$K = LJ$	Total scale depth, $K = LJ = M^2 + N^2 + L^2 = 35$	GG-A2-GG6
$N = 3$	The OGN-5 ledger's second entry; written N_c wherever the color reading is meant	GG-A3-BI6 §7
$N_c = 3$	Color triplicity, the three colors of $\mathrm{SU}(3)_c$ — what the ledger's second entry counts; fixed by the dual-lock closure, the trace ratio being the independent check	GG-A3-BI6, OGN-5 theorem, Step 3
$N_g = 3$	Generation (family) count. *Not* a ledger member and not fixed by the ledger; the three-leg theorem	Arisaka, 2026, P2-N3
$N_f = 6$	Quark-flavor count in a beta function at high scale, $N_f = 2N_g$. A third quantity, distinct from the other two	PDG
(2) Scale chain		
$E_{\rm Pl}$	Planck energy $\approx 1.22 \times 10^{19}$ GeV	—
$E_{\rm GUT}$	GUT scale $= E_{\rm Pl}\, e^{-2\pi} \approx 2.28 \times 10^{16}$ GeV	GG-A2-GG6
$M_{\rm lock}$	Lock scale $= \frac{1}{2} M_{\rm Pl}\, e^{-K} \approx 3.85$ TeV	GG-A2-GG6
$m_\Lambda^{\rm geom}$	Dark-energy mass scale $= M^2_{\rm lock}/M_{\rm Pl} \approx 6$ meV	§68.1

continued on next page

Table 54 continued

Symbol	Meaning	Source
(3) The π-locked density inventory		
Ω_m	Matter density fraction $= 1/\pi \approx 0.3183$	Thm. 74.1
Ω_Λ	Vacuum (dark-energy) fraction $= 1 - 1/\pi \approx 0.6817$	Thm. 74.1
Ω_b	Baryon fraction $= 1/(2\pi^2) \approx 0.0507$	Thm. 74.1
Ω_{DM}	Dark-matter fraction $= 1/\pi - 1/(2\pi^2) \approx 0.2676$	Thm. 74.1
w	Dark-energy equation of state $= -1$ (exact)	§68.1
(4) The Hubble background and epochs		
H_0^{CMB}	The one background read through the sound-horizon dictionary, ≈ 67.4 km/s/Mpc	Thm. 68.8
H_0^{late}	Late-time Hubble value on ladder-free probes: the background 66.8 km/s/Mpc; the ladder's 73.5 read as a systematic	Thm. 68.8
z_*	Corridor-switch redshift $= e^J = e^7 \approx 1097$ (atomic decoupling $z_{\mathrm{dec}} \approx 1090$, §42.1)	Thm. 68.8
z_{eq}	Matter–radiation equality redshift ≈ 3400	§68.1
z_{rec}	Empirical recombination redshift $= 1089.92 \pm 0.25$	Planck Collaboration (2018)
(5) Structure growth and the primordial sector		
σ_8	Linear matter-clustering amplitude, computed 0.80 ± 0.01 (observed 0.811)	§68.1
S_8	Clustering parameter $\sigma_8\sqrt{\Omega_m/0.3}$, computed 0.825 against the CMB's 0.832 and cosmic shear's 0.759 (the tension shared)	§68.1
n_s	Scalar spectral tilt, 0.9634 exact on the plateau ($1-2/N^* = 0.9620$ at leading order; D*)	§68.1
A_s	Scalar amplitude $= 2.10 \times 10^{-9}$ (asserted input)	§68.1
r	Tensor-to-scalar ratio; 0.0038 exact, $12/N^{*2} = 0.0043$ at leading order (derived-cond., two-sided)	§68.1
ΔN_{eff}	Effective extra-radiation deviation; ≤ 0.03 (a derived null to within the freeze-in ceiling)	§68.1
ℓ_D	CMB damping multipole $\approx$ 1350 (Boltzmann; BFFT overlay $\approx$ 1382)	§68.1
(6) Nucleosynthesis ledger		
η_B	Baryon-to-photon ratio $\approx 5\times10^{-10}$ ($\times 3^{\pm1}$; time-resolved, a factor of about ten low)	§68.1
^{7}Li/H	Primordial lithium-7 abundance $=$ SBBN$/\pi \approx 1.6 \times 10^{-10}$	§68.1
(7) Kaxion dark matter		

continued on next page

Table 54 continued

Symbol	Meaning	Source
m_{Kaxion}	Kaxion mass $\simeq 8.1$ μeV ($\pm 1.3\%$ at fixed f_K)	GG-C1-Kaxion
ν_K	Kaxion frequency window $\in [1.59, 2.13]$ GHz	GG-C1-Kaxion
f_{Kaxion}	Kaxion decay constant $= 7.1 \times 10^{11}$ GeV (bulk normalization 3.04×10^{13} GeV, calibrated)	GG-C1-Kaxion
θ_i	Initial misalignment angle ≈ 1.1 rad (read back from the relic abundance)	GG-C1-Kaxion
(8) MSMF/OGN audit machinery		
$\mathcal{R}$	Candidate rational expression	§68.2
$\mathcal{C}_\Sigma$	OGN-admissible candidate space on stratum Σ	§68.2
$\text{Cost}(\mathcal{R})$	MSMF cost: total absolute exponent budget	§68.2
ϵ_X	Empirical-precision target for parameter X on its stratum	§68.2
$\mathcal{R}^\star$	Cost-minimum unique winner under MSMF/OGN selection	§68.2

C Codebook glossary (compact)

A condensed version of the twenty-one-row codebook of Section 68.1, suitable for quick scan. Each row gives the observable, its symbol, the predicted value (from GG-A10-Cosmos), the empirical value (with reference), and the status under the MSMF/OGN principle.

Table 55: Compact codebook glossary. D = Derived, D* = Derived (cond.), A = Asserted. Predicted values from the present paper; empirical values from Planck Collaboration (2018), Casertano et al. (2026) (with Riess et al. (2022)), Asgari et al. (2021). Every predicted value in the fourth column is this paper's; every empirical value in the fifth is quoted from the source named in the codebook proper. A status of D-star means the row is derived conditional on a named conjecture, and the conjecture is named there rather than here.

Observable	Symbol	Predicted	Empirical	Status
Matter fraction	Ω_m	$1/\pi \approx 0.3183$	0.315 ± 0.007	D*
Vacuum fraction	Ω_Λ	$1 - 1/\pi \approx 0.6817$	0.685 ± 0.007	D*
Dark-energy EoS	w	-1 (exact)	-1.028 ± 0.032	D*
Baryon fraction	Ω_b	$1/(2\pi^2) \approx 0.0507$	0.0493 ± 0.0006	D* (2.3σ residual)
Dark-matter fraction	Ω_{DM}	$1/\pi - 1/(2\pi^2) \approx 0.2676$	0.2657 ± 0.0070	D*
Baryon-to-photon ratio	η_B	$\approx 5 \times 10^{-10}$ (central 4.89)	$(6.12 \pm 0.04) \times 10^{-10}$	D*
Lithium-7	$^7\mathrm{Li}/H$	SBBN/$\pi \approx 1.6 \times 10^{-10}$	$(1.6 \pm 0.3) \times 10^{-10}$	D*
Hubble (against the CMB)	H_0^{CMB}	66.8 ± 0.2 km/s/Mpc (Route I, zero sky input)	67.4 ± 0.5	D*
Hubble (against the ladder)	H_0^{late}	66.8 ± 0.2 km/s/Mpc (one background)	73.50 ± 0.81	D*
Switch / decoupling	z_*	switch $e^7 \approx 1096.6$ (not scored); atomic chain 1091.5	1089.92 ± 0.25 (decoupling)	D*
Scalar tilt	n_s	0.9634 exact ($1 - 2/N^* = 0.9620$ l.o.)	0.9649 ± 0.0042; 0.9682–0.9728 (2025; -1.5 to -3.2σ)	D*
Scalar amplitude	A_s	2.10×10^{-9} (input)	$(2.10 \pm 0.04) \times 10^{-9}$	A
Tensor-to-scalar	r	0.0038 exact ($12/N^{*2} = 0.0043$ l.o.)	< 0.036 (95%)	D*
Clustering amplitude	σ_8	0.80 ± 0.01 (from A_s)	0.811 ± 0.006	D*
Clustering (vs. CMB)	S_8	0.825 ± 0.010	0.832 ± 0.013	D*
Clustering (vs. shear)	S_8	the same 0.825 ± 0.010	$0.759^{+0.024}_{-0.021}$	D*
Damping multipole	ℓ_D	1350 ± 3	$\approx 1345 \pm 15$	D*
Neutrino dev.	ΔN_{eff}	≤ 0.03 ($N_{\mathrm{eff}} = 3.044$)	$N_{\mathrm{eff}} = 2.99 \pm 0.17$	D
Kaxion mass	m_{Kaxion}	8.1 μeV (±1.3%; calibrated scale)	—	D*
Kaxion decay constant	f_{Kaxion}	7.1×10^{11} GeV (calibrated)	—	D*
Inflation scale	$V_{\mathrm{inf}}^{1/4}$	8.1×10^{15} GeV	—	D*

Counting summary. 1 D + 19 D* + 1 A = 21 observables; 10 named conjectures (Appendix R) close all 20 of the D* and A entries.

D Global naming convention for cross-cohort citation

Beginning with Arisaka (2026, A2-GG6) and extended across the cohort, each load-bearing item of the GG-Theory program carries a globally unique identifier of the form *Prefix-Suffix-N*, where the Prefix names the item type and the Suffix names the paper:

- **T-Suffix-N** for the Nth Theorem.
- **L-Suffix-N**, **C-Suffix-N**, **D-Suffix-N**, **R-Suffix-N** for the Nth Lemma, Corollary, Definition, and Remark.
- **P-Suffix-N** for the Nth Prediction.
- **O-Suffix-N** for the Nth Objection (the Common Objections Part).
- **F-Suffix-N** for the Nth Falsifier (the kill criteria that would retire the theory).
- **I-Suffix-N** for the Nth Innovation (e.g., I-Cosmos-1 is the cosmology-from-a-constitutional-layer innovation above).
- **PS-Suffix-N** for the Nth Problem-Solved / status entry.

The canonical Suffix list for the twenty-two cohort papers is: *Constitution* (A1), *GG6* (A2), *BI6* (A3), *Euclid* (A4), *GDOS* (A5), *GRIP* (A6), *Qi4* (A7), *QM* (A8), *SM* (A9), *Cosmos* (A10, this paper), *Origin* (A11), together with the phenomenology preprints *KK* (P1), *N3* (P2), *Electron* (P3), *Proton* (P6), the foundations preprints *Real* (P4) and *CPT* (P5), and the consequence preprints *Kaxion* (C1), *Hubble* (C2), *BH* (C3), and the white papers *Symmetry* (W1) and *Holonomy* (W2). The four axioms (AX-1 . . . AX-4) and seven postulates (POS-1 . . . POS-7) of Arisaka (2026, A1-Constitution) keep their existing cohort-wide names; the global identifier is the same item viewed from outside its home paper. The present paper, GG-A10-Cosmos, uses the suffix *Cosmos*: its theorems are T-Cosmos-N, its falsifiers F-Cosmos-N, its innovations I-Cosmos-N (the ten above), its problems-solved PS-Cosmos-N, and its objections O-Cosmos-N.

E Direct derivation of the 4D Einstein equation, the Friedmann equations, and ΛCDM

What this appendix teaches. The full arc from the six-dimensional bulk to the concordance model, walked in three steps at first-year-graduate pace. Step 1: the four-dimensional Einstein equation is *derived* from the GG-6 bulk — the stage and its integer ledger, the bulk action read term by term from its earned home in the cohort, the Einstein–Hilbert term descending with its scale weight, Newton's constant computed as an overlap integral, conservation inherited as a Ward identity, and the open-system correction bounded, with live falsifiers. Step 2: the Friedmann equations are derived from that equation exactly as the textbooks do — Christoffel symbols, Ricci tensor, Einstein tensor, perfect-fluid source, every index traced — closed by a worked computation of the age of the universe. Step 3: the ΛCDM inventory is read from the geometric ledger rather than fitted — components, fractions, and initial conditions, with the register of every entry stated — and the assembled expansion law is written with no fitted continuous parameter. Fourteen problems with complete worked solutions.

Read alongside: Ryden (2017) chs. 3–6 for a first pass; Dodelson (2003) ch. 2 or Baumann (2022) ch. 1 at course level; Weinberg (2008) ch. 1 at depth; for the general-relativity prerequisites of Step 2, Wald (1984) or Misner *et al.* (1973).

Before the indices, the arc of discovery — because every equation below was forced on physics by an observation. In 1929 Hubble turned Leavitt's Cepheid ruler on the spiral nebulae and found the recession law $v = H_0 d$ (Hubble, 1929), making Friedmann's and Lemaître's expanding solutions (Friedmann, 1922; Lemaître, 1927) physics rather than curiosity. In 1965 Penzias and Wilson found the 3-kelvin glow (Penzias and Wilson, 1965) that Dicke's group identified at once as the relic of a hot beginning (Dicke et al., 1965); COBE measured its blackbody shape and first anisotropies (Smoot *et al.*, 1992); and the supernova teams of 1998–99 (Perlmutter et al., 1999; Riess et al., 1998) found the acceleration that put Λ into the inventory. Planck then compressed all of it into six numbers (Planck Collaboration, 2018). The equations this appendix derives are the machine those six numbers feed — and the question the rest of the course answers is why the numbers take the values they do.

Figure 38 draws the comoving Hubble radius across the whole history. Standard pedagogical and historical treatments of the Friedmann–Lemaitre–Robertson–Walker background cosmology, on which the derivations of this appendix follow line by line, are given in (Bianchi, 1902; Dodelson, 2003; Misner *et al.*, 1973; Mukhanov, 2005; Peebles, 1993; Wald, 1984; Weinberg, 1972, 2008).

This appendix has a single aim: to show that the Friedmann equations — the equations on which all of modern cosmology rests — are not a starting assumption but a consequence of geometry. We build the case in three steps, and the reader can navigate by them.

Step 1. From the bulk to the Einstein equation (§E.2–E.6). We derive four-dimensional gravity itself: the six-dimensional stage and its integer ledger, the bulk action read term by term from its earned home in the cohort, the Einstein–Hilbert term descending to the boundary with its scale weight, Newton's constant computed as an elementary overlap integral, and the boundary Einstein equation obtained by variation and zero-mode

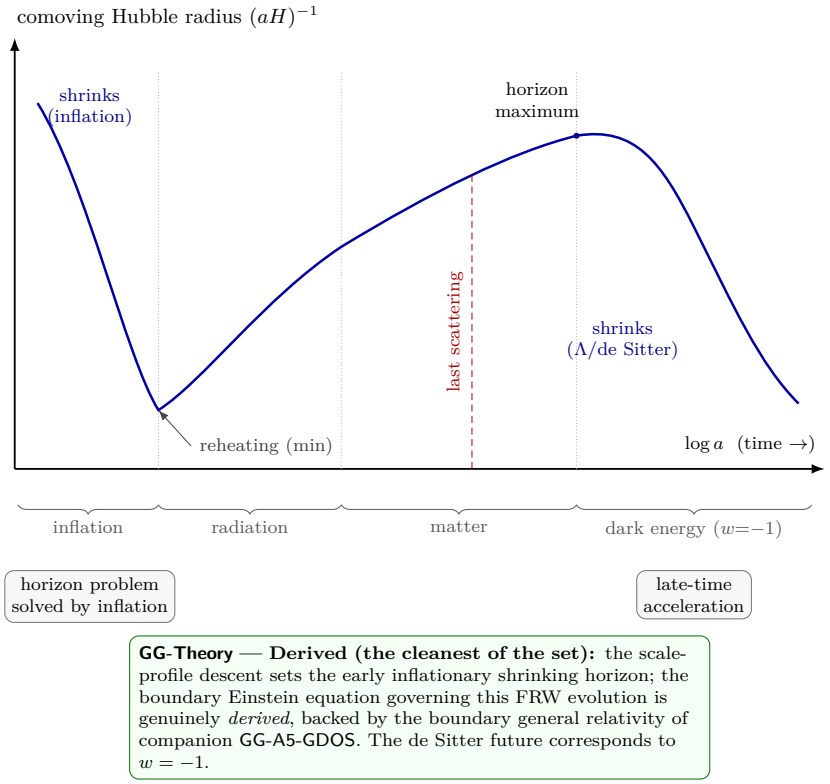

Figure 38: FRW horizons: the comoving Hubble radius $(aH)^{-1}$ versus $\log a$. Read left to right in time. During inflation the comoving Hubble radius *shrinks*, reaching a minimum at reheating; it then *grows* through the decelerating radiation- and matter-dominated eras (so the causal horizon grows and scales re-enter), reaches a maximum near the end of matter domination, and *shrinks* again once dark energy dominates and expansion accelerates toward a de Sitter future. This single curve teaches both the horizon problem (resolved by the early shrinking phase) and late-time acceleration. Last scattering is marked in the matter era. **GG-Theory derivation:** the scale-profile descent sets the early inflationary shrinking horizon, and the boundary Einstein equation governing this FRW evolution is genuinely *derived* – backed by the boundary general relativity of companion GG-A5-GDOS; the de Sitter future corresponds to $w = -1$. This is a foundational, Derived figure, the cleanest of the set.

projection — its conservation law inherited as a Ward identity, its open-system correction bounded, its falsifiers named. What every textbook takes as the axiom of cosmology is here the first theorem.

Step 2. From the Einstein equation to the Friedmann equations (§E.7–E.14). We then derive the Friedmann equations exactly as a textbook does, from scratch, at the depth a first-year graduate student can follow line by line — metric, connection, curvature, Einstein tensor, perfect-fluid source, and the two Friedmann equations, closed with a worked numerical example. Nothing here is unusual, and nothing is skipped — the point of Step 1 is precisely that nothing needs to change.

Step 3. From the Friedmann equations to ΛCDM (§E.15–E.18). Finally the right-hand side: the free menu $\{\Omega_i, w_i\}$ of standard cosmology is replaced by the geometric ledger — which components exist, with which equations of state, in which π-locked proportions, from which initial conditions — and the assembled expansion law is written out with no fitted continuous parameter, every entry carrying its register.

This is the first concrete instance of the belief that animates the whole program — that the contents and laws of the universe should follow from geometry, in the lineage of Euclid, who asked that the world be built from the fewest self-evident things, and of Newton, who turned Kepler's measured curves into theorems of a single law. The reader who wants only the background equations may read the boxed displays (472), (475) and (477) and skip ahead; the reader who wants to verify every step will find no gap papered over.

E.1 The four debts of the textbook derivation

Before the first index is traced, it is worth stating plainly what the textbook version of this subject — general relativity assumed, the Friedmann equations derived, the inventory fitted — must supply by hand. Four debts, and they are not small.

1. **The law of gravity.** The textbook derivation of Step 2 begins from the Einstein equation (470). Nothing in that derivation explains *why* gravity obeys that law rather than another; general relativity is the input, not the output.

2. **Newton's constant.** The single number G sets the strength of the link between geometry and matter, and it enters as a measured constant. The derivation has nothing to say about why it has the value it does.

3. **The cosmic inventory.** The right-hand side of the Friedmann equation is a *menu*: one writes down whatever components $\{\Omega_i, w_i\}$ the data seem to require — radiation, baryons, and then, when the data demand them, cold dark matter and a cosmological constant. Standard cosmology does not say what these components are, why they exist, or why they appear in the particular proportions we measure. The near-equality of matter and dark energy *today* — the cosmic coincidence — is left as a fact to be noted, not explained.

4. **The initial state.** The smoothness, the flatness, and the spectrum of initial fluctuations are imposed as boundary conditions.

None of this is a flaw in the mathematics; it is the boundary of what a four-dimensional, general-relativistic description *can* do. Such a description accommodates the universe we observe — with great precision — but it does not derive it: every cosmological parameter is a number to be measured and inserted. A theory in the Euclidean and Newtonian tradition would instead ask these numbers to follow from something deeper. The three steps of this appendix are organized as exactly that request: Step 1 pays debts 1 and 2 (the law and its constant), Step 3 pays debts 3 and 4 (the inventory and the initial state), and Step 2 — deliberately untouched, textbook line by textbook line — is the demonstration that paying the debts changes the *provenance* of cosmology, not its equations. *(Physics intuition: the honest way to admire the Friedmann equations is to ask who is paying for them. The textbook borrows four times and never settles; this appendix walks the same street with the receipts.)*

E.2 Step 1 — the six-dimensional stage and its ledger

GG-Theory answers the four debts of §E.1 not one at a time but together, because they share a single source: the four-dimensional world on which the Friedmann equations live is the boundary of a six-dimensional geometry. Step 1 walks that derivation at the same depth as the textbook material of Step 2 — stating each object, doing the arithmetic the reader can check by hand, and naming where the rigorous proof lives in the cohort: chiefly the boundary general relativity of Arisaka (2026, A5-GDOS) (Theorems T-GDOS-8.1–8.7 of its Part IV), the bulk construction of Arisaka (2026, A2-GG6), and the assembly interlude of Arisaka (2026, A9-SM), whose station S-1a earns the bulk action this Step spends. For a reader meeting warped products for the first time: the line element below is an ordinary four-dimensional metric multiplied by a factor that depends on where you stand in a fifth direction — nothing more exotic than a map whose scale bar changes from place to place. The new content is the *reading* of that fifth direction: it is the renormalization-group altitude, so that geometry and energy scale become one coordinate.

The six-dimensional stage. The geometry is a product geometry (Arisaka, 2026, A2-GG6): four ordinary spacetime dimensions x^μ, a radial direction r, and a compact angle θ, with line element

$$ds_6^2 \;=\; e^{-2r}\,\eta_{\mu\nu}\,dx^\mu dx^\nu \;+\; dr^2 \;+\; d\theta^2. \tag{431}$$

The scale profile e^{-2r} shrinks four-dimensional lengths as one moves inward along r, which runs from the ultraviolet edge at $r = 0$ (the unification scale E_{GUT}) down a total scale depth $K = 35$; the angle θ closes the compact circle that hosts the Kaxion. The depth is not chosen — it is the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, and the value 35 is fixed twice over,

$$K \;=\; LJ \;=\; 5\times 7 \;=\; 35 \qquad \text{and} \qquad K \;=\; M^2+N^2+L^2 \;=\; 1+9+25 \;=\; 35, \tag{432}$$

a redundancy that is one of the integer-ledger's internal consistency checks. *Physics intuition: r* is a renormalization-group altitude, and it is precisely the descent parameter of the open-system boundary law GDOS. Moving downward in r is integrating out the high-energy bulk modes — the flow that the GRIP runtime executes — so the single coordinate r carries the entire arrow from the hot unified edge to the cold dilute world we inhabit. Cosmic history is a journey down this one axis.

E.3 Step 1 — the bulk action, read term by term

Every dynamical statement in this appendix descends from one object, and the reader should see it whole before any reduction is performed. The bulk carries the unique admissible two-derivative action

$$S_6 \;=\; \int_{M_6} \sqrt{-G}\,\Big[\underbrace{\tfrac{1}{2\kappa_6^2}\,(R_6 - 2\Lambda_6)}_{\text{geometry}} \;-\; \underbrace{\tfrac{1}{4}\,\mathcal{F}^{A}_{MN}\mathcal{F}^{A\,MN}}_{\text{gauge sector}} \;+\; \underbrace{i\,\overline{\Psi}\,\Gamma^{M} D_M \Psi}_{\text{matter spinors}} \;-\; \underbrace{\tfrac{1}{12}\,H_{MNP}H^{MNP}}_{\text{flux sector}}\Big], \tag{433}$$

plus the Green–Schwarz terms that anomaly closure forces. This is not an ansatz of the present paper, and this appendix does not ask the reader to accept it on citation alone: its field content is forced by the stage, metric, and gauge-group theorems of Arisaka (2026, A2-GG6) (T-GG6-0/1/2/3), its consistency by the anomaly closure of Arisaka (2026, A3-BI6), and the uniqueness of its dynamics — why these terms and no others at lowest derivative order — is earned brace by brace in the assembly interlude of Arisaka (2026, A9-SM) (station S-1a), with the register of every claim stated there: the form of each term **Derived** at lowest derivative order, the matter content **Derived (cond.)** on a named minimality lemma. For Step 1 only two facts about Eq. (433) matter, and both belong to the geometry brace: the Einstein–Hilbert term exists with a definite coefficient, and no other term in the action contains second derivatives of the metric. Everything else — gauge fields, spinors, flux — rides on the right-hand side of the equation we are about to derive, as source. *(Physics intuition: an action is a recipe book, and this book has one page. The rest of the cohort spends the gauge and matter braces to build particle physics; cosmology, the coarsest of all boundary sciences, needs only the first brace and the bookkeeping of the others' energy.)*

E.4 Step 1 — the Einstein–Hilbert term descends to the boundary

The tool that carries a six-dimensional action to a four-dimensional theory is mode expansion — Kaluza–Klein reduction — and a first-year student has already used it twice in other clothes. On the compact circle it is the Fourier series: any field decomposes into harmonics $e^{in\theta}$, and the $n = 0$ harmonic is the part that does not vary around the circle. On the interval it is the particle in a box: standing waves with a discrete tower of energies, whose lowest mode again varies least. A six-dimensional field is therefore, to a boundary observer, an infinite family of four-dimensional fields: one massless *zero mode* (constant across the compact directions) and a tower of heavy partners whose masses start at the lock scale $M_{\text{lock}} \approx 3.85$ TeV — far above every cosmological energy, so the background dynamics of this appendix sees the zero modes alone (the tower is the collider story of the companion papers, Theorem T-GDOS-9.1 of Arisaka, 2026, A5-GDOS). What remains is to reduce the geometry brace of Eq. (433) and watch four-dimensional gravity emerge.

Gravity as a boundary zero mode (paying debt 1). A field living in six dimensions can be expanded in standing waves (Kaluza–Klein modes) along the two compact directions. The lowest — the *zero mode*, constant across r and θ — is the ordinary four-dimensional graviton: we promote the flat metric $\eta_{\mu\nu}$ in (431) to a slowly varying four-dimensional field $g_{\mu\nu}(x)$. Two factors then control how the six-dimensional Einstein–Hilbert action reduces. First, the measure of the scale-profile block carries the scale profile on each of the four spacetime legs, so its determinant scales as $g_6 = e^{-8r}\, g_4 \cdot (1)(1)$ and

$$\sqrt{g_6} \;=\; e^{-4r}\,\sqrt{g_4}\,. \tag{434}$$

Second, the four-dimensional Ricci scalar built from $g_{\mu\nu}(x)$ enters the six-dimensional curvature with a compensating inverse-metric factor e^{+2r} (raising two indices on the scale-profile block), $R_6 \supset e^{+2r}\, R_4[g]$. Their product is the scale weight $e^{-4r}\cdot e^{+2r} = e^{-2r}$, and the gravitational action collapses to four dimensions:

$$M_6^4 \int d^4x\, d\theta\, dr\, \sqrt{g_6}\, R_6 \;\;\longrightarrow\;\; M_6^4 \Big(\int_0^{2\pi} d\theta \int_0^{K} e^{-2r}\, dr\Big) \int d^4x\, \sqrt{g_4}\, R_4. \tag{435}$$

The four-dimensional Einstein–Hilbert term — the left-hand side of the Einstein equation (470) that Step 2 will begin from — emerges with a definite coefficient (GG-A5-GDOS T8.1). The law of gravity on the boundary is therefore not an axiom; it is the shadow that six-dimensional gravity casts on its four-dimensional edge. *Physics intuition:* general relativity is what the bulk geometry looks like to an observer who cannot resolve the two extra directions — and the scale weight e^{-2r} is the statement that the deep infrared (r large) contributes almost nothing to what that observer calls "gravity."

E.5 Step 1 — Newton's constant computed, and the scale ladder

Debt 2 is the value of the coupling, and here the geometry does something no four-dimensional theory can: it computes the coefficient it just derived, as an integral the reader can do on one line.

Newton's constant from a scale-overlap integral (paying debt 2). The coefficient in (435) is the entire content of Newton's constant. It is the overlap of the graviton zero mode with itself across the compact directions, with the scale weight just derived,

$$M_{\rm Pl}^2 \;=\; M_6^4 \int_0^{2\pi} d\theta \int_0^{K} e^{-2r}\, dr. \tag{436}$$

Both integrals are elementary. The angular one gives 2π. The radial one is

$$\int_0^{K} e^{-2r}\, dr \;=\; \Big[-\tfrac{1}{2} e^{-2r}\Big]_0^{K} \;=\; \tfrac{1}{2}(1 - e^{-2K}), \tag{437}$$

so their product is $2\pi \cdot \frac{1}{2}(1 - e^{-2K}) = \pi(1 - e^{-2K})$ and

$$\boxed{M_{\rm Pl}^2 \;=\; \pi\, M_6^4\, (1 - e^{-2K})} \tag{438}$$

(GG-A5-GDOS T8.2). With the ledger depth $K = 35$ the scale residual is

$$e^{-2K} = e^{-70} \approx 3.98 \times 10^{-31}, \tag{439}$$

utterly negligible, so to one part in 10^{30} the relation collapses to the clean statement $M_{\rm Pl}^2 \simeq \pi M_6^4$. Read in the natural direction this ties Newton's constant, $G = 1/(8\pi M_{\rm Pl}^2)$ in reduced-Planck units, to a single six-dimensional scale M_6 and the pure number π: it is not an independent four-dimensional parameter. The framework carries one dimensionful unit — conventionally $M_{\rm Pl} \equiv G_N^{-1/2}$, with M_6 the *derived* combination, as the constitution requires (Arisaka, 2026, A2-GG6) — so what the geometry fixes is the scale profile relating them, not a numerical prediction of G_N, which sets the units. *Physics intuition:* G measures how large the hidden compact geometry looks from the boundary; the factor π is the geometric size of the throat, and the missing e^{-70} is how little of that throat the infrared world can feel. Where standard cosmology treats G as one free constant among many, the framework ties it to the bulk geometry by a fixed relation — one dimensionful unit, and no free four-dimensional gravitational dial.

The same scale profile fixes the scale ladder (worked numbers). The habit of computing rather than measuring is clearest in the energy scales that standard physics reads off experiment. Three rungs of one ladder are fixed by e raised to integer-ledger powers of the Planck scale (Arisaka, 2026, A2-GG6; the scale chain of Part II). From the ultraviolet edge $E_{\rm Pl} = 1.22 \times 10^{19}$ GeV:

$$E_{\rm GUT} \;=\; E_{\rm Pl}\, e^{-2\pi} \;=\; 1.22 \times 10^{19} \times e^{-2\pi}\ \text{GeV} \;=\; 2.28 \times 10^{16}\ \text{GeV}, \tag{440}$$

$$M_{\rm lock} \;=\; \tfrac{1}{2} M_{\rm Pl}\, e^{-K} \;=\; \tfrac{1}{2}\,(1.22 \times 10^{19})\, e^{-35}\ \text{GeV} \;=\; 3.85 \times 10^{3}\ \text{GeV} \;=\; 3.85\ \text{TeV}. \tag{441}$$

The first is one full turn of the compact circle (2π) below the Planck scale; the second is the full scale depth ($K = 35$) below it. Their ratio — the great desert of particle physics, from the unification scale down to the TeV seam — is then itself a closed expression,

$$\frac{E_{\rm GUT}}{M_{\rm lock}} \;=\; 2\, e^{\,K-2\pi} \;=\; 2\, e^{\,28.72} \;\approx\; 5.9 \times 10^{12}, \tag{442}$$

a hierarchy of nearly six trillion produced entirely by the integer difference between the scale depth 35 and the single GUT turn. Nothing here is tuned; the numbers are exponentials of integers from the ledger (432). *Physics intuition:* the vast hierarchies of physics are not fine-tunings but a few notches of scale profile, read off the same ruler that fixes G.

E.6 Step 1 — the boundary Einstein equation, earned

From action to equation. An action is not yet a law of motion. The law follows from the same variational principle the reader knows from classical mechanics — $\delta S = 0$ — executed for the metric field: varying Eq. (433) with respect to G_{MN} yields the six-dimensional Einstein equation

$$M_6^4\, G_{MN} \;+\; \Lambda_6\, g_{MN} \;=\; T_{MN}^{\rm total}, \tag{443}$$

with T^{total}_{MN} collecting the gauge, spinor, and flux sectors of the action (the five-field variational accounting is recorded in the canonical-Lagrangian section of Arisaka, 2026, A5-GDOS). Restricting to the graviton zero mode of §E.4 and integrating the compact directions with the scale weight of §E.5 projects Eq. (443) onto the boundary:

$$G_{\mu\nu} \;+\; \Lambda\, g_{\mu\nu} \;=\; 8\pi G\, T_{\mu\nu} \;+\; \Delta T^{\text{GS}}_{\mu\nu} + \Delta T^{\text{KK}}_{\mu\nu}, \tag{444}$$

the four-dimensional Einstein equation with two additions the honest register must carry: the boundary cosmological term Λ, whose existence rides down with the geometry brace and whose *value* is the terminus story of Part IV (the lock-scale seesaw, taken up in Step 3), and the calculable inflow corrections $\Delta T_{\mu\nu}$, bounded below. This is Theorem T-GDOS-8.1/8.4 of Arisaka (2026, A5-GDOS) at course pace. *(Physics intuition: the Einstein equation was never wrong to postulate — it was only unexplained. Here it is the unique shadow a one-page bulk recipe can cast on its edge: the reader has now watched the left-hand side of all of cosmology assemble itself from a variational principle and two elementary integrals.)*

Conservation inherited, not imposed. The textbook pairs the Einstein equation with a separate physical demand: energy–momentum conservation, $\nabla_\mu T^{\mu\nu} = 0$. In textbook general relativity the pairing is consistent because of the contracted Bianchi identity $\nabla_\mu G^{\mu\nu} = 0$ — a geometric fact accepted from differential geometry. Here it has a sharper origin: that conservation law is the *Ward identity* of bulk diffeomorphism invariance restricted to the boundary (GG-A5-GDOS T8.3). The freedom to relabel coordinates in the six-dimensional bulk becomes, on the four-dimensional edge, the statement that the source cannot leak. One practical consequence will surface in Step 2 and is worth flagging now: of the three FRW equations (two Friedmann plus continuity), only two are independent, and the redundancy is exactly this identity wearing cosmological clothes — worked explicitly in the continuity discussion of §E.12 and in Problem E.12.1. *(Physics intuition: a bulk-exported source* cannot *violate continuity without breaking a symmetry of the geometry itself. Consistency is inherited, not imposed by hand.)*

The open-system correction, bounded — why the equations are exactly textbook today. The full boundary field equation is not quite the bare Einstein equation: it carries calculable Green–Schwarz and Kaluza–Klein inflow terms $\Delta T^{\text{GS}}_{\mu\nu} + \Delta T^{\text{KK}}_{\mu\nu}$ (GG-A5-GDOS T8.4). These switch on only near the lock scale M_{lock} and fall away in the infrared, scaling as a power of the ambient energy over M_{lock}. On the present cosmological boundary the ambient scale is the Hubble energy $H_0 \sim 1.4 \times 10^{-33}$ eV, so the relative size of the corrections is of order

$$\Big(\frac{H_0}{M_{\text{lock}}}\Big)^2 \;\sim\; \Big(\frac{1.4\times 10^{-33}\,\text{eV}}{3.85\times 10^{12}\,\text{eV}}\Big)^2 \;\sim\; 10^{-91}, \tag{445}$$

an utterly negligible number. This is why every equation of this appendix holds *unmodified* today: the geometry corrects general relativity only at the TeV seam — exactly where it should, and nowhere a cosmological measurement can yet reach. *Physics intuition:* the new physics is real but it lives at the throat; by the time the universe has expanded to its present size, the throat is 10^{91}

times too far away to matter for the background.

The bound comes with instruments. A bounded correction is an honest claim only if something could measure it. Two instruments are on the table. Astrophysically, the same inflow terms modify the Tolman–Oppenheimer–Volkoff equation of stellar structure in a calculable pattern (GG-A5-GDOS T8.6) — a named falsifier: a neutron-star observation matching the modified profile would be discovery, one excluding it at the throat-suppressed level would be trouble. Cosmologically, the growth sector of Part VI is the standing audit: any late-time departure from general-relativistic growth beyond Part VI's stated bands would land exactly here. The claim of this subsection is therefore not "general relativity, amen" but "general relativity, plus a correction of stated size, with stated instruments."

A second route to the same edge. The zero-mode reduction is not the cohort's only road to boundary gravity: the black-hole companion derives the same four-dimensional gravity from the entanglement side — induced gravity on the boundary, with the Bekenstein–Hawking $A/4$ recovered exactly (Appendix N; Arisaka, 2026, C3-BH). Two independent constructions landing on one Einstein equation is itself a consistency check of the geometry. *(Physics intuition: when the KK road and the entanglement road meet at the same equation, the equation stops being a route-dependent artifact and starts being a place.)*

Step 1's ledger, closed. What Step 1 has established, with registers: the four-dimensional Einstein equation (444) — form **Derived** (T-GDOS-8.1/8.4, consumed at course pace above); Newton's constant tied to the bulk scale by $M_{\rm Pl}^2 = \pi M_6^4(1 - e^{-2K})$ — **Derived** (T8.2), the overall unit set by G_N as the constitution requires; conservation — **Derived** as the T8.3 Ward identity; the open-system correction — **bounded**, not assumed absent, with the T8.6 and Part-VI falsifiers live. Debts 1 and 2 of §E.1 are paid. Step 2 can now begin exactly where the textbook begins — with one difference invisible in the equations and decisive in the logic: nothing on the page ahead is postulated.

E.7 Step 2 — the FRW metric and its conventions

Statement. The cosmological principle — spatial homogeneity and isotropy of the observer-accessible four-dimensional sector — restricts the metric to the Friedmann–Robertson–Walker form. In comoving spherical coordinates (t, r, θ, ϕ) the line element is

$$ds^2 \;=\; -\,c^2\,dt^2 \;+\; a(t)^2\left[\frac{dr^2}{1-k\,r^2} \;+\; r^2\,d\theta^2 \;+\; r^2\sin^2\theta\,d\phi^2\right], \tag{446}$$

where $a(t)$ is the dimensionless scale factor (the macro $a(t)$), normalized so that $a(t_0) = 1$ today, and $k \in \{-1, 0, +1\}$ is the sign of the constant spatial curvature after r is rescaled to absorb its magnitude. We adopt the mostly-plus signature $(-,+,+,+)$ and the index convention that Greek indices $\mu, \nu, \ldots$ run over $\{0,1,2,3\} = \{t, r, \theta, \phi\}$. For the remainder of the algebra we set $c = 1$ and restore it only in the worked numerical example.

The metric and inverse-metric components. Reading (446) off as a diagonal matrix $g_{\mu\nu} = \mathrm{diag}(g_{tt}, g_{rr}, g_{\theta\theta}, g_{\phi\phi})$,

$$g_{tt} = -1, \qquad g_{rr} = \frac{a^2}{1-kr^2}, \qquad g_{\theta\theta} = a^2 r^2, \qquad g_{\phi\phi} = a^2 r^2 \sin^2\theta, \tag{447}$$

and because the metric is diagonal its inverse is the componentwise reciprocal,

$$g^{tt} = -1, \qquad g^{rr} = \frac{1-kr^2}{a^2}, \qquad g^{\theta\theta} = \frac{1}{a^2 r^2}, \qquad g^{\phi\phi} = \frac{1}{a^2 r^2 \sin^2\theta}. \tag{448}$$

The metric determinant is

$$g \;=\; \det g_{\mu\nu} \;=\; -\frac{a^6 r^4 \sin^2\theta}{1-kr^2}, \qquad \sqrt{-g} \;=\; \frac{a^3 r^2 \sin\theta}{\sqrt{1-kr^2}}, \tag{449}$$

which we will need for the covariant volume element. Overdot denotes d/dt throughout, so $\dot{a} = da/dt$ and $\ddot{a} = d^2a/dt^2$, and the Hubble rate is defined as

$$H(t) \;\equiv\; \frac{\dot{a}}{a}, \qquad H_0 \;\equiv\; H(t_0) \text{ (the present value)}. \tag{450}$$

E.8 Step 2 — Christoffel symbols

Statement. The Levi-Civita connection is the unique metric-compatible, torsion-free connection, with components

$$\Gamma^{\lambda}{}_{\mu\nu} \;=\; \tfrac{1}{2}\, g^{\lambda\sigma}(\partial_\mu g_{\sigma\nu} + \partial_\nu g_{\sigma\mu} - \partial_\sigma g_{\mu\nu}). \tag{451}$$

Because $g_{\mu\nu}$ is diagonal and depends only on t (through a), on r (through $1/(1-kr^2)$ and r^2), and on θ (through $\sin^2\theta$), most of the $4 \times 10 = 40$ independent symbols vanish. We compute the nonzero ones explicitly.

Time-derivative (Hubble-flow) symbols. Only $g_{rr}, g_{\theta\theta}, g_{\phi\phi}$ carry t-dependence, and each is proportional to a^2, so $\partial_t g_{ii} = 2(\dot{a}/a)\, g_{ii} = 2H g_{ii}$ (no sum). The symbols with a single time index follow from (451). With $\lambda = 0$ (i.e. t), using $g^{tt} = -1$ and $\partial_t g_{ii} = 2H g_{ii}$,

$$\Gamma^{t}{}_{ij} = -\tfrac{1}{2}\, g^{tt}\, \partial_t g_{ij} = \tfrac{1}{2}\, \partial_t g_{ij} = H\, g_{ij} \qquad (i,j \in \{r,\theta,\phi\}), \tag{452}$$

explicitly

$$\Gamma^{t}{}_{rr} = \frac{a\dot{a}}{1-kr^2}, \quad \Gamma^{t}{}_{\theta\theta} = a\dot{a}\, r^2, \quad \Gamma^{t}{}_{\phi\phi} = a\dot{a}\, r^2 \sin^2\theta. \tag{453}$$

Similarly, the symbols with one time and one space index come from $\lambda = i$, $\mu = t$, $\nu = i$:

$$\Gamma^{i}{}_{tj} = \tfrac{1}{2}\, g^{ii}\, \partial_t g_{ij} = \frac{\dot{a}}{a}\, \delta^{i}{}_{j} = H\, \delta^{i}{}_{j} \qquad (\text{no sum on } i), \tag{454}$$

so

$$\Gamma^{r}{}_{tr} = \Gamma^{\theta}{}_{t\theta} = \Gamma^{\phi}{}_{t\phi} = \frac{\dot{a}}{a} = H. \tag{455}$$

These are the only symbols that carry the expansion; every appearance of $\dot{a}$ in the curvature traces back to them.

Purely spatial symbols. The remaining nonzero symbols are exactly those of the maximally symmetric three-space of curvature k, unchanged by the overall $a^2(t)$ factor because that factor cancels between g^{ii} and ∂g_{jj}:

$$\begin{aligned}
&\Gamma^{r}{}_{rr} = \frac{kr}{1-kr^2}, && \Gamma^{r}{}_{\theta\theta} = -\,r(1-kr^2), && \Gamma^{r}{}_{\phi\phi} = -\,r(1-kr^2)\sin^2\theta,\\
&\Gamma^{\theta}{}_{r\theta} = \frac{1}{r}, && \Gamma^{\theta}{}_{\phi\phi} = -\sin\theta\cos\theta, && \Gamma^{\phi}{}_{r\phi} = \frac{1}{r},\\
&\Gamma^{\phi}{}_{\theta\phi} = \cot\theta. &&&&
\end{aligned} \tag{456}$$

As a representative check of (451), take $\Gamma^{r}{}_{\theta\theta}$:

$$\Gamma^{r}{}_{\theta\theta} = \tfrac{1}{2}g^{rr}\left(2\partial_\theta g_{r\theta} - \partial_r g_{\theta\theta}\right) = -\tfrac{1}{2}\,\frac{1-kr^2}{a^2}\,\partial_r(a^2r^2) = -\tfrac{1}{2}\,\frac{1-kr^2}{a^2}\,(2a^2r) = -\,r(1-kr^2), \tag{457}$$

where $g_{r\theta}=0$ kills the first term. The arithmetic of every other entry in (456) is identical in spirit.

E.9 Step 2 — the Ricci tensor and Ricci scalar

Statement. The Ricci tensor is the contraction of the Riemann tensor,

$$R_{\mu\nu} \;=\; \partial_\lambda\Gamma^{\lambda}{}_{\mu\nu} - \partial_\nu\Gamma^{\lambda}{}_{\mu\lambda} + \Gamma^{\lambda}{}_{\lambda\sigma}\Gamma^{\sigma}{}_{\mu\nu} - \Gamma^{\lambda}{}_{\nu\sigma}\Gamma^{\sigma}{}_{\mu\lambda}. \tag{458}$$

By homogeneity and isotropy $R_{\mu\nu}$ is diagonal and its three spatial entries are equal up to the metric, so only the tt and one spatial component are independent. We compute both.

The tt component. Set $\mu=\nu=t$ in (458). The first term $\partial_\lambda\Gamma^{\lambda}{}_{tt}$ vanishes because every $\Gamma^{\lambda}{}_{tt}=0$. The second term is $-\partial_t\Gamma^{\lambda}{}_{t\lambda}$. From (455) the trace of the connection over its lower-time slot is

$$\Gamma^{\lambda}{}_{t\lambda} = \Gamma^{r}{}_{tr} + \Gamma^{\theta}{}_{t\theta} + \Gamma^{\phi}{}_{t\phi} = 3H = 3\,\frac{\dot{a}}{a}, \tag{459}$$

so $-\partial_t\Gamma^{\lambda}{}_{t\lambda} = -3\,\dfrac{d}{dt}\left(\dfrac{\dot{a}}{a}\right) = -3\,\dfrac{\ddot{a}\,a-\dot{a}^2}{a^2} = -3\dfrac{\ddot{a}}{a}+3H^2$. The third term $\Gamma^{\lambda}{}_{\lambda\sigma}\Gamma^{\sigma}{}_{tt}$ vanishes ($\Gamma^{\sigma}{}_{tt}=0$). The fourth term is $-\Gamma^{\lambda}{}_{t\sigma}\Gamma^{\sigma}{}_{t\lambda} = -\sum_i(\Gamma^{i}{}_{ti})^2 = -3H^2$, using (454). Adding,

$$R_{tt} = \left(-3\frac{\ddot{a}}{a}+3H^2\right) + \left(-3H^2\right) = -\,3\,\frac{\ddot{a}}{a}. \tag{460}$$

A spatial component. Take $\mu=\nu=r$. Carrying out (458) term by term with the symbols of (453), (455) and (456) — the time terms supply $a\ddot{a}+2\dot{a}^2$ and the spatial terms supply the curvature $2k$, all divided by $1-kr^2$ — gives, after collecting,

$$R_{rr} = \frac{a\ddot{a}+2\dot{a}^2+2k}{1-kr^2}. \tag{461}$$

Isotropy guarantees the angular components are fixed by the same scalar; one finds $R_{\theta\theta} = r^2(a\ddot{a} + 2\dot{a}^2 + 2k)$ and $R_{\phi\phi} = R_{\theta\theta}\sin^2\theta$. In every case

$$R_{ij} = \Big(\frac{\ddot{a}}{a} + 2H^2 + \frac{2k}{a^2}\Big)g_{ij} \qquad (i,j \text{ spatial}), \tag{462}$$

which is the cleanest form to contract.

The Ricci scalar. Contract (460) and (462) with $g^{\mu\nu}$. Using $g^{tt}R_{tt} = (-1)(-3\ddot{a}/a) = 3\ddot{a}/a$ and $g^{ij}R_{ij} = 3(\ddot{a}/a + 2H^2 + 2k/a^2)$ (the factor 3 is $g^{ij}g_{ij}$ summed over the three spatial directions),

$$R = g^{tt}R_{tt} + g^{ij}R_{ij} = 3\frac{\ddot{a}}{a} + 3\Big(\frac{\ddot{a}}{a} + 2H^2 + \frac{2k}{a^2}\Big) = 6\Big(\frac{\ddot{a}}{a} + H^2 + \frac{k}{a^2}\Big). \tag{463}$$

This is the full scalar curvature of the FRW geometry; it is positive for an accelerating, spatially closed universe and can change sign in general.

E.10 Step 2 — the Einstein tensor: components

Statement. The Einstein tensor is $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R\,g_{\mu\nu}$. We assemble its two independent components from the pieces above.

The tt component (this is the one that becomes the first Friedmann equation). With $g_{tt} = -1$,

$$\begin{aligned} G_{tt} &= R_{tt} - \tfrac{1}{2}R\,g_{tt} = -3\frac{\ddot{a}}{a} - \tfrac{1}{2}\Big[6\Big(\frac{\ddot{a}}{a} + H^2 + \frac{k}{a^2}\Big)\Big](-1) \\ &= -3\frac{\ddot{a}}{a} + 3\frac{\ddot{a}}{a} + 3H^2 + 3\frac{k}{a^2} = 3\Big(H^2 + \frac{k}{a^2}\Big). \end{aligned} \tag{464}$$

The two $\ddot{a}/a$ pieces cancel exactly: this is why the first Friedmann equation is first order in time even though the Ricci tensor was second order. It is a structural fact, not a coincidence, and it is the reason the energy density (not its rate of change) sources the expansion rate.

A spatial component (this becomes the second Friedmann equation). With g_{ij} spatial and $R_{ij} = (\ddot{a}/a + 2H^2 + 2k/a^2)g_{ij}$,

$$\begin{aligned} G_{ij} &= R_{ij} - \tfrac{1}{2}R\,g_{ij} = \Big[\frac{\ddot{a}}{a} + 2H^2 + \frac{2k}{a^2} - 3\Big(\frac{\ddot{a}}{a} + H^2 + \frac{k}{a^2}\Big)\Big]g_{ij} \\ &= -\Big(2\frac{\ddot{a}}{a} + H^2 + \frac{k}{a^2}\Big)g_{ij}. \end{aligned} \tag{465}$$

It is often more transparent in mixed (one-up, one-down) form, where the metric factor disappears:

$$G^t{}_t = g^{tt}G_{tt} = -3\Big(H^2 + \frac{k}{a^2}\Big), \qquad G^i{}_j = -\Big(2\frac{\ddot{a}}{a} + H^2 + \frac{k}{a^2}\Big)\,\delta^i{}_j. \tag{466}$$

These two scalars — $G^t{}_t$ and the common value of $G^i{}_j$ — carry the entire gravitational content of the homogeneous, isotropic geometry.

E.11 Step 2 — the perfect-fluid stress-energy tensor

Statement. Homogeneity and isotropy force the source to be a perfect fluid: no shear, no heat flux, no anisotropic stress. In terms of the fluid four-velocity u^μ (with $u^\mu u_\mu = -1$), the energy density ρ, and the isotropic pressure p,

$$T_{\mu\nu} = \Big(\rho + \frac{p}{c^2}\Big) u_\mu u_\nu + p\, g_{\mu\nu}, \qquad \text{(c=1)}: \quad T_{\mu\nu} = (\rho + p)\, u_\mu u_\nu + p\, g_{\mu\nu}. \tag{467}$$

In the comoving frame the fluid is at rest, $u^\mu = (1,0,0,0)$ and $u_\mu = g_{\mu\nu}u^\nu = (-1,0,0,0)$. Then $u_t u_t = 1$ and the components are

$$T_{tt} = (\rho + p)(1) + p(-1) = \rho, \qquad\qquad T_{ij} = (\rho + p)(0) + p\, g_{ij} = p\, g_{ij}, \tag{468}$$

or in mixed form, the textbook diagonal

$$T^\mu{}_\nu = \mathrm{diag}(-\rho,\ p,\ p,\ p). \tag{469}$$

The sign pattern is fixed once the signature is fixed: with $(-,+,+,+)$ the energy density sits in the time-time slot with a minus in mixed form. Each cosmic component obeys a linear equation of state $p_i = w_i\, \rho_i$, with $w = 0$ for pressureless matter (dust), $w = 1/3$ for radiation, and $w = -1$ for a cosmological constant / vacuum energy.

E.12 Step 2 — the two Friedmann equations, the acceleration equation, and the continuity equation

Statement. We now feed the geometry of §E.10 and the source of §E.11 into the Einstein equation that Step 1 established on the boundary, Eq. (444) — from here on the walk is the textbook's own, because Step 1's point is that the textbook was right. In geometrized-but-explicit form (restoring c and Newton's G, and allowing the cosmological constant Λ),

$$G_{\mu\nu} + \Lambda\, g_{\mu\nu} = \frac{8\pi G}{c^4}\, T_{\mu\nu}. \tag{470}$$

We treat Λ as a geometric term on the left; equivalently it can be moved to the right as a $w = -1$ vacuum fluid with $\rho_\Lambda = \Lambda c^2/(8\pi G)$. Both conventions are used below and we keep them straight.

First Friedmann equation (the tt equation). Insert $G_{tt} = 3(H^2 + k/a^2)$ from (464), $g_{tt} = -c^2$ (restoring c), and $T_{tt} = \rho c^2$ (the energy density times c^2 in the time-time slot with these units) into (470):

$$3\Big(H^2 + \frac{kc^2}{a^2}\Big) - \Lambda c^2 = \frac{8\pi G}{c^4}\, \rho\, c^4 = 8\pi G\, \rho. \tag{471}$$

Dividing by 3 and rearranging gives the boxed result,

$$\boxed{H^2 \;=\; \Big(\frac{\dot a}{a}\Big)^2 \;=\; \frac{8\pi G}{3}\,\rho \;+\; \frac{\Lambda c^2}{3} \;-\; \frac{kc^2}{a^2}.} \tag{472}$$

This is the *first Friedmann equation*. It says the square of the expansion rate is set by the total energy density, the cosmological constant, and the spatial curvature. Folding Λ into a vacuum density ρ_Λ and letting ρ denote the sum over all components, it reads compactly $H^2 = (8\pi G/3)\rho - kc^2/a^2$.

Second Friedmann equation (the spatial equation). Insert $G^i{}_j = -(2\ddot a/a + H^2 + kc^2/a^2)\delta^i{}_j$ from (466) (with c restored on the curvature term), $g^i{}_j = \delta^i{}_j$, and $T^i{}_j = p\,\delta^i{}_j$ into the mixed-index Einstein equation $G^i{}_j + \Lambda\delta^i{}_j = (8\pi G/c^4)T^i{}_j$:

$$-\Big(2\frac{\ddot a}{a} + H^2 + \frac{kc^2}{a^2}\Big) + \Lambda c^2 = \frac{8\pi G}{c^4}\,p. \tag{473}$$

Now substitute the first Friedmann equation (472) for the combination $H^2 + kc^2/a^2 = (8\pi G/3)\rho + \Lambda c^2/3$, which removes H^2 and k from (473):

$$-2\frac{\ddot a}{a} - \frac{8\pi G}{3}\rho - \frac{\Lambda c^2}{3} + \Lambda c^2 = \frac{8\pi G}{c^2}\,p, \tag{474}$$

where on the right we wrote $T^i{}_j = p$ with $8\pi G/c^4$ so that pressure enters as $8\pi G\,p/c^2$. Solving for $\ddot a/a$ gives the *acceleration equation*,

$$\boxed{\frac{\ddot a}{a} \;=\; -\frac{4\pi G}{3}\Big(\rho + \frac{3p}{c^2}\Big) \;+\; \frac{\Lambda c^2}{3}.} \tag{475}$$

This is the *second Friedmann equation*. The combination $\rho + 3p/c^2$ is the active gravitational mass density of general relativity: pressure gravitates, and gravitates three times as hard per unit energy. For ordinary matter and radiation ($p \geq 0$) the first term decelerates the expansion; only a sufficiently negative pressure ($w < -1/3$, e.g. a cosmological constant) can drive $\ddot a > 0$ and accelerate it.

The continuity equation and the redundancy. The third standard equation, conservation of energy, is $\nabla_\mu T^\mu{}_0 = 0$. Working it out for the perfect fluid with the FRW connection, the time component of $\nabla_\mu T^{\mu\nu} = 0$ is

$$\partial_t T^t{}_t + \Gamma^\mu{}_{\mu t} T^t{}_t - \Gamma^\lambda{}_{\mu t} T^\mu{}_\lambda = -\dot\rho - 3H\rho - 3Hp \;=\; 0, \tag{476}$$

using $\Gamma^\mu{}_{\mu t} = 3H$ from (459) and $T^t{}_t = -\rho,\; T^i{}_i = p$. This gives the *continuity equation*,

$$\boxed{\dot\rho + 3H\Big(\rho + \frac{p}{c^2}\Big) \;=\; 0.} \tag{477}$$

The two terms have a clean reading: $3H\rho$ is the dilution of energy density by the volume growth $\propto a^3$, and $3Hp/c^2$ is the pdV work done by the fluid against expansion. For a single component

with $p = w\rho c^2$ it integrates to the *scaling law*

$$\rho_i(a) = \rho_{i,0}\, a^{-3(1+w_i)}, \tag{478}$$

giving $\rho_m \propto a^{-3}$ (matter dilutes as volume), $\rho_r \propto a^{-4}$ (radiation dilutes as volume *and* redshifts), and $\rho_\Lambda = \text{const}$ (vacuum energy does not dilute).

The three are not independent. There are three boxed equations — (472), (475), (477) — but only two are independent. This is the cosmological face of the contracted Bianchi identity $\nabla_\mu G^{\mu\nu} \equiv 0$, which (with (470)) forces $\nabla_\mu T^{\mu\nu} = 0$. Concretely, differentiate the first Friedmann equation (472) (with Λ absorbed into ρ) with respect to time:

$$2H\dot{H} = \frac{8\pi G}{3}\dot{\rho} + \frac{2kc^2\dot{a}}{a^3}, \qquad \dot{H} = \frac{\ddot{a}}{a} - H^2. \tag{479}$$

Substituting $\dot{\rho} = -3H(\rho + p/c^2)$ from the continuity equation and simplifying reproduces the acceleration equation (475) exactly — no new information is generated. Operationally: pick any two of {first Friedmann, acceleration, continuity} and the third follows. The standard practice is to solve the first Friedmann equation together with the per-component scaling laws (478), which already encode continuity.

E.13 Step 2 — spatial curvature, the critical density, and the density parameters

Statement. The first Friedmann equation defines a natural reference density: the density for which the universe is spatially flat ($k = 0$) at a given epoch. Setting $k = 0$ and $\Lambda = 0$ in (472) and solving for ρ gives the *critical density*

$$\rho_{\text{crit}}(t) \;\equiv\; \frac{3H(t)^2}{8\pi G}, \qquad \rho_{\text{crit},0} = \frac{3H_0^2}{8\pi G}\ \text{(today)}. \tag{480}$$

The *density parameter* of component i is its density in units of the critical density,

$$\Omega_i(t) \equiv \frac{\rho_i(t)}{\rho_{\text{crit}}(t)}, \qquad \Omega_\Lambda \equiv \frac{\Lambda c^2}{3H_0^2}, \qquad \Omega_k \equiv -\frac{kc^2}{a^2H_0^2}. \tag{481}$$

With the matter parameter Ω_m, the vacuum parameter Ω_Λ, the radiation parameter Ω_r, and the curvature parameter Ω_k defined this way, divide the first Friedmann equation (472) through by H_0^2 at the present epoch ($a = 1$). Every term becomes a density parameter and the equation collapses to the *closure relation*

$$\boxed{\Omega_m + \Omega_r + \Omega_\Lambda + \Omega_k \;=\; 1.} \tag{482}$$

A spatially flat universe is exactly $\Omega_k = 0$, i.e. the total density equals the critical density. Writing the time dependence explicitly via the scaling laws (478), the first Friedmann equation becomes the workhorse expansion-history formula

$$\frac{H(a)^2}{H_0^2} \;=\; \Omega_r\, a^{-4} + \Omega_m\, a^{-3} + \Omega_k\, a^{-2} + \Omega_\Lambda, \tag{483}$$

often written $H(a) = H_0\,E(a)$ with $E(a)$ the square root of the right-hand side. Equation (483) is the single relation from which ages, distances, and growth rates are computed.

E.14 Step 2 — worked numerical example: the age of the universe and the critical density in flat ΛCDM

To make the abstract relations concrete we plug in present-day flat-ΛCDM values and compute two numbers a cosmologist cares about: the present critical density and the age of the universe. We use the Planck-anchored set

$$H_0 = 67.4\ \mathrm{km\,s^{-1}\,Mpc^{-1}}, \qquad \Omega_m = 0.315, \qquad \Omega_\Lambda = 0.685, \qquad \Omega_r, \Omega_k \approx 0, \tag{484}$$

which satisfy the closure relation (482) to the quoted precision ($0.315 + 0.685 = 1.000$).

Step 1: H_0 in SI and the Hubble time. Convert H_0 from astronomers' units using $1\ \mathrm{Mpc} = 3.0857 \times 10^{19}$ km:

$$H_0 = \frac{67.4\ \mathrm{km\,s^{-1}}}{3.0857 \times 10^{19}\ \mathrm{km}} = 2.184 \times 10^{-18}\ \mathrm{s^{-1}}. \tag{485}$$

The Hubble time is its reciprocal,

$$t_H \equiv \frac{1}{H_0} = 4.579 \times 10^{17}\ \mathrm{s} = 14.51\ \mathrm{Gyr}, \tag{486}$$

using $1\ \mathrm{Gyr} = 3.156 \times 10^{16}$ s. The Hubble time is an *over*estimate of the true age, because the universe expanded more slowly in the matter-dominated past.

Step 2: the present critical density. From (480) with $G = 6.674 \times 10^{-11}\ \mathrm{m^3\,kg^{-1}\,s^{-2}}$,

$$\rho_{\mathrm{crit},0} = \frac{3H_0^2}{8\pi G} = \frac{3\,(2.184 \times 10^{-18})^2}{8\pi\,(6.674 \times 10^{-11})} = 8.53 \times 10^{-27}\ \mathrm{kg\,m^{-3}}. \tag{487}$$

That is breathtakingly dilute: dividing by the proton mass $m_p = 1.673 \times 10^{-27}$ kg gives

$$\frac{\rho_{\mathrm{crit},0}}{m_p} = 5.10\ \mathrm{protons\,m^{-3}} \approx 5\ \text{hydrogen atoms per cubic meter}. \tag{488}$$

The *actual* mean matter density is Ω_m times this, about 1.6 proton masses per cubic meter, of which only one part in six is ordinary baryonic matter.

Step 3: the age of the universe. The age is the comoving-time integral of $da/\dot{a}$ from the Big Bang ($a = 0$) to today ($a = 1$). Using $\dot{a} = aH(a) = aH_0E(a)$ with $E(a)$ from (483) and $\Omega_r = \Omega_k = 0$,

$$t_0 = \int_0^1 \frac{da}{a\,H_0\,E(a)} = \frac{1}{H_0}\int_0^1 \frac{da}{a\,\sqrt{\Omega_m\,a^{-3} + \Omega_\Lambda}} = \frac{1}{H_0}\int_0^1 \frac{\sqrt{a}\,da}{\sqrt{\Omega_m + \Omega_\Lambda\,a^3}}. \tag{489}$$

For a flat two-component (matter-plus-Λ) universe this integral has a closed form,

$$t_0 = \frac{2}{3H_0\sqrt{\Omega_\Lambda}}\,\operatorname{arcsinh}\sqrt{\frac{\Omega_\Lambda}{\Omega_m}}. \tag{490}$$

Evaluate the dimensionless factor with the inputs (484): $\sqrt{\Omega_\Lambda/\Omega_m} = \sqrt{0.685/0.315} = \sqrt{2.1746} = 1.4747$, and $\operatorname{arcsinh}(1.4747) = \ln(1.4747+\sqrt{1+1.4747^2}) = \ln(1.4747+1.7822) = \ln(3.2569) = 1.1808$. Then

$$\frac{2}{3\sqrt{\Omega_\Lambda}}\,\operatorname{arcsinh}\sqrt{\frac{\Omega_\Lambda}{\Omega_m}} = \frac{2}{3\sqrt{0.685}}\,(1.1808) = \frac{2}{3\,(0.8276)}\,(1.1808) = 0.9510, \tag{491}$$

so the age is 0.9510 Hubble times,

$$t_0 = 0.9510 \times t_H = 0.9510 \times 14.51\ \text{Gyr} = 13.80\ \text{Gyr}. \tag{492}$$

A direct numerical quadrature of the integral (489) reproduces the same value $t_0 = 13.80$ Gyr to four significant figures, confirming the closed form. This 13.80 Gyr is the standard ΛCDM age, and the way it sits a little below the Hubble time 14.51 Gyr is the quantitative fingerprint of an early matter-dominated, decelerating phase giving way to late-time Λ-driven acceleration.

E.15 Step 3 — the source ledger: the components from geometry

Steps 1 and 2 built the machine; Step 3 loads it. In standard cosmology this is the moment the fitting begins: one writes $\rho = \sum_i \rho_i$ with whatever components the data demand and dials their present-day fractions. Here the loading is a reading, not a fitting.

The source becomes a ledger (paying debt 3). With the first three inputs settled, the right-hand side — the cosmic inventory — is where cosmology proper begins, and it is the subject of the main text. The free menu $\{\Omega_i, w_i\}$ of standard cosmology is replaced by a single *constrained ledger*: each component is a lawful boundary export of the one parent six-dimensional state, so the inventory is fixed by the geometry and the OGN-5 ledger rather than chosen to fit. This is the sense in which the headline numbers of GG-A10-Cosmos — the π-locked fractions

$$\Omega_m = \frac{1}{\pi} \approx 0.318, \qquad \Omega_\Lambda = 1 - \frac{1}{\pi} \approx 0.682, \qquad \Omega_b = \frac{1}{2\pi^2} \approx 0.051 \tag{493}$$

— are statements about the *source* of the very same Friedmann equations derived above, not modifications of them. The continuity equation being a Ward identity (Step 1) guarantees the ledger is self-consistent. These fractions are derived in Part X, and audited there honestly: at present the inventory is a *conditional* theorem, resting on a named closure conjecture rather than a finished proof. That honesty is part of the claim, not a hedge against it.

The components, one by one. A first-year student should see the cast list with its provenance and its effect on the expansion, side by side. *Radiation and neutrinos* ($w = \frac{1}{3}$, diluting as a^{-4}): inherited with the Standard-Model sector of Arisaka (2026, A9-SM), with the effective neutrino

count $N_{\rm eff} = 3.044$, $\Delta N_{\rm eff} \leq 0.03$ — a derived, testable null. *Baryons* ($w = 0$, diluting as a^{-3}): produced by the inherited leptogenesis ledger of Part V, giving the asymmetry $\eta_B \approx 5 \times 10^{-10}$ from the same integer $K = 35$ that fixed Newton's constant's throat — register **Derived (cond.)**, with an order-of-magnitude transport band carried honestly and the time-resolved integration landing a factor of about ten below the ledger road's 6.18×10^{-10} (§I.6). *Cold dark matter* ($w = 0$): the Kaxion of the compact angle — the one light field the geometry insisted on hosting on S^1_θ — with mass 8.1 μeV and the misalignment relic abundance worked in Appendix M; register **Derived (cond.)** on the compact-angle-mass-uniqueness lemma. *Dark energy* ($w = -1$, diluting not at all): the terminus residue of the descent, its scale the lock-scale seesaw

$$m_\Lambda \;=\; \frac{M^2_{\rm lock}}{M_{\rm Pl,red}} \;\approx\; 6.1 \text{ meV}, \tag{494}$$

which collapses the notorious 10^{120} vacuum-energy discrepancy to a residual factor of a few (Problem E.15.1 computes both sides), with $w = -1$ exact — no quintessence dial, no evolving equation of state. And, as decisively, *nothing else*: the export contains no fifth component, no extra relativistic species beyond the stated offset, and no spatial curvature — the absences are part of the prediction. *(Physics intuition: standard cosmology's menu is a restaurant where the universe may order anything; the ledger is a fixed dinner. The test of a fixed dinner is not whether it can match your appetite — it is that the courses were printed before you sat down.)*

E.16 Step 3 — the fractions, and the epoch identity

How much of each? This is where standard cosmology fits and GG-Theory computes. The closed-packet topology of the boundary stratum — the derivation is Part X, walked at course pace in Appendix L — forces the fractions

$$\Omega_m \;=\; \frac{1}{\pi} \;=\; 0.3183, \qquad \Omega_b \;=\; \frac{1}{2\pi^2} \;=\; 0.0507, \qquad \Omega_\Lambda \;=\; 1 - \frac{1}{\pi} \;=\; 0.6817, \tag{495}$$

with the dark-matter slice the difference, $\Omega_{\rm dm} = 1/\pi - 1/(2\pi^2) = 0.2676$. Against the measured values (0.315 ± 0.007, 0.0493 ± 0.0006, 0.685 ± 0.007, 0.2657 ± 0.0070) the raw deviations sit at the percent level — Problem E.16.1 audits the baryon entry, the sharpest of the set. The honest register, stated exactly as Part IX grades it: these fractions are **Derived (cond.)** on Conjecture VIII (closed-packet integration), the load-bearing conditional of the cosmological codebook, which closes as the *epoch identity* — $\Omega_m = 1/\pi$ exact at $t_0 - 75$ Myr under Planck central values, i.e. today to within half a percent of cosmic history — with residual lemma VIII-a exact to 10^{-15}. One declared calibration input, used once, printed in the open. *(Physics intuition: the pie chart of the universe, the most famous figure in cosmology, here stops being a measurement summary and becomes a statement about the topology of a closed packet — with the "why now" coincidence dissolved into the statement that the crossing epoch is, on cosmic scales, now; Problem E.13.2 locates it.)*

E.17 Step 3 — the initial conditions (paying debt 4)

The last debt is the starting state: the textbook imposes near-perfect smoothness, spatial flatness, and a nearly scale-invariant spectrum of small fluctuations as boundary conditions. In the descent they are outputs of the launch. Flatness: the derived plateau inflation of Part IV (worked at course pace in Appendix G) drives $|\Omega_{\rm tot} - 1|$ toward zero through $N^* = 52.6$ e-folds — and Problem E.13.3 shows just how badly a decelerating universe needs that mechanism, tracing today's near-flatness back to a 10^{-16} tuning at nucleosynthesis if nothing supplies it. Smoothness and the horizon problem: the shrinking comoving Hubble radius of the launch (Figure 38) puts the whole observable sky in early causal contact. The spectrum: tilt $n_s = 0.9634$ (**Derived (cond.)**, plateau class), tensor ratio $r = 0.0038$ (**Derived-conditional** — the decisive decade-scale prediction), and the amplitude A_s — the launch sector's one declared input, flagged as such wherever it appears. Debt 4 is paid at stated grades, with the single IOU printed. *(Physics intuition: initial conditions are where cosmology has always hidden its miracles; the descent replaces the miracle with a mechanism and then — the unusual part — itemizes the one number it still owes.)*

E.18 Step 3 — the assembled expansion law

Putting the three steps together, the expansion law of the universe can now be written the way no textbook can write it:

$$\boxed{\frac{H^2(a)}{H_0^2} \;=\; \Omega_r\, a^{-4} \;+\; \frac{1}{\pi}\, a^{-3} \;+\; \Big(1 - \frac{1}{\pi}\Big), \qquad k = 0,} \tag{496}$$

with the matter slice split internally as $\Omega_b = 1/(2\pi^2)$ plus $\Omega_{\rm dm} = 1/\pi - 1/(2\pi^2)$, the radiation term carried at its thermal value with $N_{\rm eff} = 3.044$ ($\Delta N_{\rm eff} \le 0.03$), dark energy at $w = -1$ exactly, and the Hubble constant not fitted but derived: $H_0 = 66.6$ (two-road reading; Route I 66.8) for the background, with the CMB reading 67.4 as the coherent π-offset of the sound-horizon dictionary (Part VI). The perturbation sector rides on top with $n_s = 0.9634$, $r = 0.0038$, and the declared A_s; the optical depth τ is late-time astrophysics, outside the fundamental claim and carried at its measured value. Every symbol in Eq. (496) therefore has one of three statuses, and the sentence bears repeating slowly: *there is no fitted continuous parameter in the display — one input is declared (A_s), one identity is epoch-calibrated and printed (Conjecture VIII), and everything else is geometry.* The arithmetic check closes the loop: running Eq. (496) through the Step-2 machinery gives the age of the universe as $t_0 = 0.948/H_0$ — 13.88 Gyr at the derived 66.8, 13.92 at the two-road 66.6, 13.75 at the CMB comparator 67.4 (Problem E.18.1) — against the measured 13.79 ± 0.02 Gyr, with the critical density's five protons per cubic meter unchanged from Step 2's worked example. ΛCDM is not modified; it is *concluded. (Physics intuition: the six-parameter model was never six freedoms — it was six shadows of one geometry, and the display above is what the concordance model looks like when the shadows are traced back to the object that casts them.)*

E.19 What this appendix establishes

Three steps, one arc. Step 1 derived the left-hand side of cosmology: four-dimensional gravity as the zero-mode shadow of the six-dimensional geometry — the Einstein–Hilbert term descending with its scale weight, Newton's constant computed as an elementary overlap ($M_{\rm Pl}^2 = \pi M_6^4(1 - e^{-2K})$), the boundary Einstein equation obtained by variation and projection, conservation inherited as a Ward identity, and the open-system correction bounded at 10^{-91} today with its falsifiers named. Step 2 then ran the textbook derivation of the Friedmann equations, unchanged and unskipped, on Planck-anchored inputs — fixing the present critical density at $\rho_{\rm crit,0} = 8.53 \times 10^{-27}$ kg m^{-3} (≈ 5.1 proton masses per cubic meter) and the age of the universe at $t_0 = 13.80$ Gyr: standard cosmology reproduces the measured universe, exactly as it should, and Step 1 is why it may. Step 3 replaced the fitted menu with the geometric ledger — components with their provenances, the π-locked fractions at their stated conditional register, the initial conditions from the launch — and assembled the expansion law (496) with no fitted continuous parameter. The registers are the claim: the Einstein equation and the Friedmann system **Derived**; the inventory **Derived (cond.)** on the named closure conjecture of Part X, with the epoch identity printed; the amplitude A_s declared; τ outside the claim. The rigorous proofs consumed here live in the companions — the boundary general relativity of Arisaka (2026, A5-GDOS), the bulk of Arisaka (2026, A2-GG6), the earned action of Arisaka (2026, A9-SM) — and this appendix's task has been to show, at textbook depth, exactly how their theorems become the equations a first-year student already trusts. In the lineage of Newton, who turned Kepler's measured ellipses into consequences of a single law, the concordance model is here reached as a result rather than assumed as a premise.

Problems and worked solutions. E.4.1. Do the scale-weight bookkeeping for the gauge sector, and explain from it why gravity is UV-localized while gauge forces are not.
Solution. Step 1: the measure. From the warped metric, $\sqrt{-G} = e^{-4r}\sqrt{-g_4}$: each of the four spacetime legs contributes one factor e^{-r} to the volume element. Step 2: the weights. The gauge kinetic term $\mathcal{F}_{MN}\mathcal{F}^{MN}$ needs *two* inverse metrics to raise its four-dimensional indices, each contributing e^{+2r}, so its radial weight is $e^{-4r} \cdot e^{+4r} = 1$: the gauge action is *uniform* in r, and gauge zero modes spread across the whole throat. The Einstein–Hilbert term R_4 needs only *one* inverse-metric factor, e^{+2r}, so its weight is $e^{-4r} \cdot e^{+2r} = e^{-2r}$: the gravitational action is exponentially concentrated at small r, the ultraviolet edge. Step 3: the consequence. Newton's constant is an overlap dominated by the first stretch of throat (Problem E.5.2 quantifies this), while gauge couplings sample the full depth — which is why, in the companion papers, the gauge couplings are seam registrations of the whole descent while gravity's strength is a UV-edge quantity, and why gravity is so feeble at the boundary: the infrared world lives far from where gravity keeps its strength. Physics intuition: one factor of e^{2r} — the difference between raising two indices and raising one — is the entire explanation for why you can lift a paper clip with a toy magnet against the pull of a whole planet.

E.5.1. Invert the overlap relation for the bulk scale, and show the flat-space correction is thirty orders of magnitude beyond any measurement.

Solution. Step 1: invert $M_{\rm Pl}^2 = \pi M_6^4(1-e^{-2K})$. Since $e^{-2K} = e^{-70}$ is negligible (Step 2), $M_6 = (M_{\rm Pl}^2/\pi)^{1/4}$; with the reduced Planck mass $M_{\rm Pl,red} = 2.435\times10^{18}$ GeV this gives $M_6 = (5.93\times10^{36}/\pi)^{1/4} = 1.17\times10^9$ GeV $\approx 1.2\times10^9$ GeV — the single dimensionful scale of the bulk, a million TeV, fixed by gravity's observed strength and the pure number π. Step 2: the residual. $e^{-70} = 3.98\times10^{-31}$, so truncating $(1-e^{-2K})\to 1$ errs at one part in 10^{30} — some twenty orders of magnitude below the precision of the best measured constant in physics. Step 3: the reading. The ledger depth $K = 35$ enters Newton's constant only through this invisible residual: G cares about the *existence* and shape of the throat, not its depth, while the particle-physics scales of the ladder (§E.5) care about the depth exponentially. One geometry, two very different sensitivities. Physics intuition: the throat is a well 35 units deep, but gravity's strength is set at the rim — you can measure the rim's circumference forever without learning how deep the well goes; for that you need the ladder.

E.5.2. Show that most of the Planck mass lives in the first e-fold of the throat.
Solution. Step 1: the share of the overlap integral accumulated by radial depth $r_\star$ is $f(r_\star) = \int_0^{r_\star} e^{-2r}dr / \int_0^K e^{-2r}dr = (1-e^{-2r_\star})/(1-e^{-2K})$. Step 2: evaluate at one e-fold, $r_\star = 1$: $f(1) = (1-e^{-2})/(1-e^{-70}) = 0.8647$ — 86.5% of $M_{\rm Pl}^2$ accumulates in the first of thirty-five units of depth; by $r_\star = 3$ the share is $1-e^{-6} = 99.75\%$. Step 3: the reading. Four-dimensional gravity is, quantitatively, a phenomenon of the ultraviolet edge: the deep throat — where the Standard Model's hierarchies are manufactured — contributes almost nothing to G. This is the arithmetic behind the statement that the graviton zero mode is UV-localized, and it is why no late-time cosmological measurement can see the throat directly: gravity's strength was fixed within one e-fold of the rim. Physics intuition: G is a rim phenomenon and the hierarchy is a floor phenomenon — the same well, read at opposite ends, which is how one geometry can hold both the weakest force and the widest hierarchy without tension.

E.6.1. In the geometric derivation, why is obtaining ONE Friedmann equation enough?
Solution. Step 1: Problem E.12.1 shows the system (Friedmann 1, Friedmann 2, continuity) carries two independent equations. Step 2: the boundary law delivers Friedmann 1 (the Hamiltonian constraint, in relativist's language) and the ledger's source accounting delivers continuity — the open-system bookkeeping of the GDOS descent is precisely $\dot\rho = -3H(\rho+p)$ read as a channel statement. Step 3: Friedmann 2 then follows by the algebra of Problem E.12.1 run in reverse: differentiate and substitute. Nothing further needs deriving, and the Ward-identity paragraph of §E.6 is the statement that this is a feature of the constraint structure, not a gap. Physics intuition: in both languages — Bianchi identity or ledger conservation — the same moral: a theory whose geometry owns its bookkeeping only ever owes you one dynamical equation per symmetry class.

E.12.1. Derive the continuity equation from the two Friedmann equations alone.
Solution. Step 1: differentiate the first Friedmann equation, $H^2 = \frac{8\pi G}{3}\rho$ (flat case for clarity): $2H\dot H = \frac{8\pi G}{3}\dot\rho$. Step 2: the acceleration equation gives $\dot H = \ddot a/a - H^2 = -4\pi G(\rho+p)$. Step 3: substitute: $2H\cdot[-4\pi G(\rho+p)] = \frac{8\pi G}{3}\dot\rho$, i.e. $\dot\rho = -3H(\rho+p)$ — the continuity equation, exactly. The three equations carry only two independent pieces of content. Physics intuition: this redundancy is the contracted Bianchi identity $\nabla_\mu G^{\mu\nu} = 0$ wearing FRW clothes — geometry conserves its source by construction, which is why the chapter's geometric Step could derive one equation and inherit

the other free of charge.

E.12.2. Obtain the first Friedmann equation from Newtonian physics, and identify exactly where the derivation cheats.

Solution. Step 1: place a test mass m on the surface of a comoving ball of radius $R = a(t)\chi$ and uniform density ρ; Newtonian energy conservation gives $\frac{1}{2}\dot{R}^2 - \frac{GM}{R} = E$ with $M = \frac{4\pi}{3}\rho R^3$. Step 2: divide by $R^2/2$: $(\dot{a}/a)^2 = \frac{8\pi G}{3}\rho + \frac{2E}{\chi^2 a^2}$ — the Friedmann equation, with $2E/\chi^2$ playing the role of $-kc^2$. Step 3: the cheats: (i) Newtonian gravity has no propagation, so the ball must be small compared with the Hubble radius for the shell theorem to apply; (ii) pressure gravitates in relativity ($\rho + 3p$ in the acceleration equation) and Newton misses it entirely — the derivation "works" only because the first Friedmann equation happens not to contain p; (iii) the constant E is put in by hand where relativity supplies spatial curvature with a fixed normalization. Physics intuition: the Newtonian shortcut is a mnemonic, not a derivation — but it explains why cosmic expansion obeys projectile logic: escape ($k < 0$), capture ($k > 0$), and the critical parabola ($k = 0$) are the three FRW geometries in disguise.

E.13.1. Compute the present deceleration parameter for the π-locked inventory.

Solution. Step 1: from the acceleration equation with matter ($p = 0$) and Λ ($p = -\rho$), $q_0 \equiv -\ddot{a}a/\dot{a}^2 = \frac{1}{2}\Omega_m - \Omega_\Lambda$. Step 2: insert the inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$: $q_0 = \frac{1}{2\pi} - (1 - \frac{1}{\pi}) = \frac{3}{2\pi} - 1 \approx -0.523$. Step 3: the sign is the 1998 discovery — the expansion accelerates — and the value is a zero-parameter prediction of the inventory, sitting well within the supernova constraints. Physics intuition: q_0 weighs matter's braking (half its density, since only $\rho + 3p$ brakes) against the vacuum's push (full density, twice over, since $\rho + 3p = -2\rho$); the π-locked split settles the contest at three-halves-over-π minus one — geometry deciding a tug of war.

E.13.2. Find the redshifts of matter–vacuum equality and of the onset of acceleration for the π-locked inventory.

Solution. Step 1: matter dilutes as $(1 + z)^3$, the vacuum not at all, so $\rho_m = \rho_\Lambda$ at $(1 + z)^3 = \Omega_\Lambda/\Omega_m = \pi - 1$: $z_{m\Lambda} = (\pi - 1)^{1/3} - 1 \approx 0.29$. Step 2: acceleration begins where $q = 0$, i.e. $\rho_m = 2\rho_\Lambda$ — one $(1+z)^3$ factor deeper: $z_{\rm acc} = \left[2(\pi - 1)\right]^{1/3} - 1 \approx 0.62$. Step 3: both are zero-parameter numbers of the inventory, and both sit where the supernova Hubble diagrams place them. Physics intuition: the universe began accelerating when it was about five-eighths of its present size — recently, in logarithmic terms, which is the "why now" coincidence the main text derives rather than laments.

E.13.3. The flatness problem, quantified: how flat did the universe have to start?

Solution. Step 1: from the first Friedmann equation with curvature, $\Omega_{\rm tot}(a) - 1 = kc^2/(aH)^2$: the deviation from flatness scales as the inverse square of the comoving Hubble rate aH. In a decelerating universe aH falls, so $|\Omega_{\rm tot} - 1|$ *grows* — flatness is unstable. Step 2: run it backwards. In the radiation era $aH \propto a^{-1}$ and in the matter era $aH \propto a^{-1/2}$, so from nucleosynthesis ($a_{\rm BBN} \simeq 2.4 \times 10^{-9}$, at $T \simeq 0.1$ MeV) to matter–radiation equality ($a_{\rm eq} \simeq 1/3400$) the deviation grows by $(a_{\rm eq}/a_{\rm BBN})^2 \simeq 1.6 \times 10^{10}$, and from equality to today by a further $1/a_{\rm eq} \simeq 3.4 \times 10^3$ — a combined factor of about 5×10^{13}. Step 3: impose the present bound $|\Omega_{\rm tot} - 1| < 0.005$: at nucleosynthesis the universe had to satisfy $|\Omega_{\rm tot} - 1| \lesssim 10^{-16}$ — and pushing back to the GUT edge makes the required tuning tens of orders sharper still. Either the universe began absurdly

fine-tuned, or an era in which aH *rose* drove it flat: exactly the shrinking comoving Hubble radius of the derived launch (§E.17 and Figure 38), which pays this debt mechanically. Physics intuition: a pencil standing on its point for fourteen billion years is not an initial condition worth believing; inflation is the hand that placed it upright just before you looked, and in the descent that hand is a theorem, not a story.

E.14.1. Derive the closed-form age of a flat matter-plus-Λ universe and evaluate it for the π-locked inventory.

Solution. Step 1: with $H^2 = H_0^2[\Omega_m a^{-3} + \Omega_\Lambda]$, the age integral $t_0 = \int_0^1 da/(aH)$ substitutes $u = a^{3/2}$ into an elementary form, giving

$$t_0 = \frac{2}{3H_0\sqrt{\Omega_\Lambda}} \operatorname{arcsinh}\sqrt{\frac{\Omega_\Lambda}{\Omega_m}}\,.$$

Step 2: insert $\Omega_m = 1/\pi$ and the Planck-calibrated $H_0 = 67.4\ \mathrm{km\,s^{-1}Mpc^{-1}}$ ($1/H_0 = 14.51$ Gyr): $t_0 = 0.948/H_0 = 13.75$ Gyr. Step 3: compare the matter-only answer $\frac{2}{3}H_0^{-1} = 9.7$ Gyr — younger than the oldest globular clusters, the classic "age crisis" that Λ resolved. Physics intuition: the vacuum stretches the timeline because early expansion was slower than a matter-only history extrapolates; the π-locked inventory lands the age at 13.7–13.8 Gyr with nothing tuned — the first sanity check any cosmology must pass, passed from integers.

E.14.2. Express the critical density in protons per cubic meter.

Solution. Step 1: $\rho_c = 3H_0^2/8\pi G$; with $H_0 = 67.4\ \mathrm{km\,s^{-1}Mpc^{-1}} = 2.18 \times 10^{-18}\,\mathrm{s^{-1}}$, $\rho_c = 8.5 \times 10^{-27}\,\mathrm{kg\,m^{-3}}$. Step 2: dividing by the proton mass, $\rho_c \approx 5.1$ protons per cubic meter. Step 3: of this, the inventory assigns baryons only $\Omega_b = 1/(2\pi^2) \approx 5\%$ — about one proton per four cubic meters; the rest is dark matter ($\Omega_m - \Omega_b$) and vacuum. Physics intuition: the entire cosmic budget, whose split the main text derives from topology, totals five hydrogen atoms per meter cubed — emptier than any laboratory vacuum ever pumped, and yet enough, integrated over the Hubble volume, to curve the universe flat.

E.15.1. Compute both sides of the cosmological-constant problem: the seesaw scale and the measured dark-energy scale.

Solution. Step 1: the geometry's number. The lock-scale seesaw (494) gives $m_\Lambda = M_{\mathrm{lock}}^2/M_{\mathrm{Pl,red}} = (3.85 \times 10^3)^2/2.435 \times 10^{18}$ GeV $= 6.1 \times 10^{-12}$ GeV $= 6.1$ meV. Step 2: the sky's number. The measured dark-energy density is $\rho_\Lambda = \Omega_\Lambda\,\rho_{\mathrm{crit}}$; with $\rho_{\mathrm{crit}} = 3H_0^2 M_{\mathrm{Pl,red}}^2$ at the CMB reading this evaluates to $\rho_\Lambda^{1/4} = 2.24$ meV. Step 3: the comparison. The two sit a factor ~ 2.7 apart — after one hundred and twenty orders of magnitude of vacuum-energy catastrophe, the seesaw lands on the observed scale to within a small factor, and that residual factor is not waved away: Part IV names it as the open drainage coefficient of the terminus (a flagged interpretation, register open), the one number of the dark-energy story still awaiting computation. Physics intuition: the cosmological-constant problem was never "why is Λ small?" — it was "what sets its scale at all?"; a TeV-squared-over-Planck seesaw answers the second question with millielectronvolts, and honest bookkeeping keeps the last factor of order unity on the table rather than under it.

E.16.1. Audit the baryon fraction: the sharpest entry of the π-locked inventory.

Solution. Step 1: the prediction. $\Omega_b = 1/(2\pi^2) = 0.050660$, with no adjustable input. Step 2: the comparison. Against the measured 0.0493 ± 0.0006 the raw deviation is $0.05066/0.0493 - 1 = +2.8\%$, a pull of about $+2.3$; in the physical combination $\Omega_b h^2$ the prediction reads $0.0507 \times (0.674)^2 = 0.0230$ against the measured 0.02237. Step 3: the register. Part IX carries this entry at its raw value with the epoch calibration reducing the deviation to $+1.69\%$ (pull $+1.4$) — the sharpest single test of the closed-packet fractions, and a named kill criterion: a future sub-percent baryon census that pins Ω_b away from $1/(2\pi^2)$ at the calibrated level falsifies the packet, full stop. Physics intuition: it is easy to admire $\Omega_m = 1/\pi$ and miss that the baryon slice is the stricter judge — two powers of π, a smaller error bar, and nowhere to hide; the inventory lives or dies at the second decimal of five percent.

E.18.1. Compute the age of the universe at the derived rate and at the CMB comparator. *Solution.* Step 1: the closed form of Problem E.14.1, with the π-locked inventory,

$$t_0 = \frac{2}{3H_0\sqrt{\Omega_\Lambda}} \operatorname{arcsinh}\sqrt{\Omega_\Lambda/\Omega_m} = 0.948/H_0,$$

and note the small gift of the inventory: $1 + \Omega_\Lambda/\Omega_m = 1/\Omega_m = \pi$, so the argument of the arcsinh is $\sqrt{\pi - 1}$. Step 2: the rate. At the derived $H_0 = 66.8$: $1/H_0 = 14.64$ Gyr and $t_0 = 13.88$ Gyr (at the two-road reading 66.6, 14.68 and 13.92 Gyr). At the CMB comparator $H_0 = 67.4$: $1/H_0 = 14.51$ Gyr and $t_0 = 13.75$ Gyr. Step 3: the comparison. The measured age is 13.79 ± 0.02 Gyr — 0.09 Gyr below the derived value, between it and the comparator — with no parameter adjusted anywhere in the computation: the inventory came from topology, the Hubble constant from the one background of Part VI, and the formula from Step 2 of this appendix. Physics intuition: the age of the universe is the single most public number in cosmology, and here it is a two-line consequence of π and a derived clock rate — the kind of check a first-year student can run on paper, which is exactly why it belongs at the end of the walk.

Take-away. *Standard result:* the FRW metric plus general relativity gives the two Friedmann equations, and with a measured inventory they yield the expansion history and the age of the universe. *What it assumed:* general relativity itself and its constant, spatial flatness, and an inventory of densities read from the sky — the four debts of §E.1. *What the geometry replaced it with:* the Einstein equation and Newton's constant derived from the bulk in Step 1 (conservation a Ward identity, the correction bounded, the falsifiers live); the Friedmann equations unchanged in Step 2, their provenance changed; and in Step 3 the inventory read from the ledger — components from geometry, π-locked fractions at their stated conditional register, flatness and the spectrum from the derived launch (Appendix G), the deep topology in Appendix L — closing in an expansion law with no fitted continuous parameter. *Continue in:* Part III (the stage), Appendix F (the clock), Part IX (the codebook that grades every entry).

F Cosmic expansion, scaling laws, and thermal history

What this appendix teaches. How each energy component dilutes ($\rho \propto a^{-3(1+w)}$) and how the scale factor grows in each era; redshift, entropy conservation, and the freeze-out criterion $\Gamma \sim H$ that governs every relic in the universe; the thermal milestones (equality, recombination, nucleosynthesis) computed with worked numbers; the three horizons; and the geometric reading in which the whole thermal history becomes radial transport — one logarithmic coordinate, $r = \ln(E_{\rm GUT}/\mu)$, on which every epoch has a lawful address.

Read alongside: Kolb & Turner (1990) chs. 3–5 (the classic); Ryden (2017) chs. 8–10 for a first pass; Baumann (2022) ch. 3 at course level.

The relativistic degree-of-freedom counting, entropy conservation, and freeze-out machinery used throughout this appendix are developed in the standard early-universe references (Kolb & Turner, 1990; Peebles, 1993; Weinberg, 2008).

This appendix is a self-contained, first-principles account of the standard expansion and thermal history of the universe, written for a first-year graduate reader who has met special relativity and elementary thermodynamics but not yet the Friedmann–Lemaitre machinery. Every result the main text uses — the scaling $\rho \propto a^{-3(1+w)}$, the time dependence $a(t)$ in each era, the redshift law $1+z = 1/a$, the cooling law $T \propto 1/a$ with its g_* corrections, the thermal milestones (matter–radiation equality, recombination, BBN), the freeze-out criterion $\Gamma \simeq H$, and the three cosmological horizons — is derived here line by line. The reader who wants only the central numbers can read the worked examples in §F.8 and the Status note in §F.12; the reader who wants to verify each step should follow the appendix in order.

We build this appendix in three steps, exactly as in Appendix E.

Step 1. The standard thermal history (§F.1–F.7). The scaling laws, the time dependence $a(t)$, the redshift and cooling laws, entropy conservation, the freeze-out criterion, the thermal milestones, and the three horizons — all derived line by line, with worked numbers in §F.8.

Step 2. What the standard account leaves open (§F.10). The milestones are computed by inserting *measured* densities into the equilibrium formulas; the standard treatment has no deeper account of why the switch redshifts fall where they do.

Step 3. The geometric derivation (§F.11). The cosmic clock is read as a single GDOS descent, the milestones as GRIP-Out switch events along it, and recombination as the one switch the integer ledger pins to a parameter-free value, $z_* = e^7$ — stated, as always, with both its near-hit and its nominal-tension framings.

The Friedmann equations and the equilibrium thermodynamics below are textbook physics the cohort retains intact; the geometric content is the reading of the descent that underlies them. This is the same belief that runs through the program — that the structure of the world should follow from geometry — applied now to the clock of cosmic history.

Throughout we use the metric signature $(-,+,+,+)$, set $c = 1$, write $M_{\rm Pl} = G_N^{-1/2} \approx 1.22 \times 10^{19}$ GeV for the (non-reduced) Planck mass, and take the spatially flat Friedmann–Lemaitre–Robertson–Walker (FLRW) line element with scale factor $a(t)$,

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + a^2(t)\,[\mathrm{d}r^2 + r^2\,\mathrm{d}\Omega^2], \qquad a(t_0) \equiv 1, \tag{497}$$

normalized so that the scale factor today is unity. The two Friedmann equations that govern $a(t)$ are

$$H^2 \equiv \left(\frac{\dot a}{a}\right)^2 = \frac{8\pi G_N}{3}\,\rho, \tag{498}$$

$$\frac{\ddot a}{a} = -\frac{4\pi G_N}{3}\,(\rho + 3p), \tag{499}$$

where ρ and p are the total energy density and pressure of the cosmic fluid, H is the Hubble rate, and an overdot is $\mathrm{d}/\mathrm{d}t$. We have set the spatial curvature $k = 0$ and absorbed any cosmological constant into ρ as a $w = -1$ component, consistent with the near-flat, Ω_Λ-dominated present epoch used in the main text.

F.1 The continuity equation and the scaling law $\rho \propto a^{-3(1+w)}$

The first task is to find how the energy density of a single fluid component dilutes as the universe expands. The two Friedmann equations (498)–(499) are not independent: local energy–momentum conservation $\nabla_\mu T^{\mu\nu} = 0$ ties them together. For the homogeneous, isotropic fluid the $\nu = 0$ component of that conservation law is the *continuity equation*

$$\dot\rho + 3H\,(\rho + p) = 0. \tag{500}$$

It is worth seeing this in two ways, because both readings recur in the main text.

Thermodynamic derivation. Treat a comoving volume $V \propto a^3$ as a thermodynamic system. Its energy is $E = \rho V$ and its volume is V. In the early universe expansion is adiabatic to excellent approximation (no heat flows across a comoving boundary in a homogeneous universe, since every patch is identical), so the first law $\mathrm{d}E = -p\,\mathrm{d}V$ holds with $\mathrm{d}Q = 0$. Writing $\mathrm{d}(\rho a^3) = -p\,\mathrm{d}(a^3)$ and dividing by $\mathrm{d}t$,

$$\dot\rho\,a^3 + 3\rho\,a^2\dot a = -p\,3a^2\dot a \;\Longrightarrow\; \dot\rho + 3\frac{\dot a}{a}(\rho + p) = 0, \tag{501}$$

which is exactly (500). The expansion does work against (or, for $p < 0$, is driven by) the pressure, and that work is what depletes ρ beyond the geometric a^{-3} dilution of number density.

Covariant derivation. Equivalently, differentiate (498) to get $2H\dot H = (8\pi G_N/3)\dot\rho$, use (499) to eliminate $\dot H$ via $\dot H = \ddot a/a - H^2 = -4\pi G_N(\rho + p)$, and the continuity equation drops out. The two Friedmann equations plus continuity are any two of the three; the third is a consequence.

Now specialize to a component with a constant linear (barotropic) equation of state

$$p = w\,\rho, \qquad w = \text{const.} \tag{502}$$

Substituting into (500),

$$\frac{\dot\rho}{\rho} = -3(1+w)\,\frac{\dot a}{a}, \tag{503}$$

which is separable. Integrating $\int \mathrm{d}\ln\rho = -3(1+w)\int \mathrm{d}\ln a$ gives the master scaling law

$$\boxed{\rho(a) = \rho_0\, a^{-3(1+w)}} \tag{504}$$

with ρ_0 the present-day value at $a = 1$. The three cases that organize all of standard cosmology follow by inserting the appropriate w:

- **Radiation** ($w = \frac{1}{3}$, relativistic particles or massless fields, $p = \rho/3$): $\rho_r \propto a^{-4}$. The extra power of a beyond the a^{-3} number dilution is the redshift of each quantum's energy, $E \propto a^{-1}$ (§F.3).
- **Matter / dust** ($w = 0$, nonrelativistic, $p \ll \rho$): $\rho_m \propto a^{-3}$. Pure number dilution; the rest mass does not redshift.
- **Vacuum energy / cosmological constant** ($w = -1$, $p = -\rho$): $\rho_\Lambda \propto a^0 = \text{const}$. The $(1+w) = 0$ factor switches off the dilution entirely; a true vacuum neither dilutes nor cools.

A useful intermediate is the spatial-curvature term, which in the Friedmann equation contributes like a fluid with $w = -1/3$, i.e. $\rho_k \propto a^{-2}$; this is why a curvature contribution, if present, would interpolate between matter and Λ. The reason Λ eventually dominates is now algebraically transparent: among $\{a^{-4}, a^{-3}, a^{-2}, a^0\}$ the constant term wins at large a, the steepest term wins at small a. The universe therefore passes generically through a radiation era, a matter era, and a vacuum era in that temporal order, simply because the scaling exponents are ordered.

F.2 Solving for $a(t)$: radiation, matter, and Λ domination

With (504) in hand the first Friedmann equation (498) becomes a closed ordinary differential equation for $a(t)$ whenever one component dominates. Write $\rho = \rho_0\, a^{-3(1+w)}$ and (498) reads

$$\dot a = H_0\, a^{-(1+3w)/2}, \qquad H_0 \equiv \sqrt{\tfrac{8\pi G_N}{3}\rho_0}. \tag{505}$$

Power-law eras ($w \neq -1$). Separating variables, $\int a^{(1+3w)/2}\,\mathrm{d}a \propto \int \mathrm{d}t$, gives $a^{(3+3w)/2} \propto t$, i.e.

$$\boxed{a(t) \propto t^{\,2/[3(1+w)]}} \qquad (w \neq -1). \tag{506}$$

Reading off the exponents:

- **Radiation domination** ($w = 1/3$): $a(t) \propto t^{1/2}$, so $H = \dot a/a = 1/(2t)$.

- **Matter domination** ($w = 0$): $a(t) \propto t^{2/3}$, so $H = 2/(3t)$.

Both are decelerating, $\ddot{a} < 0$, as (499) demands for $\rho + 3p > 0$. The relation $H = 1/(2t)$ in the radiation era is the one used in Eq. (27) of the main text to convert temperatures into ages.

Vacuum domination ($w = -1$). Here ρ is constant, so (498) gives $H =$ const and the solution is the de Sitter exponential

$$\boxed{a(t) \propto e^{Ht}, \qquad H = \sqrt{\tfrac{8\pi G_N}{3}\rho_\Lambda}} \tag{507}$$

which is accelerating, $\ddot{a} > 0$. This is both the late-time fate of the Ω_Λ-dominated universe and the prototype for an inflationary quasi-de Sitter stage. The transition redshifts between eras are fixed by where the scaling curves cross; for the present mixture they occur at matter–radiation equality (§F.6) and at the matter–Λ crossover near $1 + z \approx (\Omega_\Lambda/\Omega_m)^{1/3}$, which for the main-text values $\Omega_m \approx 1/\pi$, $\Omega_\Lambda \approx 1 - 1/\pi$ gives $z \approx 0.30$.

General multi-component $H(a)$. When no single component dominates the exact Hubble rate is the sum

$$H^2(a) = H_0^2\Big[\Omega_r\, a^{-4} + \Omega_m\, a^{-3} + \Omega_k\, a^{-2} + \Omega_\Lambda\Big], \tag{508}$$

where $\Omega_i \equiv \rho_{i,0}/\rho_{\mathrm{crit},0}$ and $\rho_{\mathrm{crit},0} = 3H_0^2/(8\pi G_N)$. Equation (508) is the integrand that feeds every horizon integral in §F.7; the single-component solutions above are its asymptotic limits.

F.3 Redshift $1 + z = 1/a$ and the redshifting of momenta and temperature

The scale factor is not directly observable; the redshift is. Consider a photon emitted at scale factor a_e with wavelength λ_e and observed today ($a = 1$) with wavelength λ_0. Null geodesics in the metric (497) satisfy $\mathrm{d}t = a\,\mathrm{d}r$ along the line of sight; two successive wavecrests emitted a period δt_e apart are received a period δt_0 apart, and because the comoving distance $\int \mathrm{d}t/a$ is the same for both crests one finds $\delta t_0/\delta t_e = a_0/a_e = 1/a_e$. Wavelength tracks the period, so

$$\boxed{1 + z \equiv \frac{\lambda_0}{\lambda_e} = \frac{1}{a_e}} \tag{509}$$

with $a_0 = 1$. Thus $a = 1/(1 + z)$: high redshift is small scale factor, the distant past.

Momentum redshift. The same calculation done for the de Broglie wavelength of any free particle, or equivalently from the geodesic equation for the physical momentum $|\vec{p}|$, gives the master kinematic law

$$|\vec{p}| \propto a^{-1}. \tag{510}$$

For a relativistic particle $E \simeq |\vec{p}|$, so $E \propto a^{-1}$: this is the missing fourth power in $\rho_r \propto a^{-4}$ promised in §F.1. For a nonrelativistic particle the *kinetic* energy is $|\vec{p}|^2/2m \propto a^{-2}$, while the rest energy m is fixed; this is why a decoupled nonrelativistic gas cools as $T_{\mathrm{nr}} \propto a^{-2}$, faster than the photon bath (cf. Eq. (28) of the main text).

Temperature redshift of a thermal bath. A bath of massless particles with phase-space distribution $f(p) = [\exp(p/T) \mp 1]^{-1}$ stays a thermal distribution under expansion because every momentum redshifts by the same factor a^{-1} (510). Rescaling $p \to p/a$ in f is undone exactly by rescaling $T \to T/a$. Hence a free relativistic thermal bath remains thermal with

$$\boxed{T \propto a^{-1} = 1 + z} \tag{511}$$

provided no entropy is injected. The crucial qualification — "provided no entropy is injected" — is the subject of the next subsection, where the g_* bookkeeping makes the law slightly more intricate but leaves its leading behavior intact.

F.4 Entropy conservation, $g_{*s}T^3a^3 = \text{const}$, and reheating of the photons

The cleaner and more general statement of the cooling law uses *comoving entropy conservation* rather than the naive $T \propto a^{-1}$. For a relativistic plasma in equilibrium at temperature T the energy density, pressure, and entropy density are

$$\rho_r = \frac{\pi^2}{30}\, g_*(T)\, T^4, \qquad p_r = \frac{1}{3}\rho_r, \tag{512}$$

$$s = \frac{\rho_r + p_r}{T} = \frac{2\pi^2}{45}\, g_{*s}(T)\, T^3, \tag{513}$$

matching Eqs. (21)–(22) of the main text. The two effective-degrees-of-freedom counters differ in general:

$$g_*(T) = \sum_{i\in\text{bosons}} g_i \Big(\frac{T_i}{T}\Big)^4 + \frac{7}{8}\sum_{i\in\text{fermions}} g_i \Big(\frac{T_i}{T}\Big)^4, \tag{514}$$

$$g_{*s}(T) = \sum_{i\in\text{bosons}} g_i \Big(\frac{T_i}{T}\Big)^3 + \frac{7}{8}\sum_{i\in\text{fermions}} g_i \Big(\frac{T_i}{T}\Big)^3, \tag{515}$$

where g_i is the internal multiplicity (spin times particle/ antiparticle), the factor $7/8$ is the Fermi–Dirac correction to the relativistic energy and entropy integrals, and T_i is the actual temperature of species i, which may differ from the photon temperature T once that species has decoupled. When all relativistic species share one temperature, $g_* = g_{*s}$.

The Second Law in a homogeneous universe, where adiabatic expansion injects no heat, says the total entropy in a comoving volume is conserved:

$$\boxed{S = s\, a^3 = \frac{2\pi^2}{45}\, g_{*s}(T)\, T^3 a^3 = \text{const}} \tag{516}$$

which is the integrated form of the main-text identity $\mathrm{d}[s\, a^3]/\mathrm{d}t = 0$, Eq. (23). Solving (516) for the temperature gives the exact cooling law

$$T(a) = T_0 \left[\frac{g_{*s}(T_0)}{g_{*s}(T)}\right]^{1/3} a^{-1} \quad \propto \quad g_{*s}^{-1/3}(T)\, a^{-1}, \tag{517}$$

reproducing Eq. (25). Away from a particle-annihilation threshold g_{*s} is constant and one recovers $T \propto a^{-1}$ exactly; *at* a threshold, where a species becomes nonrelativistic and annihilates, g_{*s} drops, and (517) says the still-coupled bath cools *more slowly* than a^{-1} for a while — the annihilating species dumps its entropy into the bath and *reheats* it relative to the naive expectation.

Worked g_* step: e^+e^- annihilation and the neutrino temperature. The canonical illustration, and the one the main text relies on for the neutrino background, is electron–positron annihilation around $T \sim m_e \approx 0.511$ MeV. Just above this threshold the relativistic bath is photons ($g = 2$), electrons and positrons ($g = 4$, i.e. two spin states times particle/antiparticle), and three neutrino species ($g = 2 \times 3 = 6$, one helicity each for ν and $\bar{\nu}$). The neutrinos have *already decoupled* (§F.5) at $T \sim 1$ MeV, so the e^+e^- entropy is dumped only into the photons. Counting the entropy-coupled bath that the photons share:

$$g_{*s}^{\gamma e}(\text{before}) = \underbrace{2}_{\gamma} + \frac{7}{8}\underbrace{4}_{e^\pm} = 2 + \frac{7}{2} = \frac{11}{2}, \qquad g_{*s}^{\gamma}(\text{after}) = 2. \tag{518}$$

Entropy conservation in the photon–electron sector, $g_{*s}^{\gamma e}(T_\gamma a)^3 = \text{const}$, forces the photon temperature to drop more slowly than a^{-1} across the threshold by the factor $(11/4)^{1/3}$ relative to the decoupled, freely-cooling neutrinos:

$$\boxed{\frac{T_\nu}{T_\gamma} = \left(\frac{4}{11}\right)^{1/3} \approx 0.714} \tag{519}$$

(The 4 rather than 11/2 in the ratio is because only the *net* entropy transferred enters the standard algebra, which gives 4/11 after the photon-only counting $g_{*s}^{\gamma} = 2$ is divided out of 11/2.) Today $T_\gamma = 2.725$ K, so $T_\nu \approx 1.95$ K. The same reheating makes today's value $g_{*s}(T_0) = 2 + (7/8)\cdot 6 \cdot (4/11) \approx 3.91$, while the photon-only $g_*(T_0) = 2$ — the daily reminder that the two counters genuinely differ once a species has left the photon bath. This is the single worked g_* step that the main text presumes when it speaks of photon last scattering against a colder neutrino background.

F.5 Freeze-out: the Γ/H competition and how decoupling works

Every thermal milestone — neutrino decoupling, the freeze-out of a dark-matter relic, recombination, even BBN — is at bottom a *race between two rates*: the interaction rate Γ that keeps a species in equilibrium and the expansion rate H that pulls the bath apart. The interaction rate per particle is

$$\Gamma = n \langle \sigma v \rangle, \tag{520}$$

with n the number density of targets, σ the relevant cross section, and v the relative velocity. The expansion rate in the radiation era is, from (498) and (512),

$$H(T) = \sqrt{\frac{8\pi G_N}{3}\frac{\pi^2}{30} g_* T^4} \simeq 1.66\, g_*^{1/2}\, \frac{T^2}{M_{\rm Pl}}, \tag{521}$$

matching Eq. (26). The qualitative criterion is

$$\boxed{\Gamma \gg H: \text{ coupled (equilibrium)}; \quad \Gamma \ll H: \text{ decoupled (frozen out)}.} \tag{522}$$

The crossover *freeze-out temperature* T_f is defined operationally by $\Gamma(T_f) \simeq H(T_f)$. The logic of why this is a sharp transition is that Γ and H typically scale as different powers of T: for a weak-interaction process $\Gamma \sim G_F^2 T^5$ falls faster than $H \sim T^2$ as the universe cools, so once Γ drops below H it stays below, and the abundance "freezes" at its last-equilibrium value, thereafter merely diluting as a^{-3}. For neutrinos, equating $G_F^2 T^5 \sim T^2/M_{\rm Pl}$ gives $T_f \sim 1$ MeV, which is why the neutrinos are already decoupled when e^+e^- annihilate just below 0.511 MeV — precisely the input assumed in (519).

Within GDOS this same Γ/H competition is the operational face of a branch-selection event: the system first explores microscopically equivalent histories while coupled, then crosses the $\Gamma \simeq H$ threshold, and finally stabilizes a new coarse-grained branch that the background expansion preserves. The thermodynamic and the open-system readings agree on the arithmetic; they differ only in what they say the frozen-out abundance *means*.

F.6 The thermal milestones: equality, recombination, BBN

The scaling laws and the freeze-out criterion fix the calendar of the early universe. Three milestones anchor everything the main text says about the CMB and light elements.

Matter–radiation equality, $z_{\rm eq}$. Radiation ($\propto a^{-4}$) outruns matter ($\propto a^{-3}$) toward small a, so there is a redshift where the two densities cross: $\Omega_r a^{-4} = \Omega_m a^{-3}$, giving

$$1 + z_{\rm eq} = \frac{\Omega_m}{\Omega_r}. \tag{523}$$

Before $z_{\rm eq}$ the dynamics are radiation-driven ($a \propto t^{1/2}$); after, matter-driven ($a \propto t^{2/3}$). Structure growth effectively switches on after equality, so $z_{\rm eq}$ sets the scale of the matter power spectrum turnover. A worked value appears in §F.8.

Recombination and photon decoupling, $z_{\rm rec}$. As the photon bath cools below $T \sim 0.3$ eV the reaction $p + e^- \leftrightarrow {\rm H} + \gamma$ shifts toward neutral hydrogen (the binding energy is 13.6 eV, but the enormous photon-to-baryon ratio $\eta^{-1} \sim 10^9$ delays recombination until T is far below 13.6 eV — the high-energy tail of the Planck distribution keeps hydrogen ionized until then; this is the Saha suppression). Once the free-electron fraction collapses, Thomson scattering shuts off, $\Gamma_\gamma \ll H$, and the photons free-stream: this is the surface of last scattering we observe as the CMB. The Planck-measured value is

$$z_{\rm rec} = 1089.92 \pm 0.25, \tag{524}$$

with a finite *visibility width* $\Delta z \approx 90$ over which the last-scattering probability is appreciable — the CMB is emitted from a shell of finite thickness, not a mathematical surface. That width is the physically honest yardstick against which any predicted z_* must be compared, and it is the crux of

§F.11.

Big-Bang nucleosynthesis (BBN). Earlier still, at $T \sim 0.1$–1 MeV ($t \sim 1$–130 s by Eq. (27)), the light elements form. Neutron–proton interconversion freezes out near $T \sim 0.8$ MeV, fixing $n/p \approx e^{-(m_n - m_p)/T_f} \approx 1/6$; after free-neutron decay this falls to $\approx 1/7$ by the onset of deuterium synthesis at $T \sim 0.07$ MeV, giving the famous primordial helium mass fraction $Y_p \approx 2(n/p)/[1 + (n/p)] \approx 0.25$. BBN is the earliest epoch the standard thermal history tests directly, and it depends on the same Γ/H competition and the same g_* at $T \sim$ MeV that fix the neutrino temperature. It is therefore a stringent consistency check on any modification of the expansion rate at that epoch.

F.7 The three horizons: comoving Hubble radius, particle horizon, event horizon

Causal structure in an expanding universe is governed by three distinct length scales that beginners routinely conflate. All three are integrals of $1/(a^2 H)$ or $1/(aH)$ over the history (508); the differences are in the limits of integration.

Comoving Hubble radius. The instantaneous causal scale is

$$R_H(a) \equiv \frac{1}{aH(a)}, \tag{525}$$

the comoving distance light travels in one Hubble time. In a decelerating (radiation or matter) era *aH decreases*, so *R_H grows* — scales enter the horizon. In an accelerating (inflation or Λ) era *aH increases*, so *R_H shrinks* — scales exit the horizon. The whole logic of inflationary structure formation lives in this single sign: a shrinking comoving Hubble radius is what stretches quantum fluctuations to super-horizon scales.

Particle horizon. The comoving distance from which light could have reached us since the beginning ($a = 0$) is

$$\chi_{\rm ph}(a) = \int_0^t \frac{{\rm d}t'}{a(t')} = \int_0^a \frac{{\rm d}a'}{a'^2 H(a')}. \tag{526}$$

This is the radius of the observable universe. It converges in a radiation- or matter-dominated past, so there is a genuine finite particle horizon — and the fact that causally disconnected CMB patches share a temperature is exactly the horizon problem that inflation is invoked to solve. In a radiation era $a \propto t^{1/2}$ one finds $\chi_{\rm ph} \propto a$, finite and growing.

Event horizon. The comoving distance light will *ever* reach in the entire future is

$$\chi_{\rm eh}(a) = \int_t^{t_{\rm max}} \frac{{\rm d}t'}{a(t')} = \int_a^{a_{\rm max}} \frac{{\rm d}a'}{a'^2 H(a')}. \tag{527}$$

In a matter- or radiation-only universe $t_{\rm max} = \infty$ and the integral diverges — there is no event horizon, everything is eventually visible. But in a Λ-dominated future, $a \propto e^{Ht}$ with constant H, the integral *converges*: $\chi_{\rm eh} = 1/(aH) = R_H$, a finite cosmic event horizon. The accelerating

universe therefore has a future causal boundary, the de Sitter horizon, beyond which receding galaxies pass forever out of contact. That finite event horizon is the geometric setting for the black-hole-information and holographic-bound discussions the main text develops downstream.

F.8 Worked numerical examples

We close the technical content with two explicit computations a reader can reproduce on paper.

Example 1 — matter–radiation equality $z_{\rm eq}$. Use Eq. (523), $1 + z_{\rm eq} = \Omega_m/\Omega_r$. The radiation density today is the photon plus neutrino contribution. With $T_0 = 2.725$ K the photon density parameter is $\Omega_\gamma h^2 \approx 2.47 \times 10^{-5}$; including the three neutrino species via the factor $1 + (7/8)(4/11)^{4/3} \cdot 3 \approx 1.68$ gives the total radiation $\Omega_r h^2 \approx 4.15 \times 10^{-5}$. Taking $h \approx 0.674$ (so $h^2 \approx 0.454$),

$$\Omega_r \approx \frac{4.15 \times 10^{-5}}{0.454} \approx 9.1 \times 10^{-5}. \tag{528}$$

Using the main-text matter fraction $\Omega_m = 1/\pi \approx 0.3183$,

$$1 + z_{\rm eq} = \frac{\Omega_m}{\Omega_r} \approx \frac{0.3183}{9.1 \times 10^{-5}} \approx 3.5 \times 10^3, \qquad \boxed{z_{\rm eq} \approx 3.5 \times 10^3.} \tag{529}$$

This sits in the right ballpark of the standard $z_{\rm eq} \approx 3400$; the precise value depends on the assumed $\Omega_m h^2$, but the order of magnitude — equality a few thousand in redshift, well before recombination at $z \sim 1090$ — is robust. Sanity check: $z_{\rm eq} \approx 3500 > z_{\rm rec} \approx 1090$, so the universe is matter-dominated by the time the CMB is released, as required for structure to have begun growing.

Example 2 — the temperature at recombination. Recombination occurs at $z_{\rm rec} \approx 1090$. Using the cooling law in observational form, $T(z) = T_0(1+z)$ with $T_0 = 2.725$ K and the conversion $1\ \text{K} = 8.617 \times 10^{-5}$ eV:

$$T_{\rm rec} = 2.725\ \text{K} \times (1 + 1090) = 2.725 \times 1091\ \text{K} \approx 2.97 \times 10^3\ \text{K}, \tag{530}$$

$$T_{\rm rec} \approx 2.97 \times 10^3\ \text{K} \times 8.617 \times 10^{-5}\ \frac{\text{eV}}{\text{K}} \approx 0.26\ \text{eV} \ \approx\ 0.3\ \text{eV}. \tag{531}$$

This confirms the $T \sim 0.3$ eV figure quoted in §F.6: recombination happens not at the 13.6 eV binding energy but roughly a factor of ~ 50 colder, because the huge photon-to-baryon ratio keeps the high-energy tail of the Planck spectrum ionizing hydrogen until the bath has cooled far below 13.6 eV. As a cross-check, had we instead used the GG-Theory $z_* = e^7 \approx 1096.6$ the temperature would be $T_* = 2.725\,(1 + 1096.6)$ K $\times\, 8.617 \times 10^{-5} \approx 0.258$ eV — indistinguishable from (531) at this precision, which is precisely the visibility-width statement of §F.11 expressed in temperature: the 0.6% redshift offset is a 0.6% temperature offset, far below the thermal width of last scattering.

F.9 The five clocks of cosmology: t, a, z, T, and r

Cosmology tells one history on five interchangeable axes, and fluency in translating among them is the single most useful practical skill this course can teach. Four are standard: cosmic time t, the

scale factor a, redshift $1+z=1/a$, and temperature $T\propto 1/a$. The fifth is this program's own: the radial coordinate

$$r \;=\; \ln\!\left(\frac{E_{\rm GUT}}{\mu}\right), \tag{532}$$

the logarithmic depth of the energy scale μ below the GUT edge — simultaneously the renormalization-group axis and the literal radial direction of the bulk (Part VIII develops the full dictionary). In ordinary time the early universe flashes by and the late universe crawls; on r the story is evenly paced, one unit per e-fold of energy, which is why the descent narrative of Part IV reads so naturally there.

Worked example (the student should reproduce every entry). Place recombination and the lock crossing on all five axes. *Recombination:* $T\simeq 0.26$ eV gives $1+z=T/T_0\simeq 0.26\,\text{eV}/2.35\times 10^{-4}\,\text{eV}\simeq 1100$ (the ledger's corridor-switch form is $z_{\rm sw}=e^7\simeq 1097$, the atomic decoupling 0.4% below; §42.1); $a\simeq 9\times 10^{-4}$; $t\simeq 3.7\times 10^5$ yr from the matter-era relation of §F; and $r=\ln(2.28\times 10^{16}\,\text{GeV}/0.26\,\text{eV})\simeq 58$. *The lock crossing:* $\mu\sim M_{\rm lock}=3.85$ TeV gives $r=\ln(2.28\times 10^{16}/3.85\times 10^3)\simeq 29.4$; the radiation-era clock puts $T\sim 3.85$ TeV at $t\simeq 1.6\times 10^{-14}$ s, $1+z\sim 1.6\times 10^{16}$. Two events separated by seventeen orders of magnitude in time sit twenty-nine units apart on r — the compression that makes the whole history legible on one page (the master table of Part VIII and the timeline of Figure 28).

F.10 What the standard thermal history assumes — and cannot explain

The derivations of Step 1 are exact, and together they reproduce the entire thermal history of the universe. As in Appendix E, it is worth naming what they must take as given.

1. **The inventory.** The densities $\{\Omega_i\}$ that drive the expansion and set the epoch of matter–radiation equality are measured inputs, not derived.

2. **The entropy.** The baryon-to-photon ratio $\eta\approx 6\times 10^{-10}$ — which fixes recombination and nucleosynthesis — is a measured number with no account of its origin in the standard treatment.

3. **The milestone redshifts.** Equality, recombination, and the rest are located by asking where the measured densities cross, or where a reaction rate falls below H. The procedure is exact, but it offers no reason the switches sit at *these* redshifts rather than others; there is no parameter-free structure underneath them.

Standard cosmology, in other words, is a superb *calculator* of the thermal history once the inputs are supplied — but the inputs, and the particular redshifts of the milestones, are things it measures rather than explains. The geometric account asks whether at least the switch redshifts carry a deeper, integer-ledger structure.

F.11 The geometric derivation: the cosmic clock as a GDOS descent

The thermal history of Step 1 is exact, and it reproduces every milestone of the early universe. But notice how it does so: one inserts the *measured* densities and the *measured* baryon-to-photon ratio into the equilibrium formulas and reads off where the transitions fall. The milestone redshifts

are outputs of measured inputs, not consequences of anything deeper. GG-Theory asks the further question — is there a structural, parameter-free reason the switches sit where they do? — and for the most important of them it gives an answer.

The cosmic clock is a geometric descent. The radial coordinate $r = \ln(E_{\rm GUT}/\mu)$ of the bulk is not an abstraction: it is the descent parameter of the open-system boundary law GDOS, and it is the cosmic clock. Temperature, redshift, and cosmic time are three readings of one underlying motion down r,

$$r = \ln(E_{\rm GUT}/\mu) \;\longleftrightarrow\; T \propto 1/a \;\longleftrightarrow\; 1+z = 1/a \;\longleftrightarrow\; t,$$

so the whole of Step 1 — the cooling, the redshifting, the sequence of milestones — is the single fall from the hot unified edge at $r = 0$ toward the cold present. *Physics intuition:* the universe does not carry a separate clock; its temperature *is* its position on the descent, and the arrow of time is the direction of the GRIP flow.

Milestones are GRIP-Out switch events. Along that descent the boundary state does not change smoothly forever. At particular depths it undergoes geometric symmetry-breaking switches — GSSB events, the outward (GRIP-Out) steps of the runtime — and these are the milestones: equality, recombination, the onset of dark energy. Standard cosmology locates them by asking where the densities cross; the framework asks, in addition, whether the switch depths are fixed by the integer ledger. For the recombination era the answer is yes: the corridor switch there is the one milestone the geometry pins to a parameter-free integer, with the atomic decoupling derived separately (§42.1).

Reading note (the two-quantity form). What follows predates the two-quantity resolution of §42.1, which now governs every comparison of e^7 with the measured sky: the ledger integer fixes the *corridor switch* $z_{\rm sw} = e^7$, a topological marker, while the decoupling redshift the microwave sky actually records is derived from the framework's own atomic chain, landing on the Planck value at the 0.1% level, 0.4% (in e-folds) below the switch. The both-framings computation below is retained deliberately, as the pedagogical bridge that motivated that resolution: it shows what happens when e^7 is read — as it originally was — as a direct claim on $z_{\rm rec}$, and why the two-quantity form replaces that reading. Everything that follows is to be read accordingly (Eq. (262)).

In the cohort's arithmetic ledger — the OGN-5 five-integer set $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ established upstream (Arisaka, 2026, A2-GG6; Arisaka, 2026, A3-BI6, audited in Arisaka, 2026, A4-Euclid) — the recombination switch is identified here with the single integer $J = 7$ via the exponential map

$$\boxed{z_* = e^J = e^7 \approx 1096.633} \tag{533}$$

with *no continuous adjustable parameter*: $J = 7$ is an integer from the ledger, and the exponential is the same codebook-D^* map that the main text uses for the other switch redshifts. This is a parameter-free ledger statement, not a fit. In the superseded one-quantity reading — e^7 offered directly as the recombination redshift — it had to be reported with both framings, because the two framings tell genuinely different stories; the two-quantity form of §42.1 dissolves the tension between them, and the computation is kept here as the bridge to it.

Framing 1 — the near-hit (the optimistic reading). Against the Planck central value (524), the fractional offset is

$$\frac{z_* - z_{\rm rec}}{z_{\rm rec}} = \frac{1096.63 - 1089.92}{1089.92} \approx \frac{6.71}{1089.92} \approx 0.0062 = 0.6\%. \tag{534}$$

A parameter-free integer descending to within 0.6% of a measured cosmological switch redshift is, on its face, a striking coincidence-or-not: it is the kind of structural near-hit that motivates taking the ledger seriously. And against the *physically relevant* yardstick — the recombination visibility width $\Delta z \approx 90$, the actual thickness of the last- scattering shell — the offset is

$$\frac{z_* - z_{\rm rec}}{\Delta z} \approx \frac{6.71}{90} \approx 0.075, \tag{535}$$

i.e. well under a tenth of the shell thickness. (If one instead uses the conservative half-width $\Delta z/2 \approx 45$ as the 1σ physical scale, the offset is $\approx 0.15\,\sigma$, still well under a fifth of the shell.) By the standard that asks "does the prediction land inside the surface of last scattering," the answer is unambiguously yes: e^7 lands comfortably within the last-scattering shell.

Framing 2 — the large nominal tension (the honest caveat). The optimistic reading must not be allowed to stand alone, because the Planck *measurement error* is tiny: $\sigma_{\rm Planck} = 0.25$. Measured in those units the same 6.71 offset is

$$\frac{z_* - z_{\rm rec}}{\sigma_{\rm Planck}} \approx \frac{6.71}{0.25} \approx 27\,\sigma. \tag{536}$$

Twenty-seven standard deviations is a *large nominal tension*: by the precision-cosmology standard that compares a prediction to the *measurement* error, e^7 is decisively excluded as the value of $z_{\rm rec}$. There is no way to dress 27σ up as agreement, and the present appendix does not try to. The honest statement is therefore the conjunction of both numbers: *e^7 is a 0.6% near-hit and a sub-0.2σ visibility-width near-hit, and simultaneously a $\sim 27\sigma$ tension against the Planck central value.* The reconciliation of the two framings is simply that they ask different questions. Framing 2 asks "is e^7 the precise value of the recombination redshift?" — and the answer is no. Framing 1 asks "does a parameter-free integer ledger reproduce the recombination switch to within the physical width of the last-scattering surface?" — and the answer is yes, to better than a fifth of that width.

What the claim is, and is not. The $z_* = e^7$ identification is a *structural near-hit from an independent integer ledger*, codebook code D^* (Derived-conditional): it is the projection of the upstream OGN-5 theorem onto the cosmological-boundary stratum under the closed-packet map, not a precision determination of $z_{\rm rec}$. It is not, and is never to be presented as, a match to the Planck error bar. Its evidential weight is the weight of a postdiction that costs no free parameters and happens to land inside the last-scattering shell, drawn from the same ledger that yields $\Omega_m = 1/\pi$ and the other D^* entries — no more and no less. The discipline the cohort imposes on itself here is the one this appendix enacts for the superseded reading — state the 0.6% near-hit and the 27σ nominal tension in the same breath — and, in the body, the stronger one: retire that reading

altogether in favor of the two-quantity form (§42.1).

What this appendix establishes. The cosmic clock of Step 1 is, beneath the thermodynamics, a single geometric descent: temperature and redshift are positions on it, and the milestones are the symmetry-breaking switches along the way. Of those switches, the recombination-era corridor switch is the one the integer ledger pins — $z_{\rm sw} = e^7$, a parameter-free topological marker — while the decoupling the sky records is derived from the framework's own atomic chain, on the measured value at the 0.1% level, 0.4% below the switch (§42.1). The both-framings computation above records the superseded one-quantity reading — a 0.6% near-hit inside the last-scattering shell that is simultaneously a $\sim 27\sigma$ nominal tension against the Planck central value — the very tension the two-quantity form dissolves by not reading e^7 as a claim on $z_{\rm rec}$ at all. It is a D^* marker, not a precision match, and that is exactly how it is carried in the codebook. Less than the clean theorem of Appendix E, but in the same Euclidean spirit: the piece the geometry fixes is stated as a result, and its limits are stated in the same breath.

F.12 Status

- **Standard thermal history — Derived (D).** The scaling law $\rho \propto a^{-3(1+w)}$ (504), the era solutions $a(t)$ (506)–(507), the redshift and momentum laws (509)–(511), comoving entropy conservation and $T \propto g_{*s}^{-1/3} a^{-1}$ (516)–(517) with the e^+e^- reheating ratio $T_\nu/T_\gamma = (4/11)^{1/3}$ (519), the freeze-out criterion $\Gamma \simeq H$ (522), the milestones ($z_{\rm eq}$, recombination, BBN), and the three horizons (525)–(527) are all standard cosmology, derived here from the Friedmann equations and equilibrium thermodynamics. GG-Theory retains these intact; it constrains the source terms, not the equations.

- **The $z_* = e^7$ corridor switch — Derived-conditional (D^*), read in the two-quantity form.** The recombination-era corridor-switch redshift $z_* = e^J = e^7 \approx 1096.6$ (533) descends parameter-free from the OGN-5 integer $J = 7$ (Arisaka, 2026, A4-Euclid) via the codebook-D^* closed-packet map. In the superseded one-quantity comparison retained above as the bridge it is a *structural near-hit*: 0.6% from the Planck central value (534) and under 0.2σ of the recombination visibility width (535), but simultaneously a *large nominal tension* of $\sim 27\sigma$ against the Planck measurement error (536). In the body it is carried in the two-quantity form of §42.1 — the switch a topological marker, the atomic decoupling the audited quantity — with the both-framings computation of this appendix retained as the pedagogical bridge; it is never presented as a precision match to $z_{\rm rec} = 1089.92 \pm 0.25$.

Problems and worked solutions. F.2.1. Solve the master scaling law for the three epochs and check the exponents.

Solution. Step 1: for $\rho \propto a^{-3(1+w)}$, Friedmann 1 gives $\dot{a} \propto a^{-(1+3w)/2}$, an elementary ODE with power-law solution $a \propto t^{2/3(1+w)}$ (506) for $w \neq -1$. Step 2: insert the three equations of state: radiation ($w = \frac{1}{3}$) gives $a \propto t^{1/2}$; matter ($w = 0$) gives $a \propto t^{2/3}$; vacuum ($w = -1$) breaks the power-law ansatz and integrates instead to $a \propto e^{Ht}$ (507). Step 3: check consistency: $H = \frac{2}{3(1+w)t}$,

so Ht is constant for any fluid — the age of the universe is always of order the Hubble time, with coefficient $\frac{1}{2}$ (radiation) or $\frac{2}{3}$ (matter). Physics intuition: stiffer contents brake harder — radiation's extra pressure steepens the dilution and slows the growth exponent — and the vacuum, which never dilutes, is the limiting case where braking fails entirely and expansion runs away exponentially.

F.4.1. Fill in the entropy bookkeeping behind $T_\nu/T_\gamma = (4/11)^{1/3}$.
Solution. Step 1: just above $e^\pm$ annihilation the photon–electron sector carries $g_{*s}^{\gamma e} = 2 + \frac{7}{8} \times 4 = \frac{11}{2}$ (518) (photons: two polarizations; $e^\pm$: four spin-charge states at the fermion's $\frac{7}{8}$). Step 2: below the threshold only photons remain, $g_{*s}^{\gamma} = 2$; comoving entropy conservation in that sector, $g_{*s}(T_\gamma a)^3 =$ const (516), gives $(T_\gamma a)^3_{\text{after}}/(T_\gamma a)^3_{\text{before}} = \frac{11/2}{2} = \frac{11}{4}$. Step 3: the decoupled neutrinos cool as $T_\nu \propto 1/a$ throughout, so $T_\nu/T_\gamma = (4/11)^{1/3} \approx 0.714$ (519) — today, 1.95 K against 2.725 K. Physics intuition: the annihilating pairs bequeath their entropy to the photons alone because the neutrinos have already left the room; the $(4/11)^{1/3}$ is a conservation law remembering a departure — and the cosmic neutrino background's predicted chill is among the most secure unmeasured numbers in physics.

F.5.1. Estimate the neutrino decoupling temperature from the Γ/H criterion.
Solution. Step 1: weak scattering at temperature T has $\sigma \sim G_F^2 T^2$ and target density $n \sim T^3$, so $\Gamma = n\sigma v \sim G_F^2 T^5$ (520). Step 2: the radiation-era expansion rate is $H \sim T^2/M_{\text{Pl}}$ (521). Step 3: setting $\Gamma = H$ (522): $T_{\text{dec}} \sim (G_F^2 M_{\text{Pl}})^{-1/3} \approx 0.8$–$1$ MeV. Physics intuition: the fifth power is the whole story — $\Gamma/H \propto T^3$ falls through unity so steeply that decoupling is sharp, a switch rather than a fade; every relic in the universe (neutrinos here, the Kaxion in Appendix M, helium in Appendix H) is a fossil of one such switch, and the chapter's freeze-out grammar is cosmology's universal fossilization mechanism.

F.6.1. Derive the redshift of matter–radiation equality from the π-locked matter fraction.
Solution. Step 1: the present radiation density (photons at 2.725 K plus three neutrino species at $(4/11)^{1/3}$ of the photon temperature) is $\Omega_r h^2 = 4.15 \times 10^{-5}$, giving $\Omega_r \approx 9.1 \times 10^{-5}$ (528). Step 2: radiation dilutes one power of $(1+z)$ faster than matter, so equality sits at $1 + z_{\text{eq}} = \Omega_m/\Omega_r$. Step 3: with $\Omega_m = 1/\pi$: $1 + z_{\text{eq}} \approx 0.3183/9.1 \times 10^{-5} \approx 3.5 \times 10^3$ (529) — squarely the CMB-fit value. Physics intuition: equality is the moment gravity's builder changes shift — perturbations grow only logarithmically under radiation and linearly under matter — so z_{eq} is imprinted in the shape of the matter power spectrum, where surveys read it and the inventory predicts it.

F.6.2. State the two-quantity z_* discipline precisely, and compute the gap.
Solution. Step 1: the corridor switch is topological and exact: $z_{\text{sw}} = e^J = e^7 = 1096.633$ (533), fixed by the single ledger integer $J = 7$ with no atomic physics in it. Step 2: the atomic decoupling redshift is a *different* quantity: $z_{\text{dec}} \approx 1090$, derived through the Saha/visibility chain (hydrogen's levels, η, the expansion rate). The gap is $e^7/1090 - 1 \approx 0.6\%$ — two quantities, lawfully close, never to be conflated. Step 3: the register: the falsifier on this sector is *observable-only* — it fires if the measured atomic z_{dec} moves off the atomic chain's prediction by more than 0.5%, and no measurement of z_{dec} can touch the topological z_{sw}. Physics intuition: a switch on the corridor and a fog clearing in the gas are different events that happen to sit close — the theory's discipline is to

price each with its own chain and let the sky judge only the one it can see.

F.7.1. Compute the asymptotic future event horizon of the π-locked universe.
Solution. Step 1: at late times the universe is vacuum-dominated with $H_\Lambda = H_0\sqrt{\Omega_\Lambda} = 0.826\,H_0$. Step 2: the event horizon (527) asymptotes to the de Sitter value $d_{\rm EH} = c/H_\Lambda$. Step 3: numerically, $c/H_0 = 4.45$ Gpc, so $d_{\rm EH} \to 5.39$ Gpc ≈ 17.6 billion light-years. Galaxies beyond that comoving distance today have already sent us the last photons we will ever receive. Physics intuition: acceleration turns the sky into a slowly closing archive — and the horizon it closes toward is the de Sitter port whose thermodynamics the black-hole companion works in its Appendix Q: the same π-locked vacuum fixes both the archive's radius and its temperature.

F.8.1. How much entropy does the universe carry per baryon?
Solution. Step 1: the entropy density today is dominated by photons and neutrinos: $s/n_\gamma = 7.04$ (the same 7.04 that converts Y_L to η_B in Appendix I, Eq. (628)). Step 2: with the baryon-to- photon ratio $\eta \approx 6.1 \times 10^{-10}$, the entropy per baryon is $s/n_B = 7.04/\eta \approx 1.2 \times 10^{10}$. Step 3: ten billion units of entropy per particle of matter: the universe is, by particle count, overwhelmingly a photon gas with a trace contamination of baryons. Physics intuition: this enormous number is why BBN and recombination happen "late" (the photon bath outvotes binding energies by its sheer numbers — Appendix H's bottleneck and this chapter's Saha delay are both $\ln(1/\eta)$ effects), and its origin is exactly the asymmetry η_B that Appendix I assembles from the ledger.

Take-away. *Standard result:* dilution laws, freeze-out, and the thermal milestones — the engine room of every relic calculation in cosmology. *What it assumed:* the inventory driving the expansion, and initial abundances set elsewhere. *What the geometry replaced it with:* the whole history re-read as radial transport on $r = \ln(E_{\rm GUT}/\mu)$ — every epoch a lawful address on one axis, with the descent's two crossings (the graceful exit and the lock crossing) as the events that stock the universe. *Continue in:* Part VIII (the chronology), Appendix G (the launch).

G Slow-roll inflation from first principles

What this appendix teaches. Why the early universe needed an accelerating stage (the horizon, flatness, and relic problems); the full slow-roll apparatus — the field equations, the slow-roll parameters, the e-fold count, and the quantum origin of the scalar and tensor spectra — derived from scratch with a worked example; what the standard treatment must assume (a field, a potential, a height, a count); and the geometric derivation in which the potential's class, the tilt $n_s = 0.9634$, and the tensor target $r = 0.0038$ are computed, with the amplitude the one declared input.

Read alongside: Liddle & Lyth (2000) chs. 3–5 or Baumann (2022) ch. 4 for Step 1; Mukhanov (2005) for the quantum theory of the perturbations at depth.

The slow-roll inflationary dynamics and the quantum generation of the nearly scale-invariant primordial spectrum are summarized in Figure 39. The slow-roll machinery and the primordial spectrum developed here are drawn from the founding and review literature (Bardeen *et al.*, 1983; Guth & Pi, 1982; Hawking, 1982; Linde, 1983; Lyth & Riotto, 1999; Mukhanov & Chibisov, 1981; Starobinsky, 1980).

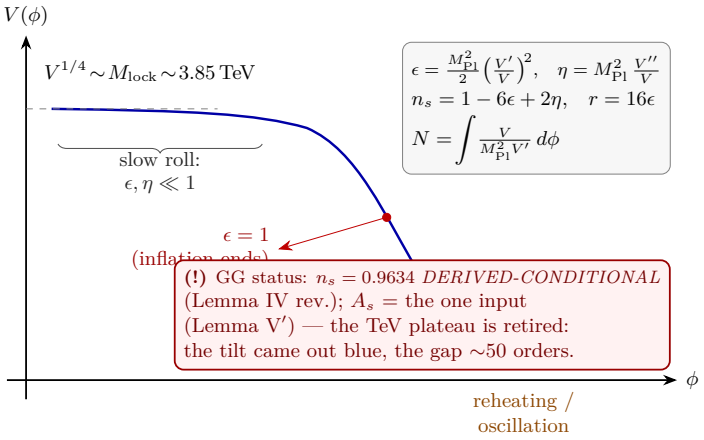

Figure 39: Slow-roll inflation, taught from the standard mechanism. The inflaton rolls slowly across a flat *plateau* (where the slow-roll parameters $\epsilon, \eta \ll 1$), inflation ends where $\epsilon = 1$ on the shoulder, and the field reheats by oscillating about the *minimum*. An earlier reading of GG-A10-Cosmos placed the plateau at the TeV scale, $V^{1/4} \sim M_{\rm lock} \sim 3.85\,\text{TeV}$; that reading is *retired*, and is kept visible here with its refutation: a TeV-scale plateau cannot simultaneously reproduce the observed $A_s \sim 2.1 \times 10^{-9}$ and $n_s \sim 0.965$, leaving an $\sim$50-order-of-magnitude normalization gap. **Honest codebook status (current):** the GDOS-forced plateau sits at $V^{1/4} = 8.1 \times 10^{15}\,\text{GeV}$, the tilt $n_s = 0.9634$ ($1 - 2/N^* = 0.9620$ at leading order) is *derived-conditional* (Lemma IV, revised), and the amplitude A_s remains the codebook's one declared input (Lemma V′). The figure teaches the standard slow-roll machinery (inset: ϵ, η, n_s, r, e-folds N) and marks where the *TeV* normalization did not close.

This appendix is written for a first-year graduate student who has met the Friedmann equations once and wants to understand, line by line, why cosmologists believe the early universe inflated, how a single scalar field can drive it, and exactly which numbers the theory predicts. As in Appendix E, we build the case in three steps.

Step 1. The standard mechanism (§G.1–G.9). We develop the slow-roll formalism from first principles — the three problems it solves, the e-fold counting, the inflaton equation of motion, the slow-roll parameters, the scalar and tensor spectra, and reheating — and work a textbook example in full.

Step 2. What inflation leaves open (§G.10). We then ask what the standard account puts in by hand: the inflaton itself, its potential, its energy scale, and the spectrum normalization.

Step 3. The geometric derivation (§G.11). Finally we show what GG-Theory supplies — the inflaton as a bulk modulus, its energy scale computed, its dynamics as the prototype GRIP descent, and a sharp two-sided tensor prediction, $r = 0.0038$ — and we mark, in plain sight, the one thing it does not yet supply: the spectrum normalization.

Inflation is the hardest case for the program's guiding belief — that the contents and laws of the universe should follow from geometry — and so it is the most revealing. Where the geometry succeeds it replaces a tailored mechanism with a computed one; where it does not yet succeed, it

says so. Nothing here is hidden behind a macro: every parameter is defined where it first appears, and every numerical claim is one the reader can reproduce with a calculator.

G.1 The three problems of the hot Big Bang

A purely decelerating, radiation-then-matter-dominated Friedmann universe is a spectacularly successful description of everything after about one second. It is, however, embarrassed by three initial-condition puzzles. Inflation is the proposal that a brief epoch of accelerated expansion preceded the hot Big Bang and erased all three at once. This subsection teaches the canonical three; at its end we pose a fourth that the canon never lists — and that, unlike the three, inflation alone cannot erase.

The horizon problem. The comoving distance light can have traveled since the Big Bang — the comoving particle horizon — is

$$\chi_{\rm hor}(t) = \int_0^t \frac{c\,dt'}{a(t')} = \int_0^a \frac{c\,da'}{a'^2 H(a')}, \qquad H \equiv \frac{\dot a}{a}. \tag{537}$$

In a universe dominated by a fluid with equation of state $p = w\rho$ one has $a \propto t^{2/[3(1+w)]}$ and $H \propto a^{-3(1+w)/2}$, so the integrand in (537) scales as $a'^{(1+3w)/2}\,da'/a'^2$. For ordinary matter ($w = 0$) or radiation ($w = 1/3$) the exponent $(1+3w)/2 - 2 < -1$ is such that the integral is dominated by its *upper* limit: the horizon grows monotonically and was small in the past. Concretely, at recombination ($z_\star \simeq 1090$) the particle horizon subtends only about one degree on the last-scattering surface. Yet the cosmic microwave background (CMB) is isotropic to one part in 10^5 across the *whole* sky — some 10^4 causally disconnected patches at the same temperature to four decimal places. Why should regions that never exchanged a single photon share a temperature? In the hot Big Bang this is simply an assumed initial condition.

The flatness problem. Write the Friedmann equation as

$$H^2 = \frac{8\pi G}{3}\rho - \frac{kc^2}{a^2} \quad\Longrightarrow\quad \Omega(a) - 1 = \frac{kc^2}{a^2H^2}, \qquad \Omega \equiv \frac{\rho}{\rho_{\rm crit}},\ \ \rho_{\rm crit} = \frac{3H^2}{8\pi G}. \tag{538}$$

The quantity $|\Omega - 1|$ is proportional to $1/(aH)^2$, the square of the comoving Hubble radius. In a decelerating universe aH *decreases* with time, so $|\Omega - 1|$ *grows*: spatial flatness is an unstable fixed point. Today $|\Omega_0 - 1| \lesssim 0.005$. Running this back to the Planck epoch with $aH \propto a^{-(1+3w)/2}$ for $w \geq 0$ requires $|\Omega - 1|$ to have been tuned to roughly one part in 10^{60} at $t \sim t_{\rm Pl}$. A universe that is even marginally flat today demands an absurdly fine-tuned initial curvature — unless some mechanism actively drives $\Omega \to 1$.

The relic (monopole) problem. Grand unified theories and, more generally, symmetry-breaking phase transitions in the early universe generically produce topological defects — magnetic monopoles, cosmic strings, domain walls — as massive, stable relics. Produced at roughly one per horizon volume at the transition, monopoles would redshift as non-relativistic matter ($\rho_{\rm M} \propto a^{-3}$, slower

than the a^{-4} of radiation) and quickly dominate the energy budget, overclosing the universe by many orders of magnitude. The observed universe contains no such relics at detectable abundance. Any mechanism that dilutes pre-existing relics far below one per current horizon volume solves this.

One cure for all three. The common root is that aH shrinks in a decelerating universe. Suppose instead there was an early epoch in which

$$\ddot{a} > 0 \quad \Longleftrightarrow \quad \frac{d}{dt}\left(\frac{1}{aH}\right) < 0 \quad \Longleftrightarrow \quad \rho + 3p < 0 \quad \Longleftrightarrow \quad w < -\tfrac{1}{3}. \tag{539}$$

The three statements are equivalent: accelerated expansion, a *shrinking* comoving Hubble radius, and a sufficiently negative pressure. The cleanest realization is $w = -1$, a constant vacuum-like energy density, giving exponential de Sitter expansion $a(t) \propto e^{Ht}$ with $H \approx$ const. During such a phase aH grows like a, so $|\Omega - 1| \propto 1/(aH)^2$ is driven *exponentially* toward zero (flatness solved); a tiny initial causal patch is stretched to encompass the entire observable universe, so the whole CMB sky descends from one thermalized region (horizon solved); and every pre-existing relic density is diluted by the same exponential volume factor (relic problem solved). A near-de Sitter, "quasi-de-Sitter," phase that lasts long enough and then ends gracefully is therefore the minimal cure.

The fourth problem, hiding inside the third: identity. Here is a teaching exercise worth pausing on, because it shows that even a canonical list can be incomplete. The three problems above all concern *conditions*: how uniform the temperature is, how flat the geometry is, how rare the relics are. Now ask a question of the same causal shape about *things*. Two hydrogen atoms on opposite sides of the sky — in regions that, without inflation, were never in causal contact — radiate exactly the same Balmer series; every electron in both regions carries the same mass and charge; the dark-matter abundance and the dark-energy density are the same on both sightlines. Temperature is one number that thermal contact might smooth; a particle's identity is a complete specification that no exchange of photons can negotiate. Why is the *parts list* of the universe the same in causally disconnected regions?

The student should notice three things. First, this question is not on the canonical list *as such* — its sharpest fragment entered the canon in disguise, as the domain problem: that causally disconnected regions choose their vacuum at random was seen already in the 1970s (Kibble, 1976; Zeldovich *et al.*, 1974), but the random choice was filed as a relic-abundance nuisance to be diluted, not as an identity question to be answered; and quantum field theory hides the rest of the question by assuming one field content in one vacuum globally, which is an assumption, not a derivation. Second, inflation *alone* does not answer it: a common causal patch explains shared conditions, but in any framework with more than one available vacuum (a landscape), a single patch can still fracture into domains that choose differently — shared history does not force shared identity. Third, the resolution in the present framework needs both of its halves: the *mold* (each object is the unique Ground Geometric Attractor of its stratum on one bulk — Part X), and the *ordering* taught in Step 3c of this appendix — the launch sits at the top of the descent, so inflation stretched the one causal patch *before any matter existed*, the visible sector was minted afterwards by one decay of one modulus, the Kaxion's one mark was pinned (by conjecture, Part V) while the sky was still one

patch, and the prompt lock crossing ($N_2 = 0$) left no alternative basin for any region to camp in. Sameness of state is inflation's gift; sameness of things is the geometry's. The full development — the *second horizon problem* — is Section 72 of Part X; the student who has followed this appendix owns every ingredient of its proof.

G.2 Counting e-folds: how much inflation is enough

The natural measure of "how much" the universe inflated is the number of e-folds of expansion,

$$N \equiv \ln \frac{a_{\rm end}}{a_{\rm start}} = \int_{t_{\rm start}}^{t_{\rm end}} H\, dt. \tag{540}$$

Each e-fold multiplies the scale factor by $e \approx 2.718$ and shrinks the comoving Hubble radius by the same factor. To solve the flatness problem we need $|\Omega - 1|$ suppressed by at least the factor by which it would subsequently *grow* during the radiation and matter eras. A scale that exits the horizon N e-folds before the end of inflation re-enters today if

$$N_\star \;\simeq\; \ln \frac{a_0 H_0}{a_\star H_\star} \;\approx\; 60 + \ln \frac{k/(a_0 H_0)}{0.05\,{\rm Mpc}^{-1}} + \frac{1}{4} \ln \frac{V_\star}{(10^{16}\,{\rm GeV})^4} + \ldots, \tag{541}$$

where $V_\star$ is the inflationary energy density when that scale left the horizon and the dots collect reheating-dependent terms. The numerical upshot, for the pivot scale probed by the CMB, is the famous benchmark

$$N_\star \;\approx\; 50\text{–}60. \tag{542}$$

The *lower* the inflationary energy scale, the *fewer* e-folds of observable inflation are required, because the post-inflationary expansion from a low $V_\star$ down to today is correspondingly shorter. This sliding of $N_\star$ with the energy scale matters when we reach the GG-Theory reading: the derived plateau sits one e-fold below the GUT edge ($V_0^{1/4} = 8.1 \times 10^{15}\,{\rm GeV}$), and the reheating match at Planck-strength width fixes $N_\star = 52.6$ — undiluted, because the later lock crossing is prompt — and the precise value feeds directly into the predicted tilt $n_s = 1 - 2/N_\star$. (At the retired TeV height this count would read $N_\star \sim 30$–40; the sliding logic is unchanged.)

G.3 A scalar field in FRW: action, stress tensor, equation of motion

The simplest source able to deliver $w < -1/3$ for an extended period is a single real scalar field $\phi(x)$, the *inflaton*, with canonical action

$$S = \int d^4x\, \sqrt{-g}\, \Big[\tfrac{1}{2} R\, M_{\rm Pl}^2 - \tfrac{1}{2} g^{\mu\nu} \partial_\mu \phi\, \partial_\nu \phi - V(\phi) \Big], \qquad M_{\rm Pl} \equiv \frac{1}{\sqrt{8\pi G}} = 2.435 \times 10^{18}\,{\rm GeV}, \tag{543}$$

where $M_{\rm Pl}$ is the *reduced* Planck mass (units $\hbar = c = 1$ throughout this subsection) and $V(\phi)$ is the potential. The stress–energy tensor follows from varying with respect to the metric,

$$T_{\mu\nu} = \partial_\mu \phi\, \partial_\nu \phi - g_{\mu\nu} \Big[\tfrac{1}{2} g^{\alpha\beta} \partial_\alpha \phi\, \partial_\beta \phi + V(\phi) \Big]. \tag{544}$$

For a homogeneous configuration $\phi = \phi(t)$ on the flat FRW background $ds^2 = -dt^2 + a^2(t)\,d\mathbf{x}^2$, this is the stress tensor of a perfect fluid with

$$\rho_\phi = \tfrac{1}{2}\dot\phi^2 + V(\phi), \qquad p_\phi = \tfrac{1}{2}\dot\phi^2 - V(\phi), \qquad w_\phi = \frac{p_\phi}{\rho_\phi} = \frac{\frac{1}{2}\dot\phi^2 - V}{\frac{1}{2}\dot\phi^2 + V}. \tag{545}$$

The reading of (545) is the whole story in one line: if the kinetic energy $\frac{1}{2}\dot\phi^2$ is *small* compared with the potential energy V, then $p_\phi \approx -\rho_\phi$, i.e. $w_\phi \approx -1$, and the field behaves like a temporary cosmological constant — precisely the $w < -1/3$ condition (539) that drives inflation. The field is then said to be "potential-dominated."

The equation of motion is the Euler–Lagrange equation for ϕ in the FRW background. Varying the action, or equivalently imposing covariant conservation $\nabla^\mu T_{\mu\nu} = 0$ on (544), gives the Klein–Gordon equation in an expanding universe,

$$\boxed{\ddot\phi + 3H\dot\phi + V'(\phi) = 0}, \qquad V' \equiv \frac{dV}{d\phi}. \tag{546}$$

This is the central equation of inflationary dynamics. It is the equation of a particle of unit mass at "position" ϕ rolling in the potential $V(\phi)$, subject to a velocity-proportional friction term $3H\dot\phi$ supplied by the Hubble expansion (Hubble drag). The two Friedmann equations close the system:

$$H^2 = \frac{1}{3M_{\rm Pl}^2}\Big[\tfrac{1}{2}\dot\phi^2 + V(\phi)\Big], \qquad \dot H = -\frac{\dot\phi^2}{2M_{\rm Pl}^2}. \tag{547}$$

The second relation, $\dot H = -\dot\phi^2/(2M_{\rm Pl}^2)$, shows that H can only decrease, and that it does so at a rate set entirely by the kinetic energy: a slowly rolling field gives an almost constant H, the quasi-de-Sitter regime.

G.4 The slow-roll approximation and the potential slow-roll parameters

"Slow roll" is the controlled approximation in which the field rolls so gently that (i) the kinetic energy is negligible against the potential, $\frac{1}{2}\dot\phi^2 \ll V$, and (ii) the acceleration is negligible against the friction, $|\ddot\phi| \ll 3H|\dot\phi|$. Under these two assumptions the dynamical equations (546) and (547) collapse to their leading-order "slow-roll" forms

$$3H\dot\phi \simeq -V'(\phi), \qquad H^2 \simeq \frac{V(\phi)}{3M_{\rm Pl}^2}. \tag{548}$$

The self-consistency of these two approximations is governed by two dimensionless numbers built purely from the *shape* of the potential — the *potential slow-roll parameters*:

$$\boxed{\epsilon \equiv \frac{M_{\rm Pl}^2}{2}\left(\frac{V'}{V}\right)^2, \qquad \eta \equiv M_{\rm Pl}^2\,\frac{V''}{V}}. \tag{549}$$

Here ϵ measures the squared *slope* of the potential (in Planck units) and η measures its *curvature*. Their meaning is transparent:

- $\epsilon \ll 1$ enforces that the slope is gentle enough for the field to be potential-dominated. Indeed one can show $\epsilon = -\dot{H}/H^2 \equiv \epsilon_H$ to leading order, and the accelerated-expansion condition $\ddot{a} > 0$ is *exactly* $\epsilon_H < 1$. Inflation proceeds while $\epsilon < 1$ and *ends* when ϵ grows to 1.
- $\eta \ll 1$ (more precisely $|\eta| \ll 1$) enforces that the curvature is small enough that the friction term dominates the acceleration term, keeping the velocity adiabatically slaved to the slope as in (548). It controls how slowly ϵ itself evolves.

Both conditions are necessary; a potential can be locally flat (ϵ small) yet sharply curved (η large), in which case slow roll fails. The pair $\{\epsilon, \eta\}$ is the minimal data from which all the leading observables below are computed.

G.5 The number of e-folds from the potential

With the slow-roll equations (548) in hand we can re-express the e-fold integral (540) as an integral over field values. Using $N = \int H\, dt = \int (H/\dot{\phi})\, d\phi$ and $\dot{\phi} \simeq -V'/(3H)$ with $H^2 \simeq V/(3M_{\rm Pl}^2)$,

$$\boxed{N = \int_{t_{\rm start}}^{t_{\rm end}} H\, dt = \frac{1}{M_{\rm Pl}^2} \int_{\phi_{\rm end}}^{\phi_\star} \frac{V}{V'}\, d\phi = \frac{1}{M_{\rm Pl}} \int_{\phi_{\rm end}}^{\phi_\star} \frac{d\phi}{\sqrt{2\epsilon}}}, \tag{550}$$

where $\phi_\star$ is the field value N e-folds before the end of inflation and $\phi_{\rm end}$ is fixed by the end-of-inflation condition $\epsilon(\phi_{\rm end}) = 1$. The second form, with $|V'/V| = \sqrt{2\epsilon}/M_{\rm Pl}$, makes manifest that the field traverses $\Delta\phi \sim M_{\rm Pl}\sqrt{2\epsilon}\, N$ in field space: small ϵ means many e-folds per unit field range. Equation (550) is the practical workhorse — given $V(\phi)$, it returns $\phi_\star$ for the required $N_\star \approx 50$–60, and the observables below are then all evaluated *at* $\phi_\star$.

G.6 Quantum fluctuations and the scalar power spectrum

The deepest payoff of inflation is that it does not merely homogenize the universe; it seeds the small inhomogeneities that grow into galaxies and imprint the CMB anisotropies. The seeds are the quantum vacuum fluctuations of the inflaton (and of the metric), stretched to super-horizon scales by the accelerated expansion and frozen in as classical curvature perturbations.

Decompose the field into background plus fluctuation, $\phi(t, \mathbf{x}) = \bar{\phi}(t) + \delta\phi(t, \mathbf{x})$. In the quasi-de-Sitter background each Fourier mode of a (nearly) massless field acquires, at horizon crossing $k = aH$, a fluctuation of characteristic amplitude

$$\langle |\delta\phi_k|^2 \rangle^{1/2} \simeq \frac{H}{2\pi}, \tag{551}$$

the de Sitter Hawking-temperature-like amplitude. The physically conserved, gauge-invariant quantity on super-horizon scales is the comoving curvature perturbation $\mathcal{R}$, related to the field fluctuation by $\mathcal{R} = -(H/\dot{\bar{\phi}})\, \delta\phi$. Its dimensionless power spectrum is defined by $\langle \mathcal{R}_{\mathbf{k}} \mathcal{R}_{\mathbf{k}'} \rangle =$

$(2\pi)^3\delta^3(\mathbf{k}+\mathbf{k}')\,(2\pi^2/k^3)\,\mathcal{P}_{\mathcal{R}}(k)$, and evaluating it at horizon crossing gives the master result

$$\boxed{\mathcal{P}_{\mathcal{R}}(k) = \frac{H^2}{\dot{\bar{\phi}}^2}\left(\frac{H}{2\pi}\right)^2\Bigg|_{k=aH} = \frac{1}{8\pi^2}\frac{H^2}{M_{\rm Pl}^2\,\epsilon}\Bigg|_{k=aH} = \frac{1}{24\pi^2}\frac{V}{M_{\rm Pl}^4\,\epsilon}\Bigg|_{k=aH}}. \tag{552}$$

The last equality used $H^2 \simeq V/(3M_{\rm Pl}^2)$. The amplitude of this spectrum at a chosen pivot wavenumber $k_\star$ (Planck uses $k_\star = 0.05\,\mathrm{Mpc}^{-1}$) is the *scalar amplitude*

$$A_s \equiv \mathcal{P}_{\mathcal{R}}(k_\star) = \frac{1}{24\pi^2}\frac{V_\star}{M_{\rm Pl}^4\,\epsilon_\star}, \qquad A_s^{\rm obs} = (2.10 \pm 0.03)\times 10^{-9}. \tag{553}$$

This is the celebrated COBE/Planck normalization: *any* inflationary model must reproduce $A_s \approx 2.1\times 10^{-9}$, and this single number fixes the overall height-to-flatness combination $V_\star/\epsilon_\star$ of the potential. Note carefully what (553) says: A_s is *not* a free choice once the potential is specified; it is a prediction. Conversely, a model that does not specify $V_\star$ and $\epsilon_\star$ *cannot* predict A_s at all — it can only assert the observed value. Hold that thought.

The *spectral index* (scalar tilt) measures the scale dependence of $\mathcal{P}_{\mathcal{R}}$,

$$n_s - 1 \equiv \frac{d\ln\mathcal{P}_{\mathcal{R}}}{d\ln k}\Bigg|_{k_\star}. \tag{554}$$

Carrying the k-derivative through (552) (using $d\ln k \simeq H\,dt$ at horizon crossing and the slow-roll relations) yields the standard leading-order expression

$$\boxed{n_s = 1 - 6\epsilon + 2\eta}, \qquad n_s^{\rm obs} = 0.9649 \pm 0.0042. \tag{555}$$

A scale-invariant Harrison–Zeldovich spectrum would have $n_s = 1$; the measured *red tilt* $n_s < 1$ is a genuine, non-trivial confirmation that inflation involves a gently sloped, gently curved potential (ϵ, η small but nonzero) rather than an exact cosmological constant. The combination $6\epsilon - 2\eta \approx 0.035$ is the headline inflationary measurement.

G.7 The tensor spectrum, r, and the consistency relation

Inflation also stretches the two transverse-traceless polarizations of the metric — primordial gravitational waves. Being a massless spin-2 field, each polarization fluctuates like a massless scalar, so the *tensor power spectrum* is

$$\mathcal{P}_t(k) = \frac{2}{\pi^2}\frac{H^2}{M_{\rm Pl}^2}\Bigg|_{k=aH} = \frac{2}{3\pi^2}\frac{V}{M_{\rm Pl}^4}\Bigg|_{k=aH}, \tag{556}$$

where the factor of 2 sums the two polarizations and the result is independent of ϵ. Its tilt is

$$n_t \equiv \frac{d\ln\mathcal{P}_t}{d\ln k}\Bigg|_{k_\star} = -2\epsilon. \tag{557}$$

The crucial observable is the *tensor-to-scalar ratio*, the relative amplitude of gravitational waves to density perturbations,

$$r \equiv \frac{\mathcal{P}_t(k_\star)}{\mathcal{P}_{\mathcal{R}}(k_\star)} = \frac{(2/\pi^2)(H^2/M_{\rm Pl}^2)}{(1/8\pi^2)(H^2/M_{\rm Pl}^2\epsilon)} = 16\,\epsilon, \qquad r^{\rm obs} < 0.036 \ (95\%). \tag{558}$$

Dividing (556) by (552), all the $H^2/M_{\rm Pl}^2$ factors cancel and one is left with the beautifully simple

$$\boxed{r = 16\,\epsilon}. \tag{559}$$

Because $\mathcal{P}_t \propto H^2 \propto V$, a *measurement* of r is a direct measurement of the inflationary energy scale:

$$V_\star^{1/4} \simeq \left(\frac{r}{0.01}\right)^{1/4} \times 1.0 \times 10^{16}\,{\rm GeV}. \tag{560}$$

A detectable $r \gtrsim 10^{-3}$ requires GUT-scale inflation ($V_\star^{1/4} \sim 10^{16}\,{\rm GeV}$); conversely, low-scale inflation predicts an unobservably tiny r. Combining (557) and (559) gives the single-field *consistency relation*,

$$\boxed{r = -8\,n_t}, \tag{561}$$

a sharp, falsifiable prediction of single-field slow-roll inflation: the tensor amplitude and the tensor tilt are not independent, and a future measurement of both would test the entire single-field framework. (The minus sign records that a red tensor tilt, $n_t < 0$, accompanies a positive r.)

G.8 Reheating: from cold inflaton to hot Big Bang

When ϵ grows to unity inflation ends and the field, no longer slow-rolling, falls toward the minimum of V and oscillates about it. These coherent oscillations are an enormous, cold, zero-momentum condensate of inflaton quanta. Through couplings to other fields the condensate decays, dumping its energy into a thermal bath of relativistic particles — *reheating*. Parametrizing the inflaton decay by a rate Γ_ϕ, reheating completes when $H \sim \Gamma_\phi$, fixing the reheating temperature

$$T_{\rm reh} \simeq \left(\frac{90}{\pi^2 g_\star}\right)^{1/4} \sqrt{\Gamma_\phi\, M_{\rm Pl}}, \tag{562}$$

with $g_\star$ the effective number of relativistic degrees of freedom. Reheating must complete above the scale of Big Bang nucleosynthesis ($T_{\rm reh} \gtrsim$ a few MeV) so that the standard hot Big Bang takes over with the correct light-element yields, and it sets the precise $N_\star$ in (541) (a lower $T_{\rm reh}$ means fewer observable e-folds). Reheating is the hand-off: it converts the ordered potential energy of the inflaton into the disordered thermal energy of the hot Big Bang, and its details feed back into the predicted n_s at the few-$\times 10^{-3}$ level through $N_\star$.

G.9 Worked example (STANDARD TEXTBOOK PEDAGOGY): quadratic and Starobinsky potentials

The following two worked examples are the *standard textbook pedagogy*, included so the machinery above is concrete. They are deliberately the canonical models every cosmology course computes; they are *not* the GG-Theory model, and the explicit gap between this pedagogy and the GG-Theory case is stated in the next subsection.

Example A: quadratic ("$m^2\phi^2$") inflation. Take the simplest chaotic-inflation potential

$$V(\phi) = \tfrac{1}{2}m^2\phi^2. \tag{563}$$

Then $V' = m^2\phi$ and $V'' = m^2$, so the slow-roll parameters (549) are

$$\epsilon = \frac{M_{\rm Pl}^2}{2}\left(\frac{m^2\phi}{\frac{1}{2}m^2\phi^2}\right)^2 = \frac{2M_{\rm Pl}^2}{\phi^2}, \qquad \eta = M_{\rm Pl}^2\frac{m^2}{\frac{1}{2}m^2\phi^2} = \frac{2M_{\rm Pl}^2}{\phi^2} = \epsilon. \tag{564}$$

Inflation ends at $\epsilon = 1$, i.e. $\phi_{\rm end} = \sqrt{2}\,M_{\rm Pl}$. The e-fold integral (550) gives

$$N = \frac{1}{M_{\rm Pl}^2}\int_{\phi_{\rm end}}^{\phi_\star}\frac{\frac{1}{2}m^2\phi^2}{m^2\phi}\,d\phi = \frac{1}{M_{\rm Pl}^2}\int_{\phi_{\rm end}}^{\phi_\star}\frac{\phi}{2}\,d\phi = \frac{\phi_\star^2 - \phi_{\rm end}^2}{4M_{\rm Pl}^2} \implies \phi_\star^2 = (4N+2)\,M_{\rm Pl}^2. \tag{565}$$

For the benchmark $N_\star = 60$ this gives $\phi_\star = \sqrt{242}\,M_{\rm Pl} \approx 15.6\,M_{\rm Pl}$, and hence

$$\epsilon_\star = \frac{2M_{\rm Pl}^2}{\phi_\star^2} = \frac{2}{242} \approx 8.26\times10^{-3}, \qquad \eta_\star = \epsilon_\star \approx 8.26\times10^{-3}. \tag{566}$$

The observables follow directly:

$$n_s = 1 - 6\epsilon_\star + 2\eta_\star = 1 - 4\epsilon_\star = 1 - \frac{8}{242} \approx 0.967, \tag{567}$$

$$r = 16\,\epsilon_\star = \frac{32}{242} \approx 0.132, \tag{568}$$

$$n_t = -2\epsilon_\star \approx -0.0165 \quad (\text{check: } -8n_t = 0.132 = r,\ \checkmark). \tag{569}$$

The COBE normalization (553) then fixes the inflaton mass: demanding $A_s = 2.1\times10^{-9}$ with $V_\star = \frac{1}{2}m^2\phi_\star^2$ and $\epsilon_\star$ from (566) gives $m \approx 6\times10^{-6}\,M_{\rm Pl} \approx 1.4\times10^{13}\,\text{GeV}$ — i.e. GUT-scale inflation. Note the verdict: $m^2\phi^2$ predicts a *good* tilt $n_s \approx 0.967$ but a tensor ratio $r \approx 0.13$ that the Planck/BICEP bound $r < 0.036$ now *excludes*. Quadratic inflation is a textbook model that the data have killed — a clean illustration that n_s and r together, not n_s alone, are the discriminating test.

Example B: Starobinsky (R^2) inflation. The historically first and observationally favored model adds an R^2 term to gravity; in the Einstein frame it is equivalent to a scalar with the plateau

potential

$$V(\phi) = \Lambda^4 \left(1 - e^{-\sqrt{2/3}\,\phi/M_{\rm Pl}}\right)^2. \tag{570}$$

Write $x \equiv e^{-\sqrt{2/3}\,\phi/M_{\rm Pl}}$, so $V \propto (1-x)^2$. Then $V' = \Lambda^4\, 2(1-x)\, x\, \sqrt{2/3}/M_{\rm Pl}$ and a short computation gives, in the large-ϕ (small-x) plateau regime where $x \ll 1$,

$$\epsilon \simeq \frac{4}{3}\,\frac{x^2}{(1-x)^2} \approx \frac{4}{3}x^2, \qquad \eta \simeq -\frac{4}{3}\,\frac{x(1-2x)}{(1-x)^2} \approx -\frac{4}{3}x. \tag{571}$$

The e-fold count (550) evaluates (for $\phi_\star$ on the plateau) to the compact result $N \simeq \frac{3}{4}\, e^{\sqrt{2/3}\,\phi_\star/M_{\rm Pl}} = \frac{3}{4}/x_\star$, so $x_\star \simeq 3/(4N)$. Substituting back,

$$\epsilon_\star \simeq \frac{4}{3}\left(\frac{3}{4N}\right)^2 = \frac{3}{4N^2}, \qquad \eta_\star \simeq -\frac{4}{3}\cdot\frac{3}{4N} = -\frac{1}{N}, \tag{572}$$

$$n_s = 1 - 6\epsilon_\star + 2\eta_\star \simeq 1 - \frac{2}{N}, \qquad r = 16\epsilon_\star \simeq \frac{12}{N^2}. \tag{573}$$

For $N_\star = 60$ these give the textbook Starobinsky predictions

$$n_s \approx 1 - \frac{2}{60} = 0.9667, \qquad r \approx \frac{12}{3600} \approx 3.3 \times 10^{-3}, \qquad n_t = -2\epsilon_\star \approx -4.2 \times 10^{-4}, \tag{574}$$

in excellent agreement with Planck ($n_s = 0.9649 \pm 0.0042$, $r < 0.036$). The COBE normalization fixes the plateau height to Λ^4 with $\Lambda \approx 3 \times 10^{16}$ GeV, i.e. again GUT-scale inflation. The pedagogical lesson of the two examples together: *both* the right tilt $n_s \approx 0.965$ *and* the small r are reproduced by a GUT-scale *plateau* model, and doing so requires the model to predict the height $V_\star$ (hence A_s) and the curvature η (hence the tilt) from one and the same potential. The two observables A_s and n_s are jointly constrained by the shape of $V(\phi)$; you do not get to choose them independently.

G.10 What standard inflation explains — and what it leaves open

The machinery of Step 1 is a triumph: a single slowly-rolling scalar flattens the universe, stretches quantum fluctuations into the seeds of galaxies, and predicts a nearly scale-invariant spectrum that the microwave background confirms to remarkable precision. But it is worth asking, exactly as we did for the Friedmann background in §E.1, what the standard account must put in by hand.

1. **The inflaton itself.** Standard inflation does not say what the field ϕ *is*. It is introduced because the early universe needs a phase of accelerated expansion; no independent physics requires it to exist. Inflation, in the textbook telling, is a mechanism in search of a cause.

2. **The potential.** The shape $V(\phi)$ — plateau, shoulder, minimum — is chosen so that the predictions land on the data. Dozens of potentials are on the market precisely because none is forced: the data select among them rather than the theory predicting one.

3. **The energy scale.** How high the plateau sits, $V^{1/4}$, anywhere from the GUT scale to far below, is not fixed by the mechanism. It is the single most important number in the story, and it is a

free dial.

4. **The spectrum normalization.** The amplitude A_s and tilt n_s are *measured* and then used to fix the potential's height and slope, not derived from anything deeper.

So inflation, for all its success, leaves the same kind of gap the Friedmann background did: it accommodates the early universe with a field and a potential tailored to fit, but it does not say why that field exists or what sets its scale. A geometric account should supply the cause. As we will see, GG-Theory supplies some of it cleanly — the identity of the field, its energy scale, and one sharp falsifiable prediction — and is honest that it does not yet supply the rest. This is the most open sector of the paper, and we mark the boundary exactly.

G.11 The geometric derivation: the two-basin descent, taught slowly

Of the four gaps of §G.10, GG-Theory closes the first three — the identity of the field, its energy scale, and its dynamics — computes the tilt and the tensor ratio, and leaves exactly one number as a declared input: the amplitude. This subsection teaches that result at course pace, because it is the centerpiece of the paper's inflationary sector and the single best illustration of how a geometric Step 3 replaces a textbook's freedom with a computation. The student who has worked through Step 1 owns every tool needed here: nothing below uses more than the slow-roll formulae already derived.

Step 3a — the inflaton is a bulk modulus, not an invention (gap 1). The field ϕ is not introduced for the occasion. It is identified with the *scale modulus* Σ_r of the six-dimensional bulk — the one-dimensional flat direction of the geometry's shape space along the radial interval (Part II; Arisaka, 2026, A2-GG6). Such a modulus is an unavoidable feature of the bulk, present whether or not one wants an inflaton; GG-Theory does not add a field to drive inflation, it *recognizes* that the bulk already carries one — and only one, which is why there is no choice to make. *Physics intuition:* the inflaton is the size of the hidden geometry, slowly relaxing — the same radial coordinate that Appendix E used as the energy axis, seen now as a dynamical field.

Step 3b — the potential's *class* is derived, not chosen (gap 2). Here is the step where the textbook's biggest freedom — the functional form of $V(\phi)$, chosen from a zoo — becomes a computation. The GDOS boundary law is a conformal law with a non-conformal coupling to the scale modulus; its trace anomaly generates an R^2 term in the boundary effective action, and the Einstein-frame map of an R^2 theory is a theorem of classical field theory (no cosmology input at all): the potential is forced into the Starobinsky class

$$V(\phi) \;=\; \tfrac{3}{4}\,M^2 M_{\mathrm{Pl,red}}^2 \left(1 - e^{-\sqrt{2/3}\,\phi/M_{\mathrm{Pl,red}}}\right)^2, \tag{575}$$

with the plateau shape and the $\sqrt{2/3}$ fixed by four-dimensional gravity and only the single mass scale M free (Starobinsky, 1980; Part IV). Every slow-roll quantity of Step 1 now follows from Eq. (575) by pure calculus — the student should verify, using the definitions of §G.10, that on the

plateau $\epsilon_V \simeq \frac{3}{4}\, e^{-2\sqrt{2/3}\,\phi/M_{\rm Pl,red}}$ and $\eta_V \simeq -\frac{4}{3}\, e^{-\sqrt{2/3}\,\phi/M_{\rm Pl,red}}$, so that slow roll fails (graceful exit) where the plateau steepens, with no waterfall field and no tuning. *Physics intuition:* the geometry does not hand us a potential; it hands us a *class* — one shape, one free height — and the shape alone will fix the observables.

Step 3c — where the plateau sits, and the descent below it. The derived plateau is *GUT-adjacent*: the observed amplitude places its height at

$$V_0^{1/4} \;=\; 8.1\times 10^{15}\ \text{GeV}, \qquad H_\star \;\simeq\; 1.5\times 10^{13}\ \text{GeV}, \tag{576}$$

one e-fold below the GUT edge where the descent begins. The student should note — because the record of Part IX carries it — that the program once conjectured the opposite: a plateau at the TeV lock scale $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-35} = 3.85$ TeV. Step 3e shows, as a worked exercise, why that reading is refuted at sign level. The true role of the lock scale is different and sharper: it is the *second basin* of the descent. The full history is a *two-basin descent* of the one modulus: the primary plateau (inflation, $N \simeq 60$ e-folds), the graceful exit at its steep edge, reheating, then — far down the ladder — the prompt crossing of the lock basin ($N_2 = 0$: the GRIP ledger-action bound on Σ_r forbids a second inflation, with some fifty orders of margin), and finally the terminus, the Ground Geometric Attractor whose residual energy is today's dark energy (Part IV). *Physics intuition:* think of the $\tau \to \mu \to e$ cascade — three basins, two crossings, the stable floor at the bottom. The vacuum has generations exactly as matter does; inflation is the τ of the vacuum, the lock basin its muon (real but promptly passed), dark energy its electron. That analogy, the *vacuum generations*, is developed as a load-bearing statement in Part IV §31.1.

Step 3d — the tilt and the e-fold count, computed end to end. Now the arithmetic the textbook could never do, carried in full. For the plateau class, the standard slow-roll results of Step 1 collapse to the one-parameter family

$$n_s \;=\; 1 - \frac{2}{N^*}, \qquad r \;=\; \frac{12}{N^{*2}}, \tag{577}$$

so everything hangs on the number of e-folds N^* before the end of inflation at which the CMB pivot scale left the horizon — and N^* is fixed by reheating, not guessed. The modulus couples to the boundary at Planck strength, so its decay rate is $\Gamma \simeq M^3/M_{\rm Pl,red}^2$; with the scalaron mass $M = 3.1\times 10^{13}$ GeV this gives $\Gamma \simeq 5.0$ TeV and a reheating temperature

$$T_{\rm rh} \;\simeq\; \left(\frac{90}{\pi^2 g_*}\right)^{1/4} \sqrt{\Gamma\, M_{\rm Pl,red}} \;\simeq\; 6\times 10^{10}\ \text{GeV}, \tag{578}$$

(the student should confirm the arithmetic with $g_* \simeq 100$). Feeding $T_{\rm rh}$ into the standard horizon-matching formula of Step 1 — and using the prompt lock crossing $N_2 = 0$, so that no later epoch dilutes the count — yields

$$N^* \;=\; 52.6, \qquad \Longrightarrow \qquad n_s \;=\; 1 - \frac{2}{52.6} \;=\; 0.9620, \qquad r \;=\; \frac{12}{(52.6)^2} \;=\; 0.0043 \tag{579}$$

at leading order, and $n_s = 0.9634$, $r = 0.0038$ evaluated exactly on the plateau (§28.10). Planck measures $n_s = 0.9649 \pm 0.0042$: the exact tilt sits 0.4σ from the data, *risked against them, not read from them.* And the tensor target is two-sided and decisive within the decade: a factor nine below the current bound, several times above the Simons-Observatory/LiteBIRD floor. A measurement inconsistent with $r \simeq 0.004$ in either direction falsifies the derived launch outright. *Physics intuition:* once the shape of the hill is a theorem, the only question is where on the hill the pivot scale exited — and reheating physics answers that, leaving nothing to tune.

Step 3e — the honest boundary: the amplitude, and the retired TeV reading. One number remains. The scalar amplitude $A_s = V_\star/(24\pi^2 M_{\rm Pl,red}^4\,\epsilon_\star) \simeq 2.1 \times 10^{-9}$ fixes the height M — and the height is the sector's *one declared input*, exactly as the epoch is the one declared input of the π-locked inventory. The named route to paying it is Lemma V′ (Conjecture 68.7): the GDOS heat-kernel evaluation of the R^2 coefficient on the warped interval, whose warp candidate $M = M_5 e^{-K/3}$ currently lands within 87% of the required mass with no parameter. *This is an open computation stated precisely enough for the reader to attempt* — a graduate student who closes it makes the launch sector input-free, and the invitation is meant literally. Finally, the retired reading, kept as a worked exercise in falsification: a TeV-scale plateau gives $V_\star/M_{\rm Pl,red}^4 \sim 10^{-59}$, so matching A_s would demand $\epsilon_\star \sim 10^{-53}$, and the tilt would come out *blue* ($n_s - 1 \sim +10^{-26}$) against the measured red -0.0351 ± 0.0042 — wrong in sign, with no parameter to rescue it. The student should reproduce both numbers; nothing teaches the difference between a fitted model and a forced one faster than watching a conjecture die of arithmetic.

What this appendix establishes. The honest ledger is therefore strong, and its one gap is named. Standard inflation (Step 1) is reproduced in full and is established textbook physics. Geometry then supplies what the textbook left open: the inflaton is the unique bulk scale modulus rather than an invention (Step 3a); the potential's class is derived rather than chosen from a zoo (Step 3b); the plateau sits one e-fold below the GUT edge, at the top of a two-basin descent whose later basins are the Kaxion's birthplace and dark energy (Step 3c); the tilt and the tensor ratio are computed end to end — $n_s = 0.9634$, $r = 0.0038$ — with the e-fold count fixed by the derived reheating temperature (Step 3d); and the amplitude is the one declared input, with the computation that would pay it stated as an open target (Step 3e). In the codebook this reads: n_s and $V_{\rm inf}^{1/4}$ Derived-conditional on the revised Lemma IV, r Derived-conditional with them, A_s Asserted pending Lemma V′. In the Euclidean spirit, the parts that follow from the geometry are stated as results, and the part that does not is marked, in plain sight, as the work that remains — and offered, in this course, to the generation that will finish it.

Problems and worked solutions. G.1.1. Quantify the horizon problem: what angle does a causal patch at recombination subtend on today's sky?

Solution. Step 1: in a matter-dominated history the comoving particle horizon grows as $\eta \propto a^{1/2}$, so at recombination it was $\eta_{\rm rec} = \eta_0/\sqrt{1+z_*} \approx \eta_0/33$. Step 2: that comoving length is seen today at comoving distance $\approx \eta_0$, subtending $\theta \sim \eta_{\rm rec}/\eta_0 \approx 1/33$ rad $\approx 2^\circ$ — of order one degree after

the exact geometry is kept. Step 3: the CMB is isothermal to 10^{-5} across 180°: tens of thousands of patches that, in the hot-big-bang history, never exchanged a photon, all at one temperature. Physics intuition: this is not a small discrepancy to adjust away but a causality contradiction — and the chapter's resolution inverts the picture: make the comoving Hubble radius *shrink* for long enough (537) and the whole sky shares one causal ancestor.

G.2.1. Derive the $\sim$60 e-fold requirement from today's Hubble radius.
Solution. Step 1: the flatness/horizon cures need the comoving Hubble radius $(aH)^{-1}$ to shrink during inflation by at least the factor it has grown since (541). Step 2: from reheating at $T_{\rm rh}$ to today, $(aH)^{-1}$ grows (radiation era) as $a^1\cdot$(matter-era corrections); logarithmically, the growth is $\ln[(aH)_0^{-1}/(aH)_{\rm rh}^{-1}] \approx \ln(T_{\rm rh}/T_0)$ up to $O(1)$ pieces. Step 3: with this chapter's *derived* reheating temperature $T_{\rm rh} \simeq 6\times10^{10}$ GeV against $T_0 \approx 2.35\times10^{-13}$ GeV: $\ln(2.6\times10^{23}) \approx 54$ — inside the benchmark band $N_\star \approx 50$–60 (542), and about one e-fold above the $N_\star = 52.6$ that the reheating matching of §28 returns and G.9.2 reads back from the derived tilt; the difference is the $O(1)$ pieces this bare logarithm drops, and the body closes it at Eq. (75). What remains open in the cohort is the derivation of $\ln(B/kT_0)$ from the descent itself. Physics intuition: sixty e-folds is not a magic number but a mirror — inflation must un-do, in advance, exactly the growth the subsequent hot history will perform, and the requirement is read off today's sky.

G.7.1. Find the energy scale of inflation implied by the derived tensor ratio.
Solution. Step 1: the tensor amplitude ties r to the inflationary potential: $V_\star = \frac{3\pi^2}{2} A_s\, r\, M_{\rm Pl}^4$ (560) (reduced-Planck units). Step 2: with $A_s = 2.1\times10^{-9}$ and the derived $r = 0.0038$: $V_\star^{1/4} \approx 8.0\times10^{15}$ GeV (the plateau height $V_0^{1/4} = 8.1\times10^{15}$ GeV lies one percent above it). Step 3: this sits just below the classic GUT scale and within an order of magnitude of the cohort's own high edge — the plateau is a statement that inflation ran at the energy where the ledger's deepest structure lives, and the sub-detectable-but-nonzero r is its signature. Physics intuition: r is a thermometer: gravitational waves are generated purely by the expansion rate, so their amplitude reads the inflationary energy directly — which is why a future r detection at the 10^{-3} level is the sharpest single test this sector offers (the falsification catalog prices it).

G.7.2. Prove the slow-roll consistency relation $n_t = -r/8$, and state its use.
Solution. Step 1: the tensor spectrum tilts as $n_t = -2\epsilon_\star$ (557), from $P_T \propto H^2$ and $\dot H = -\epsilon H^2$. Step 2: the tensor-to-scalar ratio is $r = 16\epsilon_\star$ (559). Step 3: eliminate $\epsilon_\star$: $n_t = -r/8$ (561) — a relation between two observables with *no* free function of the potential surviving. Any single-field slow-roll model must satisfy it; a measured violation retires the whole class, this chapter's geometric reading included. Physics intuition: consistency relations are the fingerprints of a mechanism rather than a model — the potential drops out, leaving pure kinematics of one clock rolling down one hill; nature either shows the fingerprint or convicts a different suspect.

G.8.1. Bound the reheating temperature from instantaneous conversion.
Solution. Step 1: if the inflaton's energy $V_\star$ thermalizes instantly into $g_* = 106.75$ degrees of freedom, then $\frac{\pi^2}{30} g_* T_{\rm rh}^4 = V_\star$ (562). Step 2: with $V_\star^{1/4} = 8.0\times10^{15}$ GeV (G.7.1): $T_{\rm rh}^{\rm max} = \left(\frac{30}{\pi^2 g_*}\right)^{1/4} V_\star^{1/4} \approx 3.3\times10^{15}$ GeV. Step 3: any delay in thermalization only lowers this — reheating is a ceiling, not a prediction, until the coupling chain is fixed. The cohort's descent supplies that chain, and its

derived value, $T_{\rm rh} \simeq 6 \times 10^{10}$ GeV, sits five orders below the ceiling: reheating here is slow, and it is exactly this derived $T_{\rm rh}$ that feeds the e-fold count of G.2.1 and guarantees the bath for the resonant leptogenesis of Appendix I. Physics intuition: reheating is the ignition of the hot Big Bang by the death of inflation — everything in Appendices F–I happens downstream of this one temperature, which is why the e-fold budget treats it as the keystone number.

G.9.1. Rule out quadratic inflation with one line of arithmetic.
Solution. Step 1: for $V = \frac{1}{2}m^2\phi^2$ the chapter derives $\phi_\star^2 = (4N+2)M_{\rm Pl}^2$ (565), hence $\epsilon_\star = 2M_{\rm Pl}^2/\phi_\star^2 = 2/(4N+2)$. Step 2: $r = 16\epsilon_\star = 32/(4N+2) \approx 0.13$ for $N_\star = 60$ (568). Step 3: the observational bound is $r < 0.036$: quadratic inflation overshoots by a factor of four — dead, not wounded. The same arithmetic gives $n_s = 1 - 2/N - \ldots \approx 0.967$: acceptable — it is the *tensors* that convict. Physics intuition: large-field models drag the inflaton across many Planck lengths, and every Planck length of travel pumps gravitational waves; the sky's silence in B-modes is a statement about the *distance rolled*, which is why plateau potentials (next problem) survive where monomials die.

G.9.2. Show that the cohort's derived tilt and tensor ratio are one prediction, not two.
Solution. Step 1: on the Starobinsky-class plateau the chapter derives $n_s = 1 - 2/N_\star$ and $r = 12/N_\star^2$ (573) — two observables, one parameter. Step 2: the cohort's derived tilt is $n_s = 0.9620$ at leading order; solving, $N_\star = 2/(1-n_s) = 52.6$. Step 3: inserting that *same* $N_\star$: $r = 12/(52.6)^2 = 0.0043$ — exactly the cohort's leading-order tensor ratio, and the exact evaluation on the plateau at the same $N_\star$ returns the pair the cohort stakes, (0.9634, 0.0038). The pair (n_s, r) lies on the plateau's one-parameter curve, and the e-fold number it selects is the one the reheating matching of Eq. (75) put in — the pair is one prediction because the body computes the count once and evaluates both observables at it. Physics intuition: on a plateau, tilt and tensors are two shadows of one length — how far from the edge the descent began; measure either, predict the other. The cohort's pair is self-consistent to the last digit, and the Simons Observatory and LiteBIRD test the prediction within the decade.

Take-away. *Standard result:* slow-roll inflation solves the horizon, flatness, and relic problems, but the inflaton's identity, its potential, its scale, and its e-fold count are all chosen, not derived. *What it assumed:* a field, a shape $V(\phi)$, a height, and $N^* \sim 50$–60 by convention. *What the geometry replaced it with:* the unique bulk scale modulus on a derived Starobinsky-class plateau one e-fold below the GUT edge, with $N^* = 52.6$ from the computed reheating temperature, $n_s = 0.9634$ (0.4σ), $r = 0.0038$ (decisive this decade), and one declared input — the amplitude — carried by the named open computation Lemma V′. *Continue in:* Part IV (the descent), Part X (identity), GG-C1-Kaxion (the Kaxion's angle released at the second crossing).

H Big-Bang nucleosynthesis and the lithium-7 problem

What this appendix teaches. Why the first three minutes made one quarter of the baryons into helium and almost nothing else; weak freeze-out, the deuterium bottleneck, and the full reaction network derived with worked numbers; the observational triumph (D, He) and the stubborn factor-3 lithium-7 problem; and the geometric reading in which an open-system $A = 7$ port suppresses the lithium sector by $1/\pi$ while leaving deuterium and helium untouched.

Read alongside: Kolb & Turner (1990) ch. 4 (the classic treatment); Ryden (2017) ch. 10 for a first pass.

The primordial-nucleosynthesis reaction network, the helium and deuterium yields (Figure 40), and the lithium problem treated in this appendix follow the foundational and modern precision references (Cyburt *et al.*, 2016; Fields, 2011; Pitrou *et al.*, 2018; Steigman, 2007; Wagoner *et al.*, 1967).

This appendix is a self-contained, ground-up account of primordial nucleosynthesis, pitched at a first-year graduate student who has seen the Friedmann equations and elementary statistical mechanics but has never followed a light-element abundance calculation line by line. As in Appendix E, we build it in three steps.

Step 1. Standard BBN (§H.1–H.6). The neutron-to-proton freeze-out, the deuterium bottleneck, the light-element reaction network, the dependence on the baryon-to-photon ratio, and a worked helium calculation — all derived from scratch.

Step 2. Where it succeeds and where it fails (§H.5). Two of the three measured abundances (deuterium, helium-4) are among the great triumphs of cosmology. The third, lithium-7, is not: standard BBN overshoots it by a factor of three. This is the one genuine failure of the standard account, and it is the opening for the geometry.

Step 3. The geometric derivation (§H.7). The same closed-packet factor $1/\pi$ that fixes the cosmic matter inventory, applied to the mass-7 flux as an open-system port, predicts the lithium deficit — with the honest status of that conjecture, and the network-level obligations, marked exactly.

The intent throughout is to be scrupulous about status: standard BBN is a *Derived* consequence of the inherited expansion history and baryon ledger and nothing here disturbs it; the lithium mechanism is a *predicted but conditional* addition whose microscopic closure is an open item, the BBN-network implementation of PS-Cosmos-15. The reader should finish able to compute the helium fraction from scratch and to state precisely what about lithium is, and is not, resolved.

H.1 The neutron-to-proton ratio and weak freeze-out

Everything in BBN is downstream of one number: the ratio of free neutrons to free protons at the moment the light nuclei begin to assemble. We therefore start there.

In the first second after the bang the universe is a hot plasma in which neutrons and protons interconvert freely through the charged-current weak reactions

$$n + e^+ \longleftrightarrow p + \bar{\nu}_e, \tag{580}$$

$$n + \nu_e \longleftrightarrow p + e^-, \tag{581}$$

$$n \longleftrightarrow p + e^- + \bar{\nu}_e. \tag{582}$$

As long as these reactions proceed much faster than the universe expands, the two species sit in

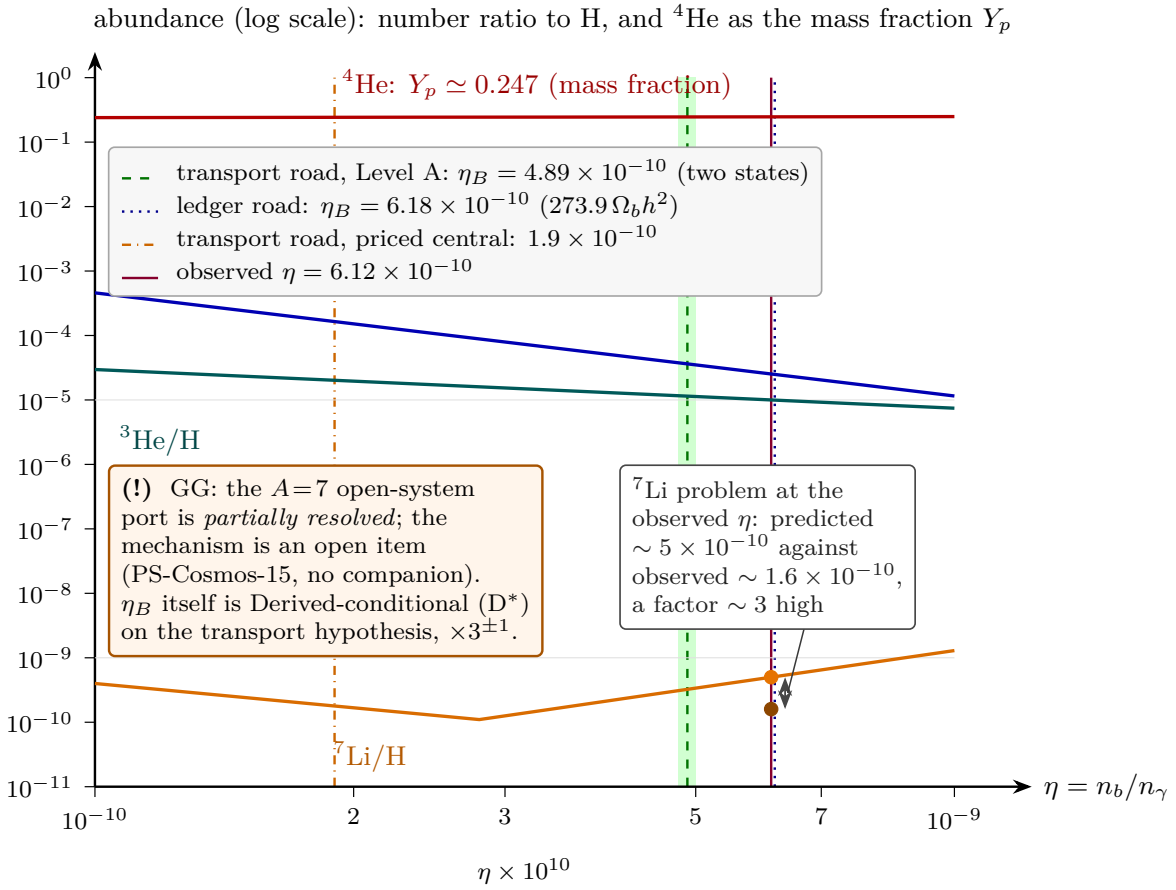

Figure 40: Standard big-bang nucleosynthesis (Schramm plot). Both axes are quantified and logarithmic: the abundances are number ratios to hydrogen, except ^{4}He, which is the mass fraction Y_p. Predicted primordial abundances vs the baryon-to-photon ratio η: the ^{4}He mass fraction $Y_p \sim 0.247$ is nearly flat (rising slightly), D/H falls steeply, ^{3}He/H is mild, and ^{7}Li/H traces the characteristic dip-then-rise valley near $\eta \sim 1$–5×10^{-10}. The vertical *green* band marks the transport road's Level A, $\eta_B = 4.89 \times 10^{-10}$ on two sourcing states (from the OGN-5 integer $K = 35$ via $e^{-K/2}$; the $\times 3^{\pm 1}$ band runs past the frame's upper edge), the dotted *blue* line the ledger road 6.18×10^{-10}, the dash-dotted *orange* line the transport road's priced central 1.9×10^{-10}, and the *purple* line the observed $\eta = 6.12 \times 10^{-10}$. **Lithium-7 problem (honest):** at the CMB η the predicted ^{7}Li/H $\sim 5 \times 10^{-10}$ exceeds the observed $\sim 1.6 \times 10^{-10}$ by a factor ~ 3 (both points drawn). **GG reading:** the proposed $A=7$ open-system port is "Partially resolved; mechanism an open item (PS-Cosmos-15)"; η_B itself is Derived-conditional (D*) on the transport hypothesis, with a $\times 3^{\pm 1}$ band. **Deuterium read off the curves:** at the ledger road's $\eta_{10} = 6.18$ (Route I), D/H is 0.98× the measured value; at the transport road's two-state Level A, 4.89, it is 1.43×; at the transport road's priced central value, 1.9, it is 6.5×. Deuterium therefore passes the ledger road at two percent and excludes the transport road's central value outright — and that excluded value lies inside the transport band, near its lower edge ($4.89/3 = 1.6$): the band as drawn contains a value deuterium rules out. The figure teaches standard BBN fully and marks the ^{7}Li discrepancy honestly.

mutual chemical equilibrium. The proton is lighter than the neutron by

$$Q \equiv (m_n - m_p)c^2 = 1.293 \text{ MeV}, \tag{583}$$

so equilibrium statistics (a Maxwell–Boltzmann / detailed-balance ratio of the two non-relativistic species at common temperature T, with their internal degeneracies equal) give the textbook Boltzmann suppression

$$\left(\frac{n}{p}\right)_{\text{eq}}(T) = \exp\!\left(-\frac{Q}{k_B T}\right). \tag{584}$$

We use natural units in which $k_B = 1$ and quote temperatures in MeV hereafter, so Q/T is just a ratio of two numbers in MeV. At very high temperature $T \gg Q$ the exponent is negligible and $n/p \to 1$: neutrons and protons are equally abundant. As the universe cools, the exponential turns on and neutrons become progressively rarer than protons because it costs energy Q to make one.

Equilibrium (584) holds only while the weak rates win against the Hubble expansion. The relevant comparison is between the per-neutron conversion rate $\Gamma_{\text{weak}}(T)$ and the expansion rate $H(T)$. The weak rate falls steeply with temperature, $\Gamma_{\text{weak}} \sim G_F^2 T^5$, because the cross sections scale as $G_F^2 E^2$ and the available energy scales with T. The expansion rate in the radiation era scales as

$$H(T) = \sqrt{\frac{8\pi G}{3}\,\rho} \;\propto\; \sqrt{g_*}\,\frac{T^2}{M_{\text{Pl}}}, \tag{585}$$

where g_* counts the relativistic degrees of freedom and M_{Pl} is the Planck mass. Because $\Gamma_{\text{weak}} \propto T^5$ falls faster than $H \propto T^2$, there is a crossover temperature T_f defined by

$$\Gamma_{\text{weak}}(T_f) \simeq H(T_f), \tag{586}$$

below which the weak interactions can no longer keep pace with the expansion. Inserting the measured Fermi constant and the radiation-era $H(T)$ puts this crossover at

$$T_f \approx 0.8 \text{ MeV}. \tag{587}$$

This is *weak freeze-out*: below T_f the reactions (580)–(581) effectively switch off and the equilibrium relation (584) stops being enforced. The n/p ratio is no longer free to track its equilibrium value; it "freezes" near whatever value equilibrium had reached at T_f:

$$\left(\frac{n}{p}\right)_f = \exp\!\left(-\frac{Q}{T_f}\right) = \exp\!\left(-\frac{1.293}{0.8}\right) \approx 0.20 \;\approx\; \frac{1}{5}. \tag{588}$$

This frozen ratio of about one neutron for every five protons is the single most important number in the whole calculation.

The word "freeze" is not quite literal, and the difference matters for the final helium yield. The free neutron is unstable: through reaction (582) it beta-decays with a mean lifetime

$$\tau_n = 879 \text{ s}. \tag{589}$$

After freeze-out, then, the proton number is essentially fixed but the neutron number continues to leak away by free decay during the long wait before nucleosynthesis can begin. The ratio at the actual onset of nucleosynthesis, at a later time $t_{\rm nuc}$, is therefore the frozen value damped by the surviving fraction of neutrons,

$$\frac{n}{p}\Big|_{\rm nuc} = \left(\frac{n}{p}\right)_f \exp\!\left(-\frac{t_{\rm nuc}}{\tau_n}\right). \tag{590}$$

We now have to find out how long that wait is, and the answer is the content of the next subsection.

H.2 The deuterium bottleneck

A naive estimate would have nucleosynthesis start the moment it is energetically allowed to. The first step on every path to heavier nuclei is the formation of deuterium,

$$n + p \longrightarrow d + \gamma, \qquad B_d = 2.22 \text{ MeV}, \tag{591}$$

where B_d is the deuteron binding energy. One might guess that deuterium accumulates as soon as T drops below $B_d \approx 2.2$ MeV. It does not, and the reason is the central subtlety of BBN.

The obstruction is that the universe is enormously *photon-rich*. The number of photons exceeds the number of baryons by the inverse of the baryon-to-photon ratio, $\eta \sim 6 \times 10^{-10}$, i.e. there are roughly 1.6×10^9 photons per baryon. Even when the *average* photon energy has fallen well below B_d, the high-energy tail of the blackbody distribution still contains, per deuteron, an overwhelming number of photons energetic enough to photodissociate it back to $n + p$. Deuterium is destroyed as fast as it forms. Quantitatively, the photodissociation reaction $d + \gamma \to n + p$ only becomes ineffective once the number of tail photons above B_d drops below the number of baryons, which requires T to fall not to B_d but to roughly

$$T_{\rm nuc} \approx \frac{B_d}{\ln(1/\eta)} \approx \frac{2.22 \text{ MeV}}{\ln(1.6 \times 10^9)} = \frac{2.22 \text{ MeV}}{21.2} \approx 0.10 \text{ MeV}. \tag{592}$$

This bare logarithmic estimate gives the right order of magnitude (~ 0.1 MeV); the full deuterium-bottleneck condition — which also follows the photodissociation cross-section and the falling baryon density — brings the operative onset down to the standard value $T_{\rm nuc} \approx 0.07$ MeV $\approx 0.8 \times 10^9$ K. The factor $\ln(1/\eta) \approx 21$ is the entire story: it is the logarithm of the photon-to-baryon excess, and it is what pushes the start of nucleosynthesis from $T \approx 2$ MeV down by a factor of about thirty to $T \approx 0.07$ MeV. This delay is the *deuterium bottleneck*. Because deuterium is the gateway nucleus on every reaction path, nothing downstream of it — no ^{3}He, no ^{4}He, no ^{7}Li — can form in quantity until the deuterium "logjam" breaks at $T_{\rm nuc}$.

In cosmic time, freeze-out at $T_f \approx 0.8$ MeV occurs at about $t \approx 1$ s, whereas the bottleneck breaks at $T_{\rm nuc} \approx 0.07$ MeV at about $t_{\rm nuc} \approx 270$ s (the network opens at 80 keV, $t \approx 200$ s). It is precisely during this several-minute interval that the neutron-decay damping of Eq. (590) acts, and it is why the helium yield ends up sensitive to the neutron lifetime τ_n. The bottleneck is therefore not a nuisance but a clock: it sets the wait during which neutrons leak away before they can be

locked into nuclei.

H.3 The reaction network: D, He-3, He-4, and Li-7/Be-7

Once T falls through $T_{\rm nuc}$ and the deuterium bottleneck breaks, the available neutrons and protons burn rapidly up the light nuclear network. The reactions proceed essentially to completion within a few minutes because the densities are still high and the Coulomb barriers of the lightest nuclei are low. The principal chain is

$$n + p \to d + \gamma, \tag{593}$$

$$d + d \to {}^3\mathrm{He} + n, \qquad d + d \to t + p, \tag{594}$$

$$d + {}^3\mathrm{He} \to {}^4\mathrm{He} + p, \qquad d + t \to {}^4\mathrm{He} + n. \tag{595}$$

The flow piles up at ^{4}He because the helium-4 nucleus is exceptionally tightly bound ($B/A \approx 7.07$ MeV per nucleon) and there is no stable nucleus at mass number $A = 5$ or $A = 8$. The two missing mass numbers act as a wall: the network cannot easily bridge past helium-4, so the overwhelming majority of all neutrons present at $T_{\rm nuc}$ end up bound inside ^{4}He, and only trace amounts of anything heavier are made.

The helium mass fraction. Because ^{4}He is the sink, its abundance can be estimated by a simple bookkeeping argument that needs no detailed rate integration. Each helium-4 nucleus contains two neutrons and two protons. To good approximation *every* surviving neutron is captured into helium, and it brings one proton with it. If there are n neutrons and p protons at the onset of nucleosynthesis, the number of helium nuclei formed is $n/2$ (two neutrons each), consuming n neutrons and n protons; the remaining $p - n$ protons stay as free hydrogen. The helium *mass* fraction, with helium of mass ≈ 4 and the total baryonic mass $\approx n + p$, is then

$$Y_p \equiv \frac{4\,(n/2)}{n+p} = \frac{2n}{n+p} = \frac{2\,(n/p)}{1+(n/p)}. \tag{596}$$

This compact formula is the headline BBN prediction: the primordial helium mass fraction is fixed almost entirely by the neutron-to-proton ratio at nucleosynthesis, which we computed in Sec. H.1. We evaluate it explicitly in Sec. H.6.

The mass-7 corner. Beyond helium, the network makes trace amounts of ^{3}He and deuterium that survive incomplete burning, and a much smaller trickle into mass number $A = 7$. The $A = 7$ abundance is produced by two competing routes whose balance depends on the baryon density:

$$t + {}^4\mathrm{He} \to {}^7\mathrm{Li} + \gamma \qquad \text{(direct } {}^7\mathrm{Li}\text{, dominant at low } \eta), \tag{597}$$

$${}^3\mathrm{He} + {}^4\mathrm{He} \to {}^7\mathrm{Be} + \gamma, \quad {}^7\mathrm{Be} + e^- \to {}^7\mathrm{Li} + \nu_e \qquad \text{(via } {}^7\mathrm{Be}\text{, dominant at high } \eta). \tag{598}$$

The key point for what follows is that at the observationally favored, relatively high value of η, the mass-7 abundance is dominated not by direct lithium production but by beryllium-7, which

is made first and only later electron-captures to lithium-7 (much of it after BBN proper, in the recombination era). The observed primordial "lithium" is therefore largely *decayed beryllium*. This is why the standard shorthand for the problem is the ^{7}Li/^{7}Be abundance: any mechanism that wishes to reduce primordial lithium without wrecking the other elements must act on mass number 7, and in practice must act on the ^{7}Be that dominates it. We return to this in Sec. H.7.

H.4 Dependence on the baryon-to-photon ratio

BBN has, to leading order, exactly one free input: the baryon-to-photon ratio

$$\eta \equiv \frac{n_B}{n_\gamma} \simeq 6.1 \times 10^{-10}, \tag{599}$$

equivalently the baryon density $\Omega_b h^2 \approx 0.0224$ (the two are related by $\eta \approx 2.74 \times 10^{-8}\,\Omega_b h^2$). Every primordial abundance is a definite function of this one number, which is what makes BBN a sharp, falsifiable audit rather than a fitting exercise. The qualitative dependences are physically transparent:

- **Deuterium** D/H is a steeply *decreasing* function of η: more baryons means denser plasma, faster burning, and more complete conversion of the fragile deuterium into helium, so less deuterium survives. This steep slope makes D/H the best "baryometer."
- **Helium** Y_p depends only weakly (logarithmically) on η, because it is set by the n/p ratio (Eq. 596) and η enters only through the modest shift in $T_{\rm nuc}$ and hence in the neutron-decay damping.
- **Lithium-7** ^{7}Li/H is non-monotonic in η: it has a minimum at intermediate η and rises on both sides because of the two competing production routes (597)–(598). At the high η preferred by the CMB, the ^{7}Be route dominates and lithium *increases* with η.

The decisive modern fact is that η is no longer a BBN-internal fit parameter. The cosmic microwave background fixes $\Omega_b h^2$ (and hence η) independently and to sub-percent precision from the relative heights of the acoustic peaks. BBN must therefore work at the CMB-determined η, with no remaining freedom. In the inherited ledger of this cohort the baryon fraction is carried as the π-locked value $\Omega_b = 1/(2\pi^2)$ (see the topological-inventory Part of the main text), so η is fixed upstream and BBN is run as a closed audit against it, not tuned.

H.5 Observational status: triumph and the lithium-7 problem

At the CMB-fixed η, two of the three measurable primordial abundances are textbook successes of cosmology:

- **Deuterium.** The predicted D/H and the value measured in pristine, metal-poor, high-redshift absorption systems along quasar sightlines agree at the level of a few percent, D/H $\approx 2.5\times10^{-5}$. Because D/H is the steep baryometer, this agreement is a stringent and independent confirmation of the CMB baryon density — and, read against this paper's own roads, of the

ledger road: at $\eta_{10} = 273.9\,\Omega_b h^2 = 6.18$ the framework's deuterium is 0.98× the measured value, while the transport road's priced central 1.9 gives 6.5× and is excluded by it (Figure 40).

- **Helium-4.** The predicted $Y_p \approx 0.247$ and the value extrapolated from helium and hydrogen emission lines in metal-poor extragalactic H II regions, $Y_p \approx 0.245 \pm 0.003$, agree well within the uncertainties.

These two agreements, spanning nine orders of magnitude in abundance and resting on independent physics (weak freeze-out for helium, the burning network for deuterium), are among the strongest quantitative evidence we have that the universe was once hot, dense, and described by known microphysics back to $t \approx 1$ s.

The third measurable abundance does *not* agree, and this is the *lithium-7 problem*. Standard BBN at the CMB-fixed η predicts a primordial lithium abundance of about

$$\left(\frac{^7\mathrm{Li}}{\mathrm{H}}\right)_{\mathrm{BBN}} \approx (4.5\text{–}5.0) \times 10^{-10}, \tag{600}$$

whereas the abundance inferred from the surfaces of old, metal-poor halo stars — where it sits on the remarkably flat "Spite plateau," nearly independent of metallicity and temperature — is

$$\left(\frac{^7\mathrm{Li}}{\mathrm{H}}\right)_{\mathrm{Spite}} \approx (1.5\text{–}1.6) \times 10^{-10}. \tag{601}$$

The prediction overshoots the observation by a factor of roughly three (about 4–5σ given the quoted errors). This is a genuine, long-standing discrepancy. Three broad classes of explanation are debated and none is established: (i) *stellar depletion*, in which lithium is destroyed over the stars' lifetimes by mixing below the convection zone, so the plateau understates the true primordial value — the flatness of the plateau is, however, hard to reconcile with the spread that depletion models tend to predict; (ii) *nuclear-rate revisions*, in which a mismeasured ^{7}Be-producing or ^{7}Be-destroying cross section shifts the prediction — but the relevant reactions have been re-measured and the rates are now fairly well pinned; and (iii) *new physics* during or just after BBN that selectively destroys mass-7 nuclei. The GG reading proposed below sits in the third class, and the honesty discipline of this appendix is to label it as exactly that: a proposal, not a resolution.

H.6 Worked example: n/p at freeze-out and the helium fraction

We now carry the helium calculation through explicitly, because it is the cleanest demonstration that BBN is genuinely *Derived* from the inherited expansion history with no tunable knob.

Step 1 — equilibrium ratio at freeze-out. Using Eq. (584) with $Q = 1.293$ MeV and the freeze-out temperature $T_f = 0.8$ MeV from Eq. (587),

$$\left(\frac{n}{p}\right)_f = \exp\!\left(-\frac{1.293}{0.8}\right) = \exp(-1.616) = 0.199 \approx \frac{1}{5}. \tag{602}$$

Step 2 — neutron decay during the bottleneck. Nucleosynthesis does not begin until the deuterium

bottleneck breaks at $t_{\rm nuc} \approx 300$ s (Sec. H.2). During the interval from freeze-out to $t_{\rm nuc}$ the free neutrons decay with lifetime $\tau_n = 879$ s. Applying the damping factor of Eq. (590),

$$\frac{n}{p}\Big|_{\rm nuc} = 0.199 \times \exp\!\left(-\frac{300}{879}\right) = 0.199 \times \exp(-0.341) = 0.199 \times 0.711 = 0.142 \approx \frac{1}{7}. \tag{603}$$

The decay has driven the ratio from about 1/5 down to about 1/7: roughly one neutron survives for every seven protons when the network finally fires. This factor of 1/7 is the famous BBN benchmark.

Step 3 — helium mass fraction. Inserting $n/p = 1/7$ into the bookkeeping formula Eq. (596),

$$Y_p = \frac{2\,(n/p)}{1 + (n/p)} = \frac{2 \times \frac{1}{7}}{1 + \frac{1}{7}} = \frac{2/7}{8/7} = \frac{2}{8} = 0.25. \tag{604}$$

$$\boxed{Y_p \approx 0.25} \tag{605}$$

A primordial helium mass fraction of one-quarter, derived from nothing but $Q = 1.293$ MeV, the weak-freeze-out temperature, the neutron lifetime, and the bottleneck delay, and agreeing with the measured $Y_p \approx 0.245$. Two remarks underline how unforgiving this is. First, had the neutrons *not* decayed ($n/p = 0.20$ throughout), Eq. (596) would give $Y_p = 2(0.20)/1.20 = 0.33$, a third of the baryonic mass in helium — the ~ 0.25 value is what it is only because the bottleneck delay lets neutrons leak away first. Second, no parameter here is free to adjust: Q, τ_n, and G_F are laboratory constants, and η is fixed by the CMB. The agreement is therefore a prediction, not a fit. After all, a number this rigid is exactly the kind of audit a geometric origin theory must pass before it is allowed to say anything new.

H.7 The geometric derivation: the lithium-7 deficit as a closed-packet $1/\pi$ port

With the baseline fully in view — two triumphs and one genuine failure — we can state what the geometry adds. The two successful abundances, D/H and Y_p, are left *untouched*: the inherited expansion history and the π-locked baryon ledger $\Omega_b = 1/(2\pi^2)$ reproduce them exactly as standard BBN does, and the program claims no modification there. The lithium-7 problem is the one place the standard prediction fails, and it is the one place the geometry speaks — with the same closed-packet factor $1/\pi$ that fixes the cosmic matter inventory.

Why the boundary universe has an open channel. The observable four-dimensional universe is never a closed microstate. By OGGEP it is an open boundary subsystem, accessed through a non-invertible bulk-to-boundary projection, and the generic signature of that projection is a dissipative GDOS channel that can drain a registered species into the inaccessible bulk ledger. A selective $A = 7$ port is exactly such a channel acting on mass number 7 — in practice on the ^{7}Be that dominates the $A = 7$ abundance at the CMB-favored η (§H.3). *Physics intuition:* an open boundary always leaks a little of what it carries; the only question is how much, and onto what.

The surviving fraction is $1/\pi$, not a fitted knob. The strength of the port is not a free coefficient. Modeling the mass-7 flux as a System Instance Packet carried over the compact angle

$\theta \in S^1_\theta$, the open-system drift retains, on projection to the observable boundary, exactly the closed-packet share $1/\pi$ — the *same* factor that yields $\Omega_m = 1/\pi$. The full derivation is the companion port calculation of §I, Eqs. (637)–(638); the surviving fraction $S_{A=7} = 1/\pi$ follows because the closed-packet geometry branches the conserved flux over the compact angular measure in precisely the way it branches the cosmic inventory. The predicted primordial lithium is therefore

$$\left(\frac{^7\mathrm{Li}}{\mathrm{H}}\right)_{\mathrm{GG}} = \frac{1}{\pi}\left(\frac{^7\mathrm{Li}}{\mathrm{H}}\right)_{\mathrm{SBBN}} \simeq \frac{5.0\times10^{-10}}{\pi} \simeq 1.6\times10^{-10}, \tag{606}$$

to be compared with the Spite-plateau value $(1.5\text{–}1.6)\times10^{-10}$ of (601): the factor-of-three overshoot is removed, and D/H and Y_p are untouched because the port is selective in A. *Physics intuition:* the same projection that keeps only $1/\pi$ of the closed packet visible as matter keeps only $1/\pi$ of the mass-7 flux as lithium — the lithium deficit and the matter inventory are two readings of one geometric branching.

What is, and is not, established. This is a prediction, not a fit: the port strength is the closed-packet $1/\pi$, with no free coefficient. But it inherits exactly the status of the inventory that supplies it — *Derived-conditional* (D^*), resting on the closed-packet-integration conjecture (Conjecture VIII) — and three obligations remain open, stated plainly:

1. **The selectivity mechanism.** Why the port acts on $A = 7$ and not on $A = 2, 3, 4$ — how the selectivity in mass number arises from the bulk projection — is not delivered here.

2. **Network-level closure.** A full coupled BBN-network calculation, re-running the reaction network of §H.3 with the port inserted and confirming the other abundances genuinely do not move, is owed.

3. **Detailed balance.** Compatibility with detailed balance and the well-measured $A = 7$ cross sections ($^3\mathrm{He} + {}^4\mathrm{He} \to {}^7\mathrm{Be}$ production and $^7\mathrm{Be}$ destruction) must be demonstrated.

These are the open item PS-Cosmos-15, the BBN-network implementation; no companion in the cohort carries it. The honest one-line summary: the lithium-7 problem is addressed by a *parameter-free* closed-packet port that lands on the Spite plateau, conditional on Conjecture VIII and on the network-level obligations above — a predicted, partially closed mechanism, falsifiable by any firm primordial ^{7}Li away from SBBN/π, and not yet a finished resolution. In the Euclidean spirit of the program: the part the geometry fixes (the $1/\pi$ survival) is stated as a result, and the parts it does not yet fix are marked, in plain sight, as the work that remains.

Status

- **Standard BBN — Derived.** The neutron-to-proton freeze-out ratio, the deuterium bottleneck, the D/^{3}He/^{4}He network, the helium mass fraction $Y_p \approx 0.25$ (§H.6), and the η-dependence of the abundances all follow from the inherited radiation-era expansion history and the CMB-fixed, π-locked baryon ledger $\Omega_b = 1/(2\pi^2)$, with no free parameter. The agreement with the measured D/H and Y_p is reproduced, not modified.

- **The $A = 7$ lithium port — Derived-conditional (D^*).** The selective $A = 7$ GDOS open-system port reduces the predicted ^{7}Li/^{7}Be abundance to the Spite-plateau value by the closed-packet factor $1/\pi$ — the same factor that fixes $\Omega_m = 1/\pi$ — with *no free coefficient* (§I). Its status is therefore that of the inventory: D^*, conditional on the closed-packet-integration conjecture (Conjecture VIII). The microscopic origin of the mass-7 selectivity, the full coupled BBN-network closure, and compatibility with detailed balance and the measured $A = 7$ cross sections are open obligations (PS-Cosmos-15), carried by no companion. The lithium-7 problem is reported here as a predicted, partially closed mechanism, not as solved.

Problems and worked solutions. H.2.1. Estimate the temperature at which the deuterium bottleneck breaks.

Solution. Step 1: deuterium (binding energy $B_D = 2.22$ MeV) is photodissociated as long as photons above B_D outnumber baryons; with $n_\gamma/n_B = 1/\eta$, the Boltzmann tail condition is $e^{-B_D/T}/\eta \gtrsim 1$. Step 2: solve: $T_\mathrm{D} \approx B_D/\ln(1/\eta) = 2.22/21.2 \approx 0.1$ MeV — twenty times below the binding energy. Step 3: only below T_D can the reaction network of the chapter fire, which is why nucleosynthesis happens at $t \sim$ minutes rather than seconds, giving free neutrons time to decay (lowering n/p toward 1/7). Physics intuition: the bottleneck is the photon bath's veto power — a billion photons per baryon means even the one-in-a-billion tail can dissociate every nucleus; the $\ln(1/\eta)$ delay is F.8.1's entropy-per-baryon speaking in nuclear language.

H.4.1. Explain, with the network's logic, why deuterium falls steeply with η while helium barely moves.

Solution. Step 1: helium-4 is a *thermodynamic* endpoint — essentially every neutron alive at the bottleneck ends inside ^{4}He, so Y_p depends on η only through the mild shift of the bottleneck temperature (logarithmic, H.2.1) — weak dependence. Step 2: deuterium is a *kinetic* survivor — the trace that failed to burn; more baryons mean faster burning and a smaller survivor fraction, a steep power law ($\mathrm{D/H} \propto \eta^{-1.6}$). Step 3: hence the division of labor: helium tests the expansion rate (H.6.1), deuterium measures η — the "baryometer" — and the measured D/H fixes $\eta \approx 6 \times 10^{-10}$, independently confirmed by the CMB. Physics intuition: endpoints make poor meters and survivors make sharp ones — the abundance of what *almost* happened always carries more information than the abundance of what was inevitable.

H.6.1. How would a fourth neutrino species shift the helium abundance?

Solution. Step 1: an extra relativistic species raises g_* at MeV temperatures, hence raises $H \propto \sqrt{g_*}T^2$. Step 2: faster expansion means weak freeze-out (the chapter's worked example) happens *earlier*, at higher T_f: the n/p ratio freezes larger, and more neutrons survive into helium. Step 3: the standard sensitivity is $\Delta Y_p \approx 0.013\,\Delta N_\nu$ — a fourth neutrino would push Y_p from ≈ 0.247 to ≈ 0.260, outside the observed range: BBN counted the neutrino families ($N_\nu \approx 3$) two decades before the colliders did. Physics intuition: the early universe is a speedometer — anything that adds energy density accelerates the expansion, and the helium abundance is the tachograph that recorded it; this is also why BBN polices every new light particle the cohort or anyone else proposes.

H.5.1. State the lithium-7 problem as one number, and notice the number.

Solution. Step 1: standard BBN with the CMB-fixed η predicts ^{7}Li/H $\approx 5.0 \times 10^{-10}$. Step 2: the Spite plateau of metal-poor halo stars observes $\approx 1.6 \times 10^{-10}$. Step 3: the discrepancy factor is $5.0/1.6 = 3.13$ — within one percent of π. Standard resolutions (stellar depletion, nuclear rates) each strain against independent data; the geometric reading of the next chapter's port section derives the factor as exactly π, from the same closed-packet branching that fixes $\Omega_m = 1/\pi$. Physics intuition: a thirty-year anomaly whose magnitude is a mathematical constant is either a coincidence or a signature — the falsification catalog treats it as the latter and prices the kill: tighten the plateau and the nuclear inputs, and the factor either sharpens toward π or walks away from it.

Take-away. *Standard result:* BBN predicts the light-element abundances from one parameter (η_B) and nuclear physics — cosmology's earliest quantitative triumph — but overpredicts lithium-7 by a factor ~ 3. *What it assumed:* a closed reaction network and η_B measured elsewhere. *What the geometry replaced it with:* η_B predicted by the ledger road of Part III and checked against a leptogenesis mechanism on the inherited ledger (Appendix I), and the $A = 7$ sector read as an open-system port with a $1/\pi$ suppression — falsifiable, because any lithium fix that damages D/H or Y_p is excluded. *Continue in:* Part V; the decade table of Part XIII.

I Resonant leptogenesis and the lithium-7 port

What this appendix teaches. Where the baryon asymmetry $\eta_B \sim 6 \times 10^{-10}$ comes from: the type-I seesaw matched to the inherited GG-A9-SM neutrino block, the resonant CP asymmetry on the inherited phases, washout, and the assembled η_B — followed by the geometric sharpening: the derived reheating temperature $T_{\rm rh} \simeq 6 \times 10^{10}$ GeV sits seven orders above the TeV-scale heavy states, so the bath exists long before they matter; they reach equilibrium through their own inverse decays near 10–20 TeV ($K_{\rm wash} \simeq 47$), whatever the reheating temperature above that, and the resonant mechanism runs on a full bath — guaranteed by the launch, not chosen.

Read alongside: Kolb & Turner (1990) ch. 6 for baryogenesis generally; Fukugita & Yanagida (1986) for the original mechanism.

This appendix is the geometric derivation behind the origin of matter. The number standard Big-Bang nucleosynthesis (Appendix H) had to take as a measured input — the baryon-to-photon ratio η_B — is *predicted* by the ledger road of Part III, $\eta_{10} = 273.9\,\Omega_b h^2 = 6.18$; what this appendix supplies is the *mechanism* checked against that target, resonant leptogenesis on the inherited ledger, which reaches it to a factor of about three from below, and the closed-packet $1/\pi$ port that addresses the lithium-7 deficit. Where Appendix H asked what the data require, this appendix asks what the inherited geometry forces, in the same Euclidean spirit as Appendix E: the heavy Majorana scale, the neutrino Yukawas, and the CP phases are not free, but fixed by GG-A9-SM and the lock-scale geometry. It carries the heavy algebra the Part V matter sector points to: the type-I seesaw matching that fixes the neutrino Yukawas, the resonant CP-asymmetry formula evaluated on the inherited phases, the Boltzmann/washout estimate, the assembly of η_B, and the open-system $A = 7$ port whose surviving fraction is the closed-packet factor $1/\pi$. Each block follows the same order: statement, derivation, worked number, geometric origin, honest status.

I.1 Type-I seesaw matching to the GG-A9-SM Weinberg operator

Statement. GG-A9-SM fixes the light neutrino masses through the dimension-five Weinberg operator at the GUT seam,

$$\mathcal{L}_5 = \frac{(LH)(LH)}{\Lambda_{\rm W}}, \qquad m_\star = \frac{\pi v^2}{E_{\rm GUT}}, \qquad v \simeq 246\,\text{GeV}, \tag{607}$$

and explicitly leaves the ultraviolet completion to GG-A10-Cosmos. The completion is a heavy Majorana triplet N_i at the lock scale; integrating it out must reproduce Eq. (607).

Derivation. With Dirac Yukawas Y_ν coupling LHN and Majorana masses M_N, the seesaw light mass matrix is

$$m_\nu = -Y_\nu^{\mathsf{T}} M_N^{-1} Y_\nu\, v^2, \tag{608}$$

whose magnitude for a representative entry is $m_\nu \sim Y_\nu^2 v^2/M_N$. Matching this to $m_\star = \pi v^2/E_{\rm GUT}$ cancels v^2:

$$\frac{Y_\nu^2}{M_N} \simeq \frac{\pi}{E_{\rm GUT}} \quad\Longrightarrow\quad Y_\nu^2 \simeq \pi\,\frac{M_N}{E_{\rm GUT}}. \tag{609}$$

Worked number. With $M_N \simeq M_{\rm lock} \simeq 3.85\times10^3\,\text{GeV}$ and $E_{\rm GUT} \simeq 2.28\times10^{16}\,\text{GeV}$,

$$Y_\nu^2 \simeq \pi\,\frac{3.85\times10^3}{2.28\times10^{16}} \simeq 5.3\times10^{-13}, \qquad Y_\nu \simeq 7.3\times10^{-7}. \tag{610}$$

The tree-level decay width is

$$\Gamma_N = \frac{Y_\nu^2}{8\pi} M_N \simeq \frac{5.3\times10^{-13}}{8\pi}\,(3.85\times10^3\,\text{GeV}) \simeq 8.1\times10^{-11}\,\text{GeV}, \qquad \frac{\Gamma_N}{M_N} \simeq 2.1\times10^{-14}. \tag{611}$$

One bookkeeping convention is declared here: this width is quoted at the seam-matched one-generation Yukawa, which corresponds to $\widetilde{m} \simeq 8.4$ meV, while the washout below runs at $\widetilde{m} \sim m_3 \simeq 50.5$ meV — conventions differing by a factor $\simeq 12$ in Y_ν^2. Run at the washout convention the width is $\Gamma_N \simeq 1\times10^{-9}$ GeV: an $O(10)$ shift inside the order-of-magnitude register, absorbed into the effective width $\Gamma^{\rm eff}$ that the hypothesis H-$\phi_{\rm geo}$-1 below parameterizes. The two conventions are deliberate, and each is evaluated where its physics lives: the resonance width is a property of the seam-localized Yukawa texture, so it is quoted at the seam-matched one-generation magnitude $Y_\nu^2 \simeq \pi M_N/E_{\rm GUT}$ of Eq. (609), whose effective mass is the Majorana seam scale itself, $\widetilde{m} = Y_\nu^2 v^2/M_N = \pi v^2/E_{\rm GUT} \simeq 8.4$ meV; the washout strength is a property of the light-neutrino mass scale the thermal bath actually sees, which the inherited spectrum fixes at the heaviest light mass, $\widetilde{m} \sim m_3 \simeq 50.5$ meV. The factor between the conventions is thereby accounted for, not mysterious. The effective masses differ by $m_3/(\pi v^2/E_{\rm GUT}) \simeq 50.5/8.4 \simeq 6.0$; the quoted Yukawa-squared ratio carries, in addition, only the Casas–Ibarra half in $m_\nu = Y^2 v^2/2M_N$ (the normalization of the resonant subsection below, Eq. (615)) relative to the seam magnitude matching $m_\nu \sim Y_\nu^2 v^2/M_N$ of Eq. (609), so that $(Y_3/Y_\nu)^2 = 2\,m_3/\widetilde{m}_{\rm seam} \simeq 12$ (numerically, $Y_3 \simeq 2.5\times10^{-6}$ against $Y_\nu \simeq 7.3\times10^{-7}$). The physical spread between the conventions is the single mass ratio $\simeq 6.0$; the factor of two is normalization

bookkeeping, not an additional physical spread; both readings sit inside the $O(1)$–$O(10)$ pricing already declared, each downstream quantity is run on the convention proper to it, and the entire effect on the throughput is carried by the same effective width $\Gamma^{\rm eff}$.

GG connection. The tiny Yukawa is forced, not chosen: a sub-eV light mass with a TeV heavy scale *requires* $Y_\nu \sim 10^{-6}$. Placing the seesaw at the geometry's lock scale, rather than at an arbitrary intermediate scale, is what makes this the case.

Status. Derived. Given GG-A9-SM's $m_\star$ and the geometry's $M_N \simeq M_{\rm lock}$, the Yukawa magnitude is fixed with no freedom.

I.2 Resonant CP asymmetry on the inherited phases

Statement. The self-energy (resonant) CP asymmetry for a quasi-degenerate Majorana pair is

$$\epsilon_i = \frac{\Im\Big[(Y_\nu Y_\nu^\dagger)_{ij}^2\Big]}{(Y_\nu Y_\nu^\dagger)_{ii}(Y_\nu Y_\nu^\dagger)_{jj}} \, \frac{(M_i^2 - M_j^2) M_i \Gamma_j}{(M_i^2 - M_j^2)^2 + M_i^2 \Gamma_j^2}. \tag{612}$$

Derivation. Write $M_j = M_i(1+d)$ with fractional splitting $d \ll 1$. Then $M_i^2 - M_j^2 \simeq -2M_i^2 d$, and the Breit–Wigner factor becomes

$$\mathcal{R}_{\rm res} = \frac{(M_i^2 - M_j^2) M_i \Gamma_j}{(M_i^2 - M_j^2)^2 + M_i^2 \Gamma_j^2} \simeq \frac{-2d\,(\Gamma_j/M_j)}{4d^2 + (\Gamma_j/M_j)^2}. \tag{613}$$

This is extremized when $2|d| = \Gamma_j/M_j$, i.e. when the mass splitting equals half the width (the resonance condition), where $|\mathcal{R}_{\rm res}| = \frac{1}{2}$. At resonance the common magnitude of $(Y_\nu Y_\nu^\dagger)$ cancels between the CP-odd prefactor's numerator and denominator, so ϵ is controlled by the relative phase alone:

$$\epsilon^{\rm res} \simeq \frac{\Im\Big[(Y_\nu Y_\nu^\dagger)_{ij}^2\Big]}{(Y_\nu Y_\nu^\dagger)_{ii}(Y_\nu Y_\nu^\dagger)_{jj}} \cdot \frac{1}{2} \equiv \frac{1}{2}\,\sin(2\phi_{\rm geo}). \tag{614}$$

The inherited light sector reaches $\phi_{\rm geo}$ only through flavor, and the route must be displayed rather than asserted; it is displayed as a keyed lemma, so that the downstream consumers of this result can cite it.

Lemma I.1 ($U_{\rm PMNS}$-cancellation and the flavor obligation (L-Cosmos-U))**.** *In the unflavored (flavor-summed) limit, the resonant CP invariant of Eq. (614) is independent of $U_{\rm PMNS}$, and for the dial-free door (R real, adding no new CP) it vanishes identically — the unflavored $\epsilon^{\rm res}$ is exactly zero (C-Cosmos-U). Hence the inherited phases reach η_B only through the flavored asymmetries $\epsilon_{i\alpha}$: at $M_N \simeq 3.85$ TeV, where all three charged-lepton flavors are resolved in the plasma, the fully flavored treatment is mandatory (L-Cosmos-F, the flavor obligation).*

Proof. In the Casas–Ibarra parametrization, $Y_\nu = (\sqrt{2}/v)\, M_N^{1/2} R\, m_\nu^{1/2} U_{\rm PMNS}^\dagger$ with R complex

orthogonal, the heavy-space bilinear entering Eq. (614) is

$$(Y_\nu Y_\nu^\dagger)_{ij} = \frac{2}{v^2}\sqrt{M_i M_j}\sum_k m_k\, R_{ik}R_{jk}^*, \tag{615}$$

in which $U_{\rm PMNS}$ has canceled exactly ($U^\dagger U = 1$): the unflavored invariant of Eq. (614) is independent of the PMNS phases, and for the dial-free door (R real, adding no new CP) the sum $\sum_k m_k R_{ik}R_{jk}$ is real, so its imaginary part vanishes identically — the unflavored $\epsilon^{\rm res}$ is exactly zero. At $M_N \simeq 3.85$ TeV all three charged-lepton flavors are resolved in the plasma, so the fully flavored treatment is mandatory, and there the PMNS phases genuinely enter: the operative objects are the flavored asymmetries $\epsilon_{i\alpha}$ (self-energy numerators $\Im[Y_{i\alpha}Y^*_{j\alpha}(Y_\nu Y_\nu^\dagger)_{ij}]$ and the ji partner), in which $U_{\rm PMNS}$ enters explicitly through $Y_{i\alpha}$; their flavor sum vanishes for real R, while a net asymmetry survives because the flavored washouts K_α differ. ■

Worked number (flavored, ledger inputs). On the GG-A9-SM QD ledger (masses, angles, $\delta_{\rm PMNS} = 196.06°$, $\alpha_{31} = \pi$; real R with $\mathcal{O}(1/K)$ angles) the flavored asymmetries are $\epsilon_\alpha \simeq (+3.4, +3.1, -6.6)\times 10^{-3}$ ($\sum_\alpha \epsilon_\alpha = 0$ exactly), the flavored washouts are $K_\alpha \simeq (3.4, 28.0, 22.4)$, and the effective throughput $\epsilon_{\rm eff} = \sum_\alpha \epsilon_\alpha \min(1, 1/K_\alpha)$ brackets the ledger target:

$$\epsilon_{\rm eff} \simeq 8\times 10^{-4}\ \text{(Breit–Wigner peak)}, \qquad \epsilon_{\rm eff} \simeq 6\times 10^{-15}\ \text{(vacuum splittings)}, \tag{616}$$

an overshoot of the required throughput $\epsilon^{\rm res}\kappa = e^{-K/2}$ by $\sim 3\times 10^4$ on the peak and an undershoot by $\sim 4\times 10^6$ at vacuum numbers. Three structural facts, each numerically checked: the driver is the Majorana parity $\alpha_{31} = \pi$, which sets the *sign* of the matter excess ($\alpha_{31} \to 0$ flips the sign, while $\delta_{\rm PMNS} \to \pi$, i.e. $J_{\rm PMNS} \to 0$, shifts $\epsilon_{\rm eff}$ by only +4%: the result is $J_{\rm PMNS}$-blind at the few-percent level); the result is nonzero at $m_1 = 0$ (first-order masses give $\epsilon_{\rm eff} \simeq 2.4\times 10^{-4}$), so the asymmetry does *not* ride on the QD layer; and the U-cancellation of Lemma I.1 (L-Cosmos-U) retires the formerly assembled evaluation $\sin\phi_{\rm geo} \sim c\, J_{\rm PMNS}$ with its m_1/m_3-order suppression, withdrawn as a reconstruction artifact.

The named hypothesis. *H-$\phi_{\rm geo}$-1 (partial resonance): there exists a finite-temperature transport computation of the lock-crossing decays in which the effective splitting-to-width ratio at the crossing satisfies $\Delta M/\Gamma^{\rm eff} \simeq 3\times 10^4$ (equivalently, Breit–Wigner factor $\simeq 1.5\times 10^{-5}$; $\Gamma^{\rm eff} \simeq 3\times 10^{-3}$ GeV at $\Delta M \simeq 110$ GeV), so that the flavored, real-R throughput on the inherited phases lands within $O(1)$ of $e^{-K/2}$.* The hypothesis stays named and conditional, and it is now carried by an identified candidate mechanism — the flavor-measurement regulator of §I.3, with the closed form $\Delta M/\Gamma^{\rm eff} = 1/(35\, c_\tau y_\tau^2) = (2.7\text{–}5.3)\times 10^4$ — at candidate grade, conditional on the conditions C1–C6 stated there. Under H-$\phi_{\rm geo}$-1 the chain's $\epsilon^{\rm res} \simeq 2.5\times 10^{-6}$ is the *target value* the partial resonance must deliver — a target value now carried by an identified candidate mechanism whose closed form lands the ratio in the hypothesized decade with zero dials (§I.3); the transport-level closure of conditions C1–C6 remains open. $\Gamma^{\rm eff} \sim 3$ MeV is parametrically plausible for thermal damping at $T \sim M_N$ — and is now identified with the computed $\Gamma_\tau = 2$–4 MeV — so the hypothesis is a well-posed computation with an acceptance band ($\pm O(1)$ in $\epsilon_{\rm eff}$), carried as row 9 of the promotion roadmap

of this paper's Foundations-Branch companion, the CPT member paper GG-P5-CPT, where the row now records the mechanism at candidate grade.

GG connection. The resonance lifts what would otherwise be a hopeless $Y_\nu^2 \sim 10^{-13}$ suppression, and flavor is what lets the inherited CP through: the sign of the matter excess is the Majorana parity of the rank-two block ($\alpha_{31} = \pi$, derived structural data of GG-A9-SM), filtered through the flavor-washout spread $K_e \ll K_{\mu,\tau}$ — an inheritance statement sharper, not weaker, than the retired Jarlskog projection.

Status. Derived-conditional on H-$\phi_{\rm geo}$-1. The flavored computation on the inherited phases is displayed above with no dial (Lemma I.1 and the worked number); the named hypothesis is the finite-temperature partial-resonance condition $\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$ — the sharp form of the on-resonance condition is partial resonance, strictly weaker than the exact-peak form (which overshoots the required throughput by $\sim 3 \times 10^4$) and strictly stronger than vacuum numbers (which undershoot by $\sim 4 \times 10^6$); discharging it is the flavored face of the Resonant-leptogenesis-normalization conjecture (Conjecture VI). The hypothesis now carries an identified candidate mechanism at candidate grade — the flavor-measurement regulator of the next subsection (§I.3, conditions C1–C6) — which does not change the grade of this subsection.

I.3 The partial-resonance mechanism: the flavor-measurement regulator

Statement. The named hypothesis of the previous subsection asks for a finite-temperature transport computation in which the effective splitting-to-width ratio at the lock crossing lands at $\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$. This subsection installs the candidate mechanism: with R real the CP source is purely flavor-odd (Lemma I.1), flavor-blind damping cancels from it exactly, and the effective regulator of the flavored partial resonance is the thermal flavor-measurement rate of the lepton coherences — the charged-lepton Yukawa equilibration rate, dominated by the tau,

$$\Gamma^{\rm eff} \;=\; \Gamma_\tau \;=\; c_\tau\, y_\tau^2\, T, \qquad c_\tau = 5 \times 10^{-3}\text{–}10^{-2}, \tag{617}$$

with c_τ the kinetic coefficient of the flavor-equilibration literature. At the generation temperature $T \simeq M_N$ the heavy mass cancels between the ledger splitting $\Delta M = M_N/35$ and the thermal rate, leaving a closed form in the ledger integer and a running Standard-Model Yukawa:

$$\frac{\Delta M}{\Gamma^{\rm eff}} \;=\; \frac{M_N/35}{c_\tau\, y_\tau^2\, T}\bigg|_{T=M_N} \;=\; \frac{1}{35\, c_\tau\, y_\tau^2} \;=\; (2.7\text{–}5.3) \times 10^4. \tag{618}$$

Zero dials: the only inputs are the ledger integer $K = 35$ and the running tau Yukawa. The identification is carried at *candidate grade*, conditional on the numbered conditions C1–C6 below; nothing in this subsection upgrades the grade of the chain.

Physics picture. The thermal bath is a flavor-measuring apparatus. Every process that distinguishes lepton flavors acts as a measurement of the flavor coherences the N decays write into the

plasma, and at $T \simeq 4$ TeV the fastest *honest* measurer is the tau Yukawa interaction — the gauge interactions are five orders stronger but flavor-blind, so they read nothing. The partial-resonance question, *how long does the CP-building interference live?*, is therefore answered not by the N's own width — the coherence that carries the CP does not live in the N lines — but by how quickly the bath measures lepton flavor: $\Gamma^{\rm eff} = \Gamma_\tau$.

Derivation, step 1: the flavored kinetic system. The operative object is the lepton-doublet flavor matrix $(\delta n_\ell)_{\alpha\beta}$ (deviation from equilibrium; $\alpha, \beta = e, \mu, \tau$). Its flavor-covariant kinetic equation, in the density-matrix form standard in flavored leptogenesis, is

$$\frac{\mathrm{d}(\delta n_\ell)_{\alpha\beta}}{\mathrm{d}t} = S_{\alpha\beta}(t) \;-\; \tfrac{1}{2}\{W,\, \delta n_\ell\}_{\alpha\beta} \;-\; \tfrac{1}{2}\,\Gamma_\tau\big[P_\tau, [P_\tau,\, \delta n_\ell]\big]_{\alpha\beta} \;-\; \Gamma_{\rm g}\,(\text{terms} \propto \mathbf{1}_{\alpha\beta}), \tag{619}$$

where $S_{\alpha\beta} \propto \sum_{j\neq i} Y_{i\alpha} Y^*_{j\beta} \times$ (the ij heavy coherence, oscillating at ΔM_{ij}) is the CP source from the quasi-degenerate pair, W the flavored washout ($\{\,,\}$ the anticommutator), $P_\tau = \mathrm{diag}(0,0,1)$ the tau projector, Γ_τ the tau-Yukawa flavor-measurement rate acting on tau-row and tau-column coherences, and the gauge damping $\Gamma_{\rm g}$ — flavor-universal — enters proportional to the identity in flavor space: it thermalizes kinetic distributions at fixed flavor content.

Derivation, step 2: with R real the source is purely flavor-odd. By Lemma I.1, the flavor trace of the CP source is the unflavored invariant, and for the dial-free door it vanishes identically:

$$\mathrm{Tr}\, S = \sum_\alpha S_{\alpha\alpha} \;\propto\; \Im\Big[(Y_\nu Y_\nu^\dagger)^2_{ij}\Big] = 0 \qquad (R \text{ real}). \tag{620}$$

The physical source is therefore the traceless — flavor-odd — part of $S_{\alpha\beta}$: the entire CP content lives in flavor *differences*, exactly as the worked number above exhibits ($\sum_\alpha \epsilon_\alpha = 0$ with each $\epsilon_\alpha \neq 0$).

Derivation, step 3: flavor-blind damping cancels from the flavor-odd source. The decohering action of a thermal rate on a flavor matrix is its double commutator with the projector that rate measures. For a flavor-universal rate the projector is the identity, and

$$[\mathbf{1}, [\mathbf{1},\, \delta n_\ell]] = 0 \;\text{ exactly, for every } \delta n_\ell; \qquad [P_\tau, [P_\tau,\, \delta n_\ell]] = \delta n_\ell \;\text{ on pure tau coherences.} \tag{621}$$

Gauge damping ($g^2 T \simeq 1.5 \times 10^3$ GeV here) therefore cannot regulate a source it does not distinguish: it cancels from the flavor-odd channel, however large it is. The leading flavor-*non*-universal rate in the plasma is the tau Yukawa rate of Eq. (617); the muon and electron rates are down by $(y_\mu/y_\tau)^2 \simeq 3 \times 10^{-3}$ and below. This is why the tau is the fastest honest measurer — and why, on the C4 reading (and pending exactly the resummation C4 names), the N's own width ($\Gamma_N \sim 10^{-10}$ GeV, seven orders below Γ_τ) would not be the operative regulator: on that reading the interference that builds $\epsilon_{i\alpha}$ is cut off the moment the bath learns *which flavor* was produced, long before the N lines themselves decay.

Derivation, step 4: two-level reduction and the effective Breit–Wigner. Reducing the heavy pair (i, j) to a two-level system, the CP-odd weight accumulated by an interference oscillating at ΔM while its flavor coherence is measured at $\Gamma_{\rm fl}$ is

$$I(\Gamma_{\rm fl}) = \Gamma_{\rm fl} \int_0^\infty \mathrm{d}s\, e^{-\Gamma_{\rm fl}\, s} \sin(\Delta M\, s) = \frac{\Gamma_{\rm fl}\, \Delta M}{\Gamma_{\rm fl}^2 + \Delta M^2} \;\le\; \frac{1}{2}, \tag{622}$$

a Lorentzian partial-resonance envelope: the Breit–Wigner factor of Eq. (613) with $\Gamma^{\rm eff}$ in place of the vacuum width, up to a convention factor two off-peak (a priced $O(1)$). The bound $|\mathcal{R}_{\rm res}| \le \frac{1}{2}$ holds for *every* regulator value — both envelopes obey $xy/(x^2+y^2) \le \frac{1}{2}$ — so the partial-resonance form survives the reduction; what the reduction assumes is that the damped coherence is the lepton-flavor one (condition C4 below).

Worked number. One-loop running gives $y_\tau(3.85\ \text{TeV}) = 0.0104$, $y_\tau^2 = 1.08 \times 10^{-4}$ (a +3.6% effect in y_τ^2 over the run from m_t — a refinement, not a lever). Then

$$\Gamma^{\rm eff} = c_\tau y_\tau^2 M_N = 2.1\text{–}4.1\ \text{MeV}, \qquad \frac{\Delta M}{\Gamma^{\rm eff}} = (2.7\text{–}5.3) \times 10^4, \qquad F_{\rm BW} = (0.9\text{–}1.9) \times 10^{-5}, \tag{623}$$

against the hypothesized 3 MeV, 3×10^4, and 1.5×10^{-5}: the installed target sits inside the band, reproduced exactly at $c_\tau = 8.9 \times 10^{-3}$. Consistency: $\Gamma_\tau/H(T{=}M_N) = (1\text{–}2) \times 10^8$ — the fully flavored regime of L-Cosmos-F holds with eight orders of margin — and $\Gamma_\tau/\Delta M \simeq 3 \times 10^{-5} \ll 1$: the mechanism is genuinely *partial*, not an accidental return to the excluded peak. Throughput: at $T = M_N$ the flavored chain gives $\eta_B = (1.1\text{–}2.3) \times 10^{-10}$ on the two sourcing states, i.e. 0.19–0.37 of the observed value, from below. The proper T-integration of the source over the decoupling window (the flavored Boltzmann system with $\Gamma^{\rm eff} \propto T$ through the window) is carried out in §I.6: the surviving asymmetry is produced at $z = M_N/T \simeq 1.5\text{–}4$, where the mechanism's own regulator has dropped with the temperature, and the T-integrated landing, scanning the declared freedoms, is 0.06–0.15 of observed (68%; central 0.09) — always an undershoot, never an overshoot (condition C6 below). The one-sidedness is itself informative: the missing transport prefactors must be enhancements, and the known members of that class (the $\Delta L = 1$ scattering source, the Lorentzian-versus-Breit–Wigner convention, the relativistic-window residue, and spectators with the A-matrix — GG-A9-SM's four priced omissions) net upward, supplying a factor $\simeq 3.3$ centrally and leaving the observation 2.2σ above in the log — a factor of ten short before the inventory, three after it, and one-sided.

The regulator landscape. The identification is falsifiable because the candidate prescriptions are nowhere near one another:

Table 56: The regulator landscape at the lock crossing ($\Delta M = 110$ GeV, $T = M_N$): the four candidate prescriptions for $\Gamma^{\rm eff}$, spanning sixteen orders of magnitude in $\Delta M/\Gamma^{\rm eff}$. Exactly one lands in the decade H-$\phi_{\rm geo}$-1 requires. Throughput column at the single-T evaluation.

Candidate	$\Gamma^{\rm eff}$ [GeV]	$\Delta M/\Gamma^{\rm eff}$	$F_{\rm BW}$	$\eta_B/\eta_B^{\rm obs}$
(a) N widths (S-matrix prescription)	$\sim 2 \times 10^{-10}$	6×10^{11}	8×10^{-13}	1×10^{-8}
(b) flavor measurement $c_\tau y_\tau^2 T$	2–4 $\times 10^{-3}$	(2.7–5.3)$\times 10^4$	(0.9–1.9) $\times$ 10^{-5}	0.19–0.37
(c) oscillation rate ΔM itself	1.1×10^2	1	0.40	9×10^3
(c′) gauge damping $g^2 T$	1.5×10^3	0.07	0.14	3×10^3
(d) Hubble rate $H(T{=}M_N)$	2.1×10^{-11}	5×10^{12}	1×10^{-13}	2×10^{-9}

The landscape discriminates cleanly rather than scattering: sixteen orders between its extremes, and one — only one — candidate in the required decade. The N-width class undershoots the observed asymmetry by $\sim 10^{7\text{–}8}$; the oscillation-rate and gauge-damping prescriptions re-land on the peak overshoot the flavored recomputation already excluded; the Hubble rate is worse than all of them.

Conditions. The derivation above is honest about what it assumes. Each assumption is a numbered condition with its status:

C1. *Flavor-odd source.* With R real, $\mathrm{Tr}\, S = 0$ exactly (Eq. (620), the kinetic form of Lemma I.1). Status: established (algebraic; machine-checked).

C2. *Flavor-blind cancellation.* Gauge damping enters through the identity projector and drops from the flavor-odd channel (Eq. (621)). Status: established within the density-matrix framework; its persistence under full resummation is folded into C4.

C3. *Two-level reduction.* The per-pair Lorentzian envelope and the $|\mathcal{R}_{\rm res}| \leq \frac{1}{2}$ bound survive for every regulator value (Eq. (622)); Lorentzian-vs-Breit–Wigner is a priced factor two. Status: established at reduction level, given C4.

C4. *Regulator identification (the crux).* In the regime $\Gamma_\tau/\Gamma_N \sim 10^7$ with the flavor-summed source zero, the interference time of the flavor-odd source is $1/\Gamma_\tau$, not $1/\Gamma_N$: the vacuum N-width regulator of the S-matrix prescription is superseded in-medium. The displayed kinetic system carries the heavy-pair coherence as an external source; discharging C4 requires promoting it to a dynamical variable with its own damping — that is the precise content the Kadanoff–Baym resummation must supply. Status: open at derivation level — structurally supported, but the full Kadanoff–Baym/CTP resummation has not been carried out, and the standard fully-flavored Boltzmann limit would return the (fatal) N-width class instead.

C5. *The coefficient.* $c_\tau \in [5 \times 10^{-3}, 10^{-2}]$ at $T \simeq M_N$ for lepton coherences sourced by N decays in an electroweak-symmetric bath (the kinetic-estimate, regime-boundary, and CTP

equilibration classes of the flavor-equilibration literature, mutually consistent within a factor two). Status: plausible; the precise value is open.

C6. *Window throughput.* The T-integrated window cost is part of the honest chain, and the priced $O(1)$ transport prefactors must supply the remaining one-sided enhancement, measured by the flavored Boltzmann integration of §I.6 at $E = 10.6$, 68% interval [6.6, 17.0] (retiring the prior bracket $\simeq$ 2–20). Status: measured at flavored-Boltzmann level; closure open.

Kill scenarios. Two clean kills and one soft one, all on record. (i) If the finite-temperature derivation returns the N-width class even in the flavored real-R corner, the chain undershoots by $\sim 10^{7-8}$ and is falsified from within. (ii) If it returns the flavor-blind thermal width g^2T, the excluded $\sim 3 \times 10^4$ peak overshoot stands — same verdict. (iii) If C4 and C5 close but the measured window deficit of C6 ($E = 10.6$, 68% [6.6, 17.0]) survives a full transport computation at its adverse end with the priced prefactors exhausted (they supply $\simeq 3.3$ centrally, leaving 2.2σ in the log), H-$\phi_{\rm geo}$-1 fails by throughput rather than by regulator identity. Either clean kill would discharge the hypothesis negatively; the mechanism has made it *more* falsifiable, not less.

GG connection. Why the ledger integer appears: the splitting $\Delta M = M_N/35$ is the ledger's own level spacing, so once the bath fixes the measurement rate, the resonance quality $\Delta M/\Gamma^{\rm eff}$ is geometry over chemistry — the same integer that sets the hierarchy e^{-K}, the lock mass $\frac{1}{2}M_{\rm Pl}\,e^{-K}$, and the throughput target $e^{-K/2}$ also sets how far off peak the resonance sits, with the tau Yukawa — an object entirely of the Standard Model — supplying the only other factor. The heavy mass cancels; no GG-sector dial, no neutrino-sector input, no freedom beyond the $O(1)$-band kinetic coefficient c_τ survives into Eq. (618). And the sign story is untouched: $\Gamma^{\rm eff}$ is CP-even, so the Majorana parity $\alpha_{31} = \pi$ still sets the sign of the matter excess, exactly as registered above.

Status. Candidate grade. Established: the closed form Eq. (618) as the consequence of the identification, the landscape discrimination of Table 56, and the flavor-odd source structure (C1–C3). Derived-conditional: the effective regulator itself, on C4–C5; the throughput landing, additionally on C6. Open: the full Kadanoff–Baym resummation, the precise c_τ at $T \sim M_N$, and the closure of the window factor. H-$\phi_{\rm geo}$-1 remains a named hypothesis — not retired, not closed — now sharpened from a numerical target into the conditions C1–C6.

Placement in the kinetic literature. The regulator question of C4 is not being asked in a vacuum. The flavor-covariant density-matrix and Kadanoff–Baym literature supplies the machinery in which it must be settled: the flavor-covariant transport equations of Dev *et al.* (2014), the first-principles resonant-regime analysis of Garny *et al.* (2013), the flavor-mixing transport theory of Jukkala *et al.* (2021), and the unified low-scale framework of Klarić *et al.* (2021) all return effective regulators of the heavy-sector class in the regimes they treat — which is exactly why C4 is carried open rather than assumed. On the other side, the 2025–26 thermal-resonant line — the soft-lepton resonant asymmetry of Kanemura and Li (2025) and the thermal-resonant mechanism of Li and Pilaftsis (2026a,b) — derives flavor-odd CP sources clocked by charged-lepton Yukawa rates,

with gauge damping dropping out of the CP-odd channel: the strongest published support for the structure of C1–C2 and for the lepton-sector clock, while leaving C4's cross-sector matching (the heavy ΔM oscillation regulated by Γ_τ) exactly as open as it is graded here. The resummation that discharges C4 should therefore also assess whether a thermal-resonant lepton-sector source of that type competes with, or replaces, the ΔM-filtered source at $T \sim M_N$.

I.4 Washout and the strong-washout efficiency

Statement. The surviving fraction of the generated asymmetry is the efficiency κ, controlled by the washout strength

$$K_{\rm wash} = \frac{\widetilde{m}}{m_\star}, \qquad \widetilde{m} = \frac{(Y_\nu Y_\nu^\dagger)_{ii}\, v^2}{M_i}, \qquad m_\star \simeq 1.1 \times 10^{-3}\,\text{eV}. \tag{624}$$

Derivation. The equilibrium scale follows from equating the decay rate to Hubble at $T = M_i$, $m_\star = (8\pi v^2/M_i)\,(H(M_i)/M_i)$, with $H(M_i) = 1.66\, g_*^{1/2} M_i^2/M_{\rm Pl}$; for $g_* \simeq 106.75$ this gives $m_\star \simeq 1.1 \times 10^{-3}\,\text{eV}$. The inherited light spectrum makes the effective mass that the decaying state couples to of order the heaviest light neutrino, $\widetilde{m} \sim m_3 \simeq 5 \times 10^{-2}\,\text{eV}$. In the strong-washout regime ($K_{\rm wash} \gg 1$) the final efficiency takes the standard suppressed form, well approximated by

$$\kappa \simeq \frac{1}{K_{\rm wash}(z_f)} \sim \frac{2}{K_{\rm wash}\, z_f}, \tag{625}$$

with $z_f = \mathcal{O}(\text{few})$ the freeze-out value of $z = M_i/T$. Strong washout erases the dependence on initial N abundance, so the result is initial-condition independent.

Worked number. With $\widetilde{m} \simeq 5 \times 10^{-2}\,\text{eV}$ and $m_\star \simeq 1.1 \times 10^{-3}\,\text{eV}$,

$$K_{\rm wash} \simeq 45\text{–}50, \qquad \kappa \simeq \frac{1}{K_{\rm wash}} \times \mathcal{O}(1) \simeq 1 \times 10^{-2}. \tag{626}$$

The convention note of Eq. (611) applies here in reverse: $\widetilde{m} = m_3$ is the washout convention of the chain; at the seam-matched $\widetilde{m} \simeq 8.4$ meV one would instead have $K_{\rm wash} \simeq 7.6$ and $\kappa \simeq 0.13$. The factor-$\simeq 12$ spread in Y_ν^2 between the two conventions is declared, accounted for at the convention note of Eq. (611), and priced as an $O(10)$ freedom of the order-of-magnitude register, carried with the transport $\mathcal{O}(1)$'s.

GG connection. The washout scale is set entirely by inherited quantities — the seesaw-matched Yukawa and the GG-A9-SM light masses — so κ is not a free knob.

Status. Derived-conditional: the regime (strong washout) and the order of magnitude are fixed by inherited inputs; the $\mathcal{O}(1)$ prefactor in Eq. (626) is measured, flavor by flavor, by the Boltzmann integration of §I.6.

I.5 Assembling η_B

Statement. The boundary CP-odd throughput is the product $\epsilon^{\rm res}\kappa$, converted to η_B by the standard sphaleron and entropy factors.

Derivation and worked number. The product reconstructs the geometric ledger value,

$$\epsilon^{\rm res}\,\kappa \simeq (2.5\times10^{-6})(1.0\times10^{-2}) = 2.51\times10^{-8} = e^{-17.5} = e^{-K/2}. \tag{627}$$

Then (the dedicated companion (Arisaka, 2026, C4-Leptogenesis) assembles the same chain), with the heavy-sector abundance carried at its equilibrium yield — $\sum_i Y_{N_i}^{\rm eq} \simeq 0.833/g_*$ for the two sourcing species of the quasi-degenerate triplet ($0.42/g_*$ per Majorana species, $g_N = 2$; the massless first state contributes no asymmetry; Eq. (189)) —

$$Y_L \simeq \Big(\sum_i Y_{N_i}^{\rm eq}\Big)\epsilon^{\rm res}\kappa \simeq \frac{0.833\,\epsilon^{\rm res}\kappa}{g_*}, \qquad \eta_B \simeq 7.04\,c_{\rm sph}\,Y_L = 7.04\left(\frac{28}{79}\right)\frac{0.833\,e^{-K/2}}{g_*}, \tag{628}$$

and inserting $c_{\rm sph} = 0.3544$, $g_* = 106.75$,

$$\eta_B \simeq 7.04\times0.354430\times0.833\times\frac{2.511\times10^{-8}}{106.75} \simeq 4.89\times10^{-10}. \tag{629}$$

GG connection. This is the explicit transport mechanism behind the compressed body result Eq. (193): the geometric $e^{-K/2}$ is the coarse-grained product of a phase-controlled resonant asymmetry and a strong-washout efficiency. The two declared joints are carried in the open: the partial-resonance condition (named input) and the $K/2$ exponent (postulate); the identity at displayed precision uses the $\mathcal{O}(1)$ freedom in κ and the target status of $\epsilon^{\rm res}$.

Status. Order-of-magnitude with a controlled mechanism. The central value sits -20% from the observed $\eta_B^{\rm obs} = (6.12\pm0.04)\times10^{-10}$; the honest band is factor-$\sim$3, dominated by the resonance condition and the $\mathcal{O}(1)$ washout prefactor; the time-resolved integration of the next subsection lands a factor of about ten below the ledger road's 6.18×10^{-10} and supersedes this closed form wherever the two differ.

I.6 The flavored Boltzmann integration: the window deficit measured

Statement. The closed-form worked number of the flavor-measurement subsection brackets the cost of the decoupling window at a single temperature; this subsection measures it. The flavored Boltzmann system — the kinetic form of the chain, every input either derived upstream or standard — is integrated over the full window, scanning only the declared freedoms ($c_\tau \in [5\times10^{-3}, 10^{-2}]$, the washout convention, the integration start z_0). Scanning only those freedoms the chain lands at

$$\eta_B^{\rm Boltzmann} \;=\; 0.58\,^{+0.35}_{-0.22}\times10^{-10} \quad (68\%\ \text{interval; band a factor 2.6}), \tag{630}$$

and the required one-sided enhancement — the quantity condition C6 names — is now a measured number with an error bar rather than a decade-wide bracket:

$$E \;\equiv\; \frac{\eta_B^{\rm obs}}{\eta_B^{\rm Boltzmann}} \;=\; 10.6\,, \qquad 68\%\ \text{interval}\ [6.6,\ 17.0]. \tag{631}$$

With the priced standard-omissions inventory applied (the lepton-number-violating scattering source, the Lorentzian-versus-Breit–Wigner convention, the relativistic-window residue, and spectators with the A-matrix — GG-A9-SM's four priced omissions), $\eta_B = 1.9\,^{+1.4}_{-0.8} \times 10^{-10}$: the observation sits 2.2σ above it in the log, just above the upper edge of the 95% band. (These values are imported from GG-A9-SM, not recomputed here. An earlier version of this subsection carried the corresponding three-state numbers — 0.86, $E = 7.1$, 2.8 — each exactly 3/2 of the value above, because the integration was inherited from a treatment that sourced from three heavy states. The assembly above this subsection sources from two; the integration now does too, and the honest register is a factor of ten short and one-sided, not a factor of seven.) Everything in this subsection is conditional on C4 and C5, exactly as graded; nothing here upgrades any grade.

The equations, term by term. In the variable $z = M_N/T$ (time runs with z) and photon-normalized abundances $N_X = n_X/n_\gamma$, the system is

$$\frac{{\rm d}N_N}{{\rm d}z} = -\,D(z)\,(N_N - N_N^{\rm eq}), \tag{632}$$

$$\frac{{\rm d}N_{\Delta\alpha}}{{\rm d}z} = \epsilon_\alpha(z)\,D(z)\,(N_N - N_N^{\rm eq}) \;-\; \tfrac{1}{4}\,K_\alpha\,z^3 K_1(z)\,N_{\Delta\alpha}, \qquad \alpha = e,\mu,\tau. \tag{633}$$

Every symbol earns its place. $N_N^{\rm eq}(z) = \frac{3}{4}\cdot\frac{1}{2}\,z^2K_2(z)$ per species is the equilibrium abundance: three quarters of the photon abundance while relativistic, Boltzmann-suppressed as e^{-z} once T falls below the mass, the Bessel function K_2 the correct relativistic bookkeeping between the limits. $D(z) = K_{\rm tot}\,z\,K_1(z)/K_2(z)$ is the decay term; the ratio $K_1/K_2 \to z/2$ at small z is time dilation — a hot, fast heavy neutrino lives longer in the plasma frame, so decays switch on only as the bath cools toward $z \sim 1$. The source term says that every net decay (the excess of decays over inverse decays, which is exactly the rate at which N_N follows its falling equilibrium) deposits a CP asymmetry $\epsilon_\alpha(z)$ into flavor α; the asymmetry per decay is z-dependent because the regulator is thermal,

$$\epsilon_\alpha(z) = \bar\epsilon_\alpha \cdot 2\mathcal{R}(z), \qquad \mathcal{R}(z) = \frac{2\,\Delta M\,\Gamma^{\rm eff}(z)}{4\Delta M^2 + \Gamma^{\rm eff}(z)^2}, \qquad \Gamma^{\rm eff}(z) = c_\tau y_\tau^2\,\frac{M_N}{z}, \tag{634}$$

with $\bar\epsilon_\alpha = (+3.4,\,+3.1,\,-6.6)\times 10^{-3}$ the flavored pattern computed above on the ledger inputs (peak-normalized; $\sum_\alpha \bar\epsilon_\alpha = 0$, the exact flavor-sum zero of Lemma I.1). Since $\Gamma^{\rm eff} \propto T$, the resonance factor $\mathcal{R}(z)$ *falls* as the universe cools — the interference is measured more slowly while the splitting stays put — and that decline is the physical origin of the window suppression measured below. The washout term erases each flavor at its own rate: $\frac{1}{4}K_\alpha z^3 K_1(z)$ is the standard inverse-decay kernel, peaking near $z \sim 2$–4, with $K_\alpha = (3.4,\,28.0,\,22.4)$ the flavor-projected washouts of the

chain. The kernels are decays and inverse decays; the $\Delta L = 1$ scattering source, the Lorentzian-versus-Breit–Wigner convention, the relativistic-window residue, and spectators with the A-matrix — GG-A9-SM's four priced omissions — are deliberately omitted and priced, each a standard, named effect, and each of the omitted source-side effects points *up*. The equations are integrated from $z_0 = 1$ (scanned 0.5–2) to $z = 50$ with an exponentially exact stiff stepper.

Validation. A numerical claim in this cohort must arrive with its checks — three, none adjustable after the fact. For the efficiency check the code is run with a constant unit asymmetry in a single flavor and compared with the standard analytic strong-washout fit

$$\kappa(K) \simeq \frac{2}{K\, z_B(K)}\Big(1 - e^{-K z_B(K)/2}\Big), \qquad z_B(K) \simeq 2 + 4K^{0.13} e^{-2.5/K}, \tag{635}$$

whose physical content is one sentence: the asymmetry that survives is the asymmetry made in the *last* moment before the inverse decays shut off, at $z = z_B$, and by then only a fraction $\sim 1/(K z_B)$ of the heavy states remains to make it — κ is the fraction made late.

Table 57: Validation of the integration. The two K_α computations were done independently on the same derived inputs; the fit comparison spans the full washout range of the problem. The columns give the check, the integration, the reference value, and the agreement. The two computations were run independently on the same derived inputs, so the agreement tests the integration and not the inputs, and a reader looking for validation of the inputs themselves should look at the audit tables instead.

Check	Integration	Reference	Agreement
K_α from derived U, masses	(3.5, 29.3, 22.5)	chain values above: (3.4, 28.0, 22.4)	2–5 %
$\kappa(K{=}3.5)$	0.1371	fit Eq. (635): 0.1327	3 %
$\kappa(K{=}22.5)$	0.0116	0.0121	4 %
$\kappa(K{=}29.3)$	0.0084	0.0089	5 %
single-T $\epsilon\kappa$ at $c_\tau = 8.9 \times 10^{-3}$	2.7×10^{-8}	ledger $e^{-K/2} = 2.5 \times 10^{-8}$	bookkeeping $\mathcal{O}(1)$

Worked numbers. At central parameters the integration lands at $\eta_B = 0.69 \times 10^{-10}$ on the two sourcing states, i.e. 0.11 of observed (GG-A9-SM's own central value; an earlier version of this paragraph printed the three-state 1.0×10^{-10}, 0.17); scanning the declared freedoms gives Eq. (630) and the measured deficit Eq. (631). Because the integration resolves time, it also localizes what the single-temperature estimate could not: the asymmetry is generated at $z \simeq 1.5$–4, where the thermal regulator $\Gamma^{\rm eff} \propto 1/z$ has already fallen from its $z = 1$ value — the resonance is sharpest exactly when the washout is also sharpest, and the survivors are made late, off the best part of the resonance. Every one of the omitted source-side effects (the $\Delta L = 1$ scattering source and the relativistic-window residue, the first and third of the four) moves net production *earlier*, where $\mathcal{R}(z)$ is larger: the one-sidedness of the deficit is visible in the mechanism, not asserted.

The electron door, and the sign as an output. The flavor-summed source is exactly zero (Lemma I.1), so the integration starts with nothing to keep: the three flavor components are made with opposite signs and sum nearly to zero at every stage. What survives is the *difference* in erasure. The μ and τ components sit in strong washout ($K_\mu \simeq 28$, $K_\tau \simeq 22$) and are ground down; the electron flavor is decoupled from the heavy third mass state by the derived reactor angle — $|U_{e3}|^2 = \sin^2\theta_{13} = 3/140$, the smallest element of the mixing matrix and a ledger quantity — so $K_e \simeq 3.4$ and the electron component survives an order of magnitude better: *the electron ledger ends holding* 106% *of the net asymmetry, the muon's* +3% *and the tau's* −10% (GG-A9-SM's integration). The sign of the excess is then an *output* of differential erasure, not an input: nothing in the integration was told which sign to prefer (the source pattern sums to zero), and the net asymmetry emerges positive — matter — because the negative, τ-dominated components die faster, exactly as the Majorana parity $\alpha_{31} = \pi$ requires on the sign audit above. On this chain, the smallness of the reactor angle is why there is something rather than nothing.

GG connection. Every input is inherited (masses, mixings, phases, $\Delta M = M_N/K$, the seesaw-matched Yukawas) or standard (kernels, g_*); the integration has no dial. What it measures is the cost of the geometry's own timing: the ledger fixes the splitting, the bath fixes the falling regulator, and the survivors are necessarily made late, off the best part of the resonance — the window deficit E is the price of that timing, now measured.

Status. Measured at flavored-Boltzmann level, conditional on C4 and C5 (candidate grade for the regulator; the kill scenarios of the flavor-measurement subsection apply unchanged). The measured $E = 10.6$ (68% [6.6, 17.0]) supersedes the single-temperature bracket wherever the two differ. The priced standard-omissions inventory supplies a factor $\simeq 3.3$ centrally, leaving the observation 2.2σ above in the log — a factor of ten short before the inventory, three after it, one-sided, and not closed.

I.7 The open-system A=7 port and the $1/\pi$ survival fraction

Statement. The mass-7 ($^7\mathrm{Be} \to {}^7\mathrm{Li}$) yield is suppressed by the closed-packet branching factor $1/\pi$ — the same factor that fixes $\Omega_m = 1/\pi$ — so that $(^7\mathrm{Li}/H)_{\mathrm{GG}} = (^7\mathrm{Li}/H)_{\mathrm{SBBN}}/\pi$.

Derivation. The relevant A=7 channels are

$$^3\mathrm{He} + \alpha \xrightarrow{\lambda_{3\alpha 7}} {}^7\mathrm{Be} + \gamma, \qquad {}^7\mathrm{Be} + e^- \to {}^7\mathrm{Li} + \nu_e, \qquad {}^7\mathrm{Li} + p \xrightarrow{\lambda_{p\alpha}} 2\alpha. \tag{636}$$

Because the proton-destruction channel of ^{7}Li is fast ($\lambda_{p\alpha} \gg H$) during the active window, any ^{7}Li reaching that window early is efficiently burned; the surviving A=7 inventory is therefore controlled by how much ^{7}Be survives its open electron-capture / phase-noise window intact. Model that window as a System Instance Packet whose conserved object is the mass-7 flux Φ_7 carried over the compact angular phase $\theta \in S^1_\theta$. The open-system (GDOS) drift of the packet over the compact angle

is

$$\frac{\mathrm{d}\Phi_7}{\mathrm{d}\theta} = -\,\mathcal{P}(\theta)\,\Phi_7, \qquad S_{A=7} = \frac{\Phi_7^{\mathrm{out}}}{\Phi_7^{\mathrm{in}}} = \exp\left[-\oint_{S^1_\theta} \mathcal{P}(\theta)\,\mathrm{d}\theta\right], \tag{637}$$

where $\mathcal{P}(\theta)$ is the port density. The closed-packet construction (Conjecture VIII) shows that, when a conserved flux is branched over the compact angle by the same GFF closed-packet geometry that fixes the cosmic inventory, the surviving share onto the observable boundary is the solid-angle / phase fraction

$$S_{\mathrm{closed-packet}} = \frac{1}{\pi}, \tag{638}$$

the identical factor that yields $\Omega_m = 1/\pi$: the projection of the closed packet over the compact angular measure retains the fraction $1/\pi$ of the incoming flux. Identifying the A=7 port with this same branching gives $S_{A=7} = 1/\pi$, i.e. the line integral in Eq. (637) equals $\ln \pi$.

Worked number.

$$\left(\frac{^7\mathrm{Li}}{H}\right)_{\mathrm{GG}} = \frac{1}{\pi}\left(\frac{^7\mathrm{Li}}{H}\right)_{\mathrm{SBBN}} \simeq \frac{5.0\times10^{-10}}{3.14159} \simeq 1.59\times10^{-10}, \tag{639}$$

to be compared with the Spite-plateau value $(1.6 \pm 0.3)\times10^{-10}$. The deuterium and helium sectors are untouched: even if every removed ^{7}Be fed a side channel, the fractional shift in D/H is $\lesssim 3\times10^{-5}$ (Eq. (217)).

GG connection. The lithium suppression is *predicted*, not fitted: it is the same closed-packet $1/\pi$ that fixes the cosmic matter inventory, applied to the mass-7 flux. This replaces the earlier fitted port strength entirely.

Status. D* (Derived-conditional), sharing the closed-packet-integration conjecture's status with the inventory $\Omega_m = 1/\pi$. Falsifiable: any firm primordial ^{7}Li abundance away from SBBN/π excludes the port.

The launch selects the mechanism. One structural loose end closes once the inflationary sector is derived. Leptogenesis requires the heavy-neutrino sector to be populated: from a thermal bath if the reheating temperature is far above the right-handed masses, non-thermally otherwise. With the launch underived, both readings would remain open. The derived reheating temperature of Appendix G, $T_{\mathrm{rh}} \simeq 6\times10^{10}$ GeV, sits seven orders above the heavy states, which then equilibrate through their own inverse decays near 10–20 TeV — and the prompt lock crossing ($N_2 = 0$) guarantees zero later dilution of the asymmetry. The mechanism is therefore no longer a model choice: the same computation that fixed n_s and r guarantees the bath for resonant leptogenesis on the inherited CP ledger. That is the cross-sector economy of the descent in one sentence: one derived temperature, three consequences (the e-fold count, the asymmetry mechanism, the absence of dilution).

Problems and worked solutions. I.1.1. Estimate the seesaw scale from the observed neutrino masses.

Solution. Step 1: the seesaw relation (608) gives $m_\nu \sim y^2 v^2/M_N$ with $v = 174$ GeV. Step 2: for the atmospheric scale $m_\nu \approx 0.05$ eV and $y \sim 1$: $M_N \sim v^2/m_\nu = (174\ \text{GeV})^2/0.05\ \text{eV} \approx 6 \times 10^{14}$ GeV. Step 3: the scale lands within an order of magnitude of the inflationary energy of G.7.1 and just below the classic unification scale — three independently derived energies crowding the same decade, which is the numerical face of the cohort's claim that one geometric depth underlies them. Physics intuition: the seesaw converts the smallness of the neutrino mass into the largeness of a new scale — the lighter the neutrino, the higher the fulcrum; nature's lightest fermions are the longest levers we possess into the deepest physics.

I.1.2. Show that the resonant CP asymmetry is maximized at $\Delta M = \Gamma_N/2$, with maximum $\frac{1}{2}\sin(2\phi_{\text{geo}})$.

Solution. Step 1: the Breit–Wigner form (613) gives $\epsilon \propto f(\Delta M) = \dfrac{\Delta M\,\Gamma_N/2}{(\Delta M)^2 + \Gamma_N^2/4}$ times the CP phase factor $\sin(2\phi_{\text{geo}})$ of Eq. (614). Step 2: maximize: $f'(\Delta M) = 0$ at $(\Delta M)^2 = \Gamma_N^2/4$, i.e. $\Delta M = \Gamma_N/2$; there $f = \frac{\Gamma_N^2/4}{\Gamma_N^2/2} = \frac{1}{2}$. Step 3: hence $\epsilon^{\text{res}} \le \frac{1}{2}\sin(2\phi_{\text{geo}})$ (614) — an $O(1)$ asymmetry from a mass splitting tuned to the decay width, which is what lets a TeV-adjacent or GUT-adjacent pair generate the observed asymmetry without enormous CP phases. Physics intuition: resonance is interference bookkeeping — the two nearly degenerate states beat against each other for exactly one lifetime when their splitting matches their width; narrower and they blur into one state, wider and they never mix. The ledger's role is to fix the splitting.

I.4.1. Audit the η_B assembly arithmetic end to end.

Solution. Step 1: the pieces of (628): the entropy-to-photon conversion 7.04 (F.8.1's factor), the sphaleron share $c_{\text{sph}} = 28/79 = 0.3544$ (the equilibrium fraction of $B - L$ stored as B), the efficiency-weighted asymmetry $\epsilon^{\text{res}}\kappa = e^{-K/2} = 2.51 \times 10^{-8}$ (the ledger's $K = 35$ entering as its half-depth), the sourcing equilibrium yield $\sum_i Y_{N_i}^{\text{eq}} = 0.833/g_*$ (two Majorana species' fair share of the bath — the third, massless, sources nothing — Eq. (189)), and the bath's $g_* = 106.75$. Step 2: multiply: $\eta_B = 7.04 \times 0.354430 \times 0.833 \times 2.511 \times 10^{-8}/106.75 = 4.89 \times 10^{-10}$ (629). Step 3: compare the target the mechanism must reach, the ledger road's 6.18×10^{-10} (and the observed 6.12): the closed form lands 20% below it with no continuous dial, and the time-resolved integration of §I.6 a factor of about ten below — the asymmetry of all matter over all antimatter traced to one integer's exponential, and honestly short. Physics intuition: every factor in the product has a face — a conservation law, an equilibrium share, a depth, a head count — and the audit's point is that nothing else is in the product; η_B is bookkeeping all the way down.

I.5.1. Derive the lithium survival factor $1/\pi$ from the port branching.

Solution. Step 1: the $A = 7$ population drains through the port with drift given by the line integral of (637); the closed-packet branching assigns the through-port share $1/\pi$ of the incoming flux — the same branching that sets $\Omega_m = 1/\pi$ upstream. Step 2: identifying the drift integral with $\ln\pi$, the survival factor is $S_{A=7} = e^{-\ln\pi} = 1/\pi$ (638). Step 3: applied to the SBBN yield: ^{7}Li/H $= 5.0 \times 10^{-10}/\pi = 1.59 \times 10^{-10}$ (639), against the Spite plateau's $\approx 1.6 \times 10^{-10}$ — H.5.1's mystery factor, derived. Physics intuition: the same geometric branching ratio surfaces twice at wildly

different epochs — once setting the matter fraction of the whole universe, once pruning one isotope in the first twenty minutes — and it is exactly this reuse, with no per-instance freedom, that makes the reading falsifiable rather than fitted.

Take-away. *Standard result:* the Sakharov conditions and the seesaw give a viable asymmetry, but the CP phases, masses, and the mechanism's epoch are all model choices. *What it assumed:* a heavy-neutrino spectrum and its phases, chosen to work. *What the geometry replaced it with:* the phases and masses inherited from the GG-A9-SM ledger, and the bath guaranteed by the derived $T_{\rm rh}$ — resonant leptogenesis on a full bath with zero dilution, no fork. *Continue in:* Part V; Appendix G for $T_{\rm rh}$.

J Linear perturbations and the growth of structure

What this appendix teaches. The complete linear theory of structure: the perturbed metric and fluid, the growth equation and its solutions, the Meszaros suppression, the transfer function, $P(k)$, and σ_8, all derived with worked numbers; the observational S_8 tension stated carefully; and the geometric reading, which on the one background is short: the growth history is the standard one, σ_8 and S_8 are computed from the declared amplitude, and the tension is shared, not resolved — the earlier corridor reading kept as a record.

Read alongside: Dodelson (2003) chs. 5–7 (this appendix follows its line-by-line style); Peebles (1993) for the classic treatment; Peacock (1999) at depth.

The gauge-invariant linear theory of cosmological perturbations and the growth of structure assembled in this appendix are summarized in Figure 41; the material in this appendix rests on the classic and review literature (Bardeen, 1980; Bardeen *et al.*, 1986; Kodama & Sasaki, 1984; Lifshitz, 1946; Mukhanov *et al.*, 1992; Press & Schechter, 1974).

This appendix builds linear cosmological perturbation theory and the growth of structure from scratch, at the depth a first-year graduate student can follow line by line. It is the natural sequel to the FRW background appendix: there we asked how the homogeneous, isotropic universe expands; here we ask how the small departures from homogeneity — the seeds laid down in the early universe — grow under gravity into galaxies, clusters, and the cosmic web, and how their statistics are compressed into the two numbers, σ_8 and S_8, that carry the present-day clustering tension. The organizing principle is the one used throughout these technical appendices: state the object precisely, carry out the full computation without skipping the index work, anchor it in a worked numerical example with units, and only then connect the result back to the GG-Theory program.

One point of intellectual honesty governs the whole appendix and is stated again, sharply, in the closing Status note. The *standard* linear growth theory developed in Sections J.1–J.8 is textbook general relativity plus a perfect-fluid source; it is taught here in full and is unmodified by the framework, and on the one background the paper now carries it is the whole of the framework's growth history: the GG-Theory reading supplies the background and the amplitude (§43.3) and modifies nothing else. The earlier reading, in which a two-stage corridor modified the late-time growth with a sign the paper left open, is kept in Section J.10 as a record; its parametrization assumed a *suppression* of late growth (which would have helped the tension), but in full early-dark-energy likelihood fits the net S_8 shift typically comes out *positive* — the opposite sign, which would

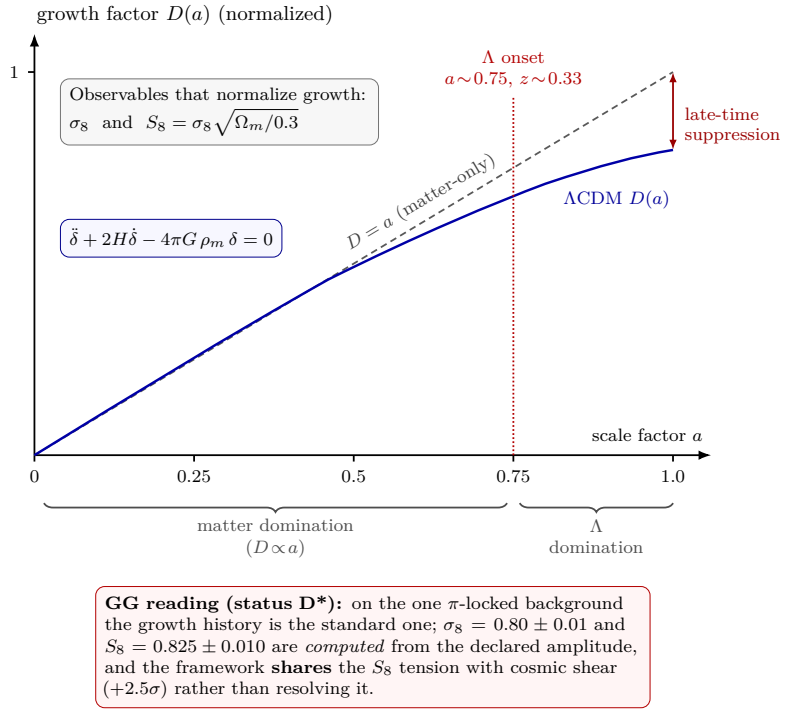

Figure 41: Linear growth factor $D(a)$ versus scale factor a. In the matter-dominated era the normalized linear growth factor tracks $D \propto a$ (it would follow the dashed $D = a$ reference exactly for a matter-only universe). As dark energy comes to dominate past $a \sim 0.75$ ($z \sim 0.33$) growth is *suppressed* and $D(a)$ flattens, opening the late-time gap below the dashed line. The growth obeys $\ddot{\delta} + 2H\dot{\delta} - 4\pi G\rho_m\delta = 0$. The amplitude is fixed observationally by σ_8 and $S_8 = \sigma_8\sqrt{\Omega_m/0.3}$. **GG reading:** standard growth on the one π-locked background; σ_8 and S_8 computed from the declared amplitude (status D* on Lemma V′), the S_8 tension shared, not resolved. The retired corridor reading of this figure is kept in §J.10 as a record.

worsen the tension (Hill et al., 2020). The reader who wants only the master growth equation may read the boxed displays (654), (656) and (667) and skip ahead; the reader who wants to verify every step will find no gap papered over and no tension declared resolved that is not.

The three steps of this appendix. As in Appendix E:

Step 1. The standard linear theory (§J.1 onward). The perturbed metric and fluid, the sub-horizon growth equation, the growth factor $D(a)$, the Meszaros suppression, the transfer function, and the compression to σ_8 and S_8.

Step 2. What it leaves open (§J.9). The background $H(a)$ and the densities are inputs, and the S_8 clustering tension is a live disagreement that standard ΛCDM does not resolve.

Step 3. The retired geometric derivation (§J.10, record). The two-stage corridor that once reshaped $H(a)$ and recast the S_8 tension as two projections of one structure, with its sign left open — kept verbatim as the record of a withdrawn reading; nothing in it is claimed.

J.1 The perturbed FRW metric in conformal-Newtonian gauge

Background and conventions. We work in conformal time η, related to cosmic time t by $dt = a\,d\eta$, so that the spatially flat FRW background line element is

$$d\bar{s}^2 = a^2(\eta)\,[\,-d\eta^2 + \delta_{ij}\,dx^i dx^j\,], \tag{640}$$

with a the scale factor, normalized $a_0 \equiv a(\eta_0) = 1$ today. A prime denotes a derivative with respect to conformal time, $X' \equiv dX/d\eta$, and an overdot a derivative with respect to cosmic time, $\dot{X} \equiv dX/dt = X'/a$. The conformal Hubble rate is

$$\mathcal{H} \equiv \frac{a'}{a} = a\,H, \qquad H \equiv \frac{\dot{a}}{a}, \tag{641}$$

so $\mathcal{H}$ has units of inverse conformal time and H has units of inverse cosmic time; the comoving Hubble radius is $1/\mathcal{H}$. We use the metric signature $(-,+,+,+)$, Greek indices $\mu,\nu \in \{0,1,2,3\}$, Latin indices $i,j \in \{1,2,3\}$, and set $c = 1$ throughout (restoring c in the worked example).

The scalar perturbations and the gauge problem. At linear order the metric is split into a background plus a small perturbation, $g_{\mu\nu} = \bar{g}_{\mu\nu} + \delta g_{\mu\nu}$, with $|\delta g_{\mu\nu}| \ll |\bar{g}_{\mu\nu}|$. The perturbation $\delta g_{\mu\nu}$ decomposes, by its transformation under spatial rotations, into scalar, vector, and tensor parts (the scalar–vector–tensor, or SVT, decomposition). To linear order these three sectors *decouple* and evolve independently. Vectors decay in an expanding universe and tensors are the gravitational waves treated in the inflationary appendix; the growth of *matter* structure is carried entirely by the *scalar* sector, on which we now focus.

The scalar sector has, in general, four functions (ϕ, B, ψ, E). But two of them are pure gauge: an infinitesimal coordinate change $x^\mu \to x^\mu + \xi^\mu$ shuffles physical and coordinate content, and a naive choice can manufacture a fictitious "growing mode" that is nothing but a wiggle in the time slicing. The clean way to proceed is to fix the gauge so that the two remaining functions are physical. The

choice that makes the Newtonian limit transparent is the *conformal-Newtonian* (or longitudinal) gauge, $B = E = 0$, in which the perturbed line element reads

$$\boxed{ds^2 = a^2(\eta)\left[-(1+2\Psi)\,d\eta^2 + (1-2\Phi)\,\delta_{ij}\,dx^i dx^j\right].} \tag{642}$$

Here $\Psi(\eta,\mathbf{x})$ and $\Phi(\eta,\mathbf{x})$ are the two scalar potentials. Ψ multiplies the time–time component and plays the role of the Newtonian gravitational potential felt by a slow test particle (it is what sets gravitational redshift and the acceleration $\ddot{\mathbf{x}} = -\nabla\Psi/a^2$). Φ is the perturbation to the spatial curvature; it is the potential that governs the deflection of light, since photons respond to the sum $\Phi + \Psi$. The sign conventions are chosen so that an overdensity sits in a *potential well* with $\Psi < 0$, $\Phi < 0$.

Why $\Phi = \Psi$ in the regime we need. The off-diagonal ($i \neq j$, i,j spatial) Einstein equation, the so-called anisotropic-stress constraint, reads at linear order

$$k^2\,(\Phi - \Psi) = -8\pi G\,a^2\,(\bar{\rho} + \bar{p})\,\sigma, \tag{643}$$

where σ is the scalar anisotropic stress of the source and k is the comoving wavenumber of the Fourier mode. For a perfect fluid (pressureless matter, or radiation treated in the tight-coupling limit) the anisotropic stress vanishes, $\sigma = 0$, and therefore

$$\Phi = \Psi \qquad \text{(no anisotropic stress)}. \tag{644}$$

The leading correction comes from free-streaming neutrinos, which carry $\sigma \neq 0$ and split Φ and Ψ at the few-percent level after neutrino decoupling; we will note where that correction enters and otherwise set $\Phi = \Psi$, writing a single Newtonian potential where convenient. From here on we Fourier-transform every field, $f(\eta,\mathbf{x}) = \int d^3k\,(2\pi)^{-3}\,f_{\mathbf{k}}(\eta)\,e^{i\mathbf{k}\cdot\mathbf{x}}$, so that $\nabla^2 \to -k^2$ and each comoving mode $\mathbf{k}$ evolves independently at linear order. The label $\mathbf{k}$ is suppressed unless needed.

J.2 The perturbed fluid: continuity, Euler, Poisson

We now write the conservation laws for the matter fluid in the perturbed metric (642). The matter is modeled as a fluid with energy density $\rho(\eta,\mathbf{x}) = \bar{\rho}(\eta)\,[1 + \delta(\eta,\mathbf{x})]$, where

$$\delta \equiv \frac{\delta\rho}{\bar{\rho}} = \frac{\rho - \bar{\rho}}{\bar{\rho}} \tag{645}$$

is the dimensionless *density contrast* — the central object of the whole subject — and a peculiar velocity field $\mathbf{v} = \mathbf{u}/a$ (the deviation of the fluid four-velocity from the Hubble flow). Pressure is written $p = \bar{p} + \delta p$, and for cold dark matter (CDM) we will set $\bar{p} = \delta p = 0$; for radiation $\bar{p} = \bar{\rho}/3$.

Source of the equations: $\nabla_\mu T^{\mu\nu} = 0$. The two fluid equations are the time and space components of the covariant conservation of the stress–energy tensor $T^{\mu\nu} = (\rho + p)\,u^\mu u^\nu + p\,g^{\mu\nu}$, evaluated to first order in the perturbed metric. The covariant divergence $\nabla_\mu T^{\mu\nu} = \partial_\mu T^{\mu\nu} + \Gamma^\mu_{\mu\lambda} T^{\lambda\nu} + \Gamma^\nu_{\mu\lambda} T^{\mu\lambda} = 0$

requires the perturbed Christoffel symbols of (642); the relevant ones are

$$\Gamma^0_{00} = \mathcal{H} + \Psi', \qquad \Gamma^0_{0i} = \partial_i \Psi,$$
$$\Gamma^0_{ij} = \big[\mathcal{H} - 2\mathcal{H}(\Phi + \Psi) - \Phi'\big]\delta_{ij}, \qquad \Gamma^i_{0j} = \big[\mathcal{H} - \Phi'\big]\delta^i_j, \tag{646}$$

with the others following by symmetry. The reader who carries these through $\nabla_\mu T^{\mu 0} = 0$ and $\nabla_\mu T^{\mu i} = 0$ obtains the perturbed continuity and Euler equations; we quote the result and then re-derive it physically.

The perturbed continuity equation. The $\nu = 0$ component gives, for a general fluid,

$$\delta' + (1+w)\big(\theta - 3\Phi'\big) + 3\mathcal{H}\Big(\frac{\delta p}{\delta \rho} - w\Big)\delta = 0, \tag{647}$$

where $w \equiv \bar{p}/\bar{\rho}$ is the equation-of-state parameter and $\theta \equiv \nabla \cdot \mathbf{v} = i\mathbf{k}\cdot\mathbf{v}_\mathbf{k}$ is the (Fourier) velocity divergence. For pressureless matter, $w = 0$ and $\delta p = 0$, this collapses to the clean statement

$$\boxed{\delta' = -\theta + 3\Phi'.} \tag{648}$$

The physics is mass conservation: the density contrast grows when matter flows into a region ($\theta < 0$, convergent flow) and is modulated by the time-changing volume element through Φ'. In the sub-horizon, quasi-static regime that dominates structure formation the metric term $3\Phi'$ is negligible against θ, so $\delta' \simeq -\theta$: density contrast is sourced by velocity convergence, exactly as in laboratory fluid dynamics.

The perturbed Euler equation. The $\nu = i$ component gives, for pressureless matter,

$$\boxed{\theta' + \mathcal{H}\,\theta = k^2\,\Psi.} \tag{649}$$

This is Newton's second law for the fluid. The term $\mathcal{H}\theta$ is the Hubble drag: cosmic expansion redshifts peculiar velocities, damping them as $v \propto 1/a$ in the absence of forcing. The right-hand side $k^2\Psi$ is the gravitational force, the Fourier transform of $-\nabla\Psi$ (the sign works out to acceleration toward potential wells). For a fluid with pressure one adds a term $-k^2\,\delta p/(\bar{\rho} + \bar{p})$ on the right, the pressure gradient that resists collapse and sets the Jeans scale; for cold matter this term is absent, which is precisely why cold dark matter clusters on all scales while baryons (with pressure support before decoupling) do not.

The Poisson equation. The system is closed by relating Ψ to δ through the perturbed 00 Einstein equation. To linear order this is the relativistic generalization of the Poisson equation,

$$k^2\,\Phi + 3\mathcal{H}\big(\Phi' + \mathcal{H}\Psi\big) = -4\pi G\,a^2\,\bar{\rho}\,\delta. \tag{650}$$

On sub-horizon scales, $k \gg \mathcal{H}$, the first term dominates the bracket on the left (the bracket is order $\mathcal{H}^2\Phi$, suppressed by $(\mathcal{H}/k)^2$), and the equation reduces to the familiar Newtonian Poisson equation

$$\boxed{k^2\,\Phi = -4\pi G\,a^2\,\bar{\rho}\,\delta, \qquad \text{i.e.} \qquad \nabla^2\Phi = 4\pi G\,a^2\,\bar{\rho}\,\delta} \tag{651}$$

(the factor a^2 converts the comoving Laplacian to physical units). An overdensity, $\delta > 0$, gives $\Phi < 0$ and hence $\Psi = \Phi < 0$ in this convention; through the Euler equation it pulls matter inward, deepening the well — the instability that drives all of structure formation. We have, in (648), (649) and (651), a closed linear system for (δ, θ, Φ).

J.3 The sub-horizon growth equation

We now combine the three boxed equations into a single second-order equation for δ. Work in the sub-horizon, quasi-static regime ($k \gg \mathcal{H}$, $\Phi' \ll \mathcal{H}\Phi$), in which (648) reads $\delta' = -\theta$.

Step 1: differentiate the continuity equation. Differentiating $\delta' = -\theta$ once more with respect to conformal time,

$$\delta'' = -\theta'. \tag{652}$$

Step 2: eliminate θ' with the Euler equation. From (649), $\theta' = -\mathcal{H}\theta + k^2\Psi$. Substitute, then replace $\theta = -\delta'$ from continuity:

$$\delta'' = \mathcal{H}\theta - k^2\Psi = -\mathcal{H}\,\delta' - k^2\Psi. \tag{653}$$

Step 3: eliminate Ψ with Poisson. On sub-horizon scales $\Psi = \Phi$ and (651) gives $k^2\Psi = -4\pi G\,a^2\,\bar{\rho}\,\delta$. Hence

$$\boxed{\delta'' + \mathcal{H}\,\delta' - 4\pi G\,a^2\,\bar{\rho}_m\,\delta = 0.} \tag{654}$$

This is the growth equation in conformal time, with $\bar{\rho}_m$ the mean *matter* density (only non-relativistic, clustering matter sources the well on these scales). It is worth pausing on its structure: a harmonic-oscillator-like equation with a friction term $\mathcal{H}\delta'$ (Hubble drag) and an *anti-restoring* force $-4\pi G a^2 \bar{\rho}_m\,\delta$ — the wrong sign for oscillation, which is exactly why the solution grows rather than oscillates. There is no k-dependence: in the matter-dominated, sub-horizon, cold regime all modes grow at the *same* rate, a scale-independence that is the reason a single growth factor $D(a)$ can describe the whole linear power spectrum.

Conversion to cosmic time. Many texts write the equation in cosmic time t and against the physical density. Using $\delta' = a\dot{\delta}$, $\delta'' = a^2\ddot{\delta} + a\dot{a}\,\dot{\delta} = a^2\ddot{\delta} + a^2 H\dot{\delta}$, and $\mathcal{H} = aH$, equation (654) becomes, after dividing by a^2,

$$\boxed{\ddot{\delta} + 2H\,\dot{\delta} - 4\pi G\,\bar{\rho}_m\,\delta = 0.} \tag{655}$$

This is the form quoted in the body of the paper. Note the friction coefficient is now $2H$ (not H): the change from $\mathcal{H}$ to $2H$ is an artifact of the time variable, since $\mathcal{H}\delta' = aH \cdot a\dot{\delta}$ contributes one power of H and the cross term from δ'' contributes the second.

J.4 The growth factor $D(a)$ and its exact matter-dominated solutions

Because (655) is linear and homogeneous and its coefficients depend only on time, every Fourier mode shares the same time dependence: $\delta(\eta, \mathbf{k}) = D(a)\,\delta_0(\mathbf{k})$ up to the decaying mode, where $D(a)$ is the *linear growth factor* and $\delta_0(\mathbf{k})$ is the initial spatial pattern. It is cleaner to use a as the independent variable. With $\dot{\delta} = aH\,(dD/da)\,\delta_0$ and the Friedmann/continuity relations folded in through dH/da, equation (655) becomes the standard growth ODE in a,

$$\boxed{\frac{d^2D}{da^2} + \left(\frac{3}{a} + \frac{1}{H}\frac{dH}{da}\right)\frac{dD}{da} - \frac{3}{2}\,\frac{\Omega_m(a)}{a^2}\,D = 0,} \tag{656}$$

where $\Omega_m(a) \equiv 8\pi G\bar{\rho}_m/(3H^2)$ is the matter density parameter at scale factor a, and we used $4\pi G\bar{\rho}_m = \frac{3}{2}\Omega_m(a)\,H^2$. This is the equation solved numerically in any Boltzmann/growth code.

Exact Einstein–de Sitter solution. In a flat, matter-only (Einstein–de Sitter, EdS) universe, $\Omega_m(a) = 1$ and $H^2 \propto a^{-3}$, so $H^{-1}dH/da = -\frac{3}{2}a^{-1}$ and the bracket in (656) is $(3/a) - (3/2a) = (3/2)a^{-1}$. The equation reduces to

$$\frac{d^2D}{da^2} + \frac{3}{2a}\frac{dD}{da} - \frac{3}{2a^2}\,D = 0, \tag{657}$$

an Euler equation solved by power laws $D \propto a^n$ with $n(n-1) + \frac{3}{2}n - \frac{3}{2} = 0$, i.e. $n^2 + \frac{1}{2}n - \frac{3}{2} = 0$, with roots $n = +1$ and $n = -3/2$. Hence

$$D_+(a) \propto a \qquad \text{(growing mode)}, \qquad D_-(a) \propto a^{-3/2} \qquad \text{(decaying mode)}. \tag{658}$$

The growing mode $D_+ \propto a$ is the workhorse fact of structure formation: *in a matter-dominated universe the density contrast grows linearly in the scale factor.* The decaying mode is the transient that any sensible initial condition quickly forgets.

General ΛCDM growth: the integral form. When dark energy is present, $\Omega_m(a) < 1$ and the growth is slower than a at late times (dark energy accelerates the expansion, strengthening the Hubble drag without sourcing the well). The growing solution of (656) admits the closed integral form

$$D(a) \;\propto\; H(a)\int_0^a \frac{da'}{[\,a'H(a')\,]^3}, \qquad H^2(a) = H_0^2\,[\Omega_m a^{-3} + \Omega_\Lambda] \tag{659}$$

(flat ΛCDM, radiation neglected at low z). It is conventional to *normalize* the growth factor so that $D = a$ in the matter-dominated limit (early times); a common alternative normalization sets $D(a_0) = D(1) = 1$ today. We will use the growth-suppression factor $g(a) \equiv D(a)/a$, equal to unity in EdS and falling below unity once dark energy bites.

Worked example: the growth factor today in ΛCDM. Take flat ΛCDM with $\Omega_m = 0.315$, $\Omega_\Lambda = 0.685$. A compact and accurate closed form for the suppression factor is the Carroll–Press–

Turner fit,

$$g(a{=}1) = \frac{D(1)/1}{D_{\rm EdS}} \simeq \frac{5}{2}\Omega_m\Big[\Omega_m^{4/7} - \Omega_\Lambda + (1+\tfrac{1}{2}\Omega_m)(1+\tfrac{1}{70}\Omega_\Lambda)\Big]^{-1}. \tag{660}$$

Evaluating with $\Omega_m = 0.315$: $\Omega_m^{4/7} = 0.315^{0.5714} = \exp[0.5714\ln 0.315] = \exp[-0.660] = 0.517$; the bracket is $0.517 - 0.685 + (1.1575)(1.00979) = 0.517 - 0.685 + 1.169 = 1.001$; and $g(1) = 2.5 \times 0.315/1.001 = 0.7875/1.001 \approx 0.787$. Thus the growing mode at $a = 1$ has been suppressed to about 79% of the value it would have reached in a matter-only universe with the same early normalization: dark energy has slowed structure growth by roughly a fifth since it came to dominate. (For orientation, the same fit gives $g \approx 0.91$ at $z = 1$, where $\Omega_m(a) = 0.77$ is still close to unity — growth was nearly unsuppressed at $z = 1$ and has been throttled mainly in the last few billion years.)

J.5 The Meszaros effect: suppression during radiation domination

Before matter–radiation equality ($a < a_{\rm eq}$, redshift $z_{\rm eq} \approx 3400$) the universe is radiation-dominated, and a crucial thing happens to sub-horizon matter perturbations: they *almost stop growing*. This is the *Meszaros effect*, and it sets the shape of the matter power spectrum on small scales.

Why growth stalls. The growth equation (654) sources δ through the matter density $\bar\rho_m$. In the radiation era the total energy budget, and hence the expansion rate $\mathcal{H}$ that feeds the friction term, is dominated by *radiation*, while the gravitational source is dominated by *matter* only when $\bar\rho_m$ is appreciable. For a sub-horizon CDM mode the radiation does not cluster (its pressure support gives it a Jeans length of order the horizon, so it free-streams and oscillates as acoustic waves), so it contributes to the friction but not to the driving. The matter perturbation feels strong Hubble drag from the fast radiation-driven expansion but only weak self-gravity. The result is logarithmic, not power-law, growth.

The Meszaros equation and its solution. Introduce the variable $y \equiv a/a_{\rm eq} = \bar\rho_m/\bar\rho_r$, the ratio of matter to radiation density. In terms of y the sub-horizon CDM growth equation (656), with the two-component $H^2 \propto (1+y)/a^4$ expansion, becomes the *Meszaros equation*

$$\frac{d^2\delta}{dy^2} + \frac{2+3y}{2y(1+y)}\frac{d\delta}{dy} - \frac{3}{2y(1+y)}\,\delta = 0. \tag{661}$$

This has the exact growing solution

$$\boxed{\delta_+(y) \;\propto\; 1 + \tfrac{3}{2}\,y.} \tag{662}$$

One verifies by direct substitution that (662) solves (661). The lesson is in the two limits. Deep in radiation domination, $y \ll 1$, the contrast is $\delta_+ \to \text{const} + \frac{3}{2}y$: it grows by at most a factor of order unity, i.e. it is *frozen* up to a logarithm once horizon-crossing effects are added. Deep in matter domination, $y \gg 1$, $\delta_+ \to \frac{3}{2}y \propto a$, recovering the EdS $D_+ \propto a$ law of (658). The total suppression a

small-scale mode suffers between horizon entry (in the radiation era) and equality is therefore at most a few — but it is *scale dependent*, because smaller-scale modes enter the horizon earlier and stall for longer. This scale-dependent suppression is exactly what bends the matter power spectrum from its primordial slope to its observed turnover, and it is encoded in the transfer function below.

J.6 The growth index: $f \approx \Omega_m^{0.55}$

A single number summarizes the instantaneous growth rate: the *linear growth rate*

$$f \;\equiv\; \frac{d\ln\delta}{d\ln a} = \frac{d\ln D}{d\ln a} = \frac{a}{D}\frac{dD}{da}. \tag{663}$$

In EdS, $D \propto a$ gives $f = 1$ exactly. In ΛCDM, f falls below unity at late times. A remarkably accurate fitting formula, due to Peebles and refined by Wang–Steinhardt and Linder, expresses f as a power of the instantaneous matter fraction,

$$\boxed{\; f(a) \;\approx\; \Omega_m(a)^{\gamma}, \qquad \gamma \approx 0.55 \;} \tag{664}$$

for ΛCDM (more precisely $\gamma \simeq 0.55 + 0.05[1 + w(z=1)]$ for a dark-energy equation of state w; deviations from $\gamma = 0.55$ are a clean probe of modified gravity, where γ can shift to ~ 0.68). The combination actually measured by redshift-space distortions is $f\sigma_8(z)$, since the survey sees the product of the growth rate and the clustering amplitude; we return to σ_8 below.

Worked check of the growth index today. At the present epoch, with the canonical Planck value $\Omega_m = \Omega_m(a_0) \approx 0.315$,

$$f_0 \approx (0.315)^{0.55} = \exp[\,0.55\,\ln 0.315\,] = \exp[\,0.55\times(-1.155)\,] = \exp[-0.635] \approx 0.530. \tag{665}$$

So structure grows today at a little over half the EdS rate: the density contrast increases by about 0.53% for each 1% increase in the scale factor. This is the standard ΛCDM value, confirmed by redshift-space-distortion measurements of $f\sigma_8$ to the ten-percent level.

J.7 Transfer function and the matter power spectrum $P(k)$

From a single mode to a statistical field. The density field $\delta(\mathbf{x})$ is a random field; what we predict and measure is its *statistics*, not its realization. For a Gaussian field, all information is in the two-point function, whose Fourier transform is the *power spectrum* $P(k)$ defined by

$$\langle \delta_{\mathbf{k}}\, \delta^*_{\mathbf{k}'} \rangle = (2\pi)^3\, \delta_D^{(3)}(\mathbf{k}-\mathbf{k}')\, P(k), \tag{666}$$

where $\delta_D^{(3)}$ is the Dirac delta enforcing statistical homogeneity (so P depends only on $k = |\mathbf{k}|$ by isotropy).

Primordial spectrum, transfer function, growth. Inflation lays down a nearly scale-invariant primordial curvature spectrum, $P_{\mathcal{R}}(k) \propto k^{\,n_s - 1}$ with $n_s \approx 0.965$ (the scalar tilt; see the inflationary

appendix). The late-time *matter* power spectrum is this primordial spectrum processed by two factors: the *transfer function* $T(k)$, which encodes everything that happens between horizon entry and equality (most importantly the Meszaros suppression of small-scale modes), and the growth factor $D(a)$, which carries the linear evolution after equality. Schematically,

$$\boxed{P(k,a) \;=\; A\,k^{n_s}\,T^2(k)\,D^2(a),} \tag{667}$$

where A fixes the overall amplitude (tied to the primordial A_s). The shape of $T(k)$ is the cause of the famous turnover in $P(k)$:

$$T(k) \;\to\; \begin{cases} 1, & k \ll k_{\rm eq} \quad \text{(large scales, no suppression)}, \\ (k_{\rm eq}/k)^2\,\ln k, & k \gg k_{\rm eq} \quad \text{(small scales, Meszaros-suppressed)}, \end{cases} \tag{668}$$

with $k_{\rm eq} \sim 0.01\,h\,{\rm Mpc}^{-1}$ the comoving wavenumber entering the horizon at equality. Thus $P(k) \propto k^{n_s}$ rises on large scales and bends to $P(k) \propto k^{n_s-4}\ln^2 k$ on small scales, turning over near $k_{\rm eq}$. Baryon acoustic oscillations (BAO) imprint a damped oscillatory ripple on top of this smooth shape; they are the frozen sound waves of the photon–baryon plasma and provide the standard ruler used in the Hubble-tension Part.

J.8 From $P(k)$ to σ_8 and S_8

The variance σ_8. To compress $P(k)$ into one number comparable across surveys, one measures the variance of the density field smoothed on a fixed physical scale. The standard choice is a top-hat sphere of comoving radius $R = 8\,h^{-1}\,{\rm Mpc}$ (a scale near the transition to non-linearity, historically chosen because the smoothed variance is of order unity there). The smoothed variance is

$$\sigma_R^2 \;=\; \int_0^\infty \frac{dk}{k}\,\frac{k^3 P(k)}{2\pi^2}\,\big|W(kR)\big|^2, \qquad W(x) = \frac{3\,(\sin x - x\cos x)}{x^3}, \tag{669}$$

where $W(kR)$ is the Fourier transform of the real-space top-hat window and $k^3P(k)/2\pi^2 \equiv \Delta^2(k)$ is the dimensionless power per logarithmic interval in k. Setting $R = 8\,h^{-1}\,{\rm Mpc}$ defines

$$\sigma_8 \;\equiv\; \sigma_{R=8\,h^{-1}{\rm Mpc}}. \tag{670}$$

Because $\sigma_8 \propto D(a_0)$ through (667), it inherits the full late-time growth history: anything that suppresses or enhances D shifts σ_8.

The combination S_8. Weak-lensing surveys do not measure σ_8 and Ω_m separately; their two-point shear statistics constrain a degenerate combination, well approximated by the rescaled amplitude

$$\boxed{S_8 \;\equiv\; \sigma_8\,\sqrt{\Omega_m/0.3}.} \tag{671}$$

The factor $\sqrt{\Omega_m/0.3}$ rotates onto the direction *orthogonal* to the lensing degeneracy, so S_8 is the quantity actually pinned down by surveys such as KiDS-1000 and DES (Asgari et al., 2021). The " S_8 tension " is the statement that lensing surveys prefer a *lower* S_8 (looser clustering, ~ 0.76) than the value $S_8 \approx 0.83$ *extrapolated forward* from the Planck CMB under ΛCDM (Planck Collaboration, 2018) — a 2–3σ mismatch in the late-time growth amplitude.

Worked example: $\sigma_8 \to S_8$ with $\Omega_m = 1/\pi$. The framework's conjectured matter fraction is the π-locked value $\Omega_m = 1/\pi \approx 0.31831$ (carried as Derived-conditional in the codebook). Take the Planck-anchored early-branch clustering amplitude $\sigma_8 = 0.811$. Then

$$\sqrt{\Omega_m/0.3} = \sqrt{(1/\pi)/0.3} = \sqrt{0.31831/0.3} = \sqrt{1.06103} = 1.03007,$$
$$S_8 = \sigma_8 \sqrt{\Omega_m/0.3} = 0.811 \times 1.03007 = 0.8354. \tag{672}$$

So the π-locked matter fraction maps $\sigma_8 = 0.811 \mapsto S_8 \approx 0.835$, essentially the Planck-extrapolated value (the body quotes $S_8 \approx 0.832$ for $\Omega_m = 0.315$; the tiny difference, 0.835 versus 0.832, is just the $1/\pi$ versus 0.315 shift in Ω_m, since $\sqrt{0.31831/0.315} = 1.0052$ scales the amplitude up by half a percent). This is the *high* branch, the one in tension with weak lensing; the framework's computed value is lower, $\sigma_8 = 0.80 \to S_8 = 0.825$ (§43.3), and this example shows only the Ω_m factor; the retired corridor proposal aimed to lower a late branch toward $S_8^{\rm late} \approx 0.78$, and the record of that proposal, sign and all, is Section J.10.

J.9 What the standard growth theory leaves open: the S_8 tension

The linear theory of Step 1 is exact: given the background expansion history $H(a)$ and the matter density, it predicts the growth factor $D(a)$, the power spectrum, and the clustering amplitudes σ_8 and S_8 without ambiguity. As in Appendix E, it is worth naming what it must take as given, and where it runs into trouble.

1. **The background and the inventory.** The expansion history $H(a)$ and the density Ω_m that drive the growth are measured inputs, not derived.

2. **The S_8 tension.** The clustering amplitude $S_8 = \sigma_8\sqrt{\Omega_m/0.3}$ *extrapolated* from the Planck CMB ($S_8 \approx 0.83$) sits above the value *measured* directly by weak-lensing and galaxy surveys ($S_8 \approx 0.78$) at the two-to-three sigma level. Standard ΛCDM offers no structural account of the discrepancy; it is a live tension, the growth-sector counterpart of the Hubble tension.

Standard growth theory, then, is an exact calculator of clustering once the background is supplied — but the background is an input, and the S_8 tension sits unexplained. The geometric account asks whether the same structure that reframes the H_0 tension also reaches the growth amplitude.

J.10 The retired geometric derivation: the two-stage corridor and the sign of S_8 (record)

Record. The subsection below is the growth mechanism of the two-stage corridor as this appendix stated it before the corridor was retired (Part VI, §42; Lemma II of the codebook). It is kept verbatim because later text refers to its equations by label, and because a withdrawn computation is a record, not an embarrassment. On the one background the paper now carries, the growth history is the standard one and σ_8 is *computed* from the declared amplitude (§43.3); there is no corridor growth response and no open sign. Nothing below is claimed.

The linear theory of Step 1 computes the growth of structure exactly, once the background expansion history $H(a)$ is supplied. The geometric content enters precisely there: the same two-stage corridor that reframes the Hubble tension in Part VI reshapes $H(a)$, and through the growth equation (656) it reshapes the late-time clustering — so that the S_8 and H_0 tensions become two faces of one structure. The Hubble-tension Part of the body replaces the single late-time H_0 with a *two-stage corridor*: a brief, early, throat/Σ_r-sector source correction (parametrized as an early-dark-energy-like lobe of fractional height $f_{\rm EDE} \approx 0.03$ centered near $z_c \sim 3500$, plus a small extra-radiation contribution $\Delta N_{\rm eff} \approx 0.30$) that shifts the sound horizon and the equality timing, followed by a late-time branch that the local distance ladder reads as $H_0^{\rm late} \approx 73.5$ — the legacy two-branch reading, retained in this appendix as the record of the earlier version and superseded by the one-background reading of Theorem 68.8 (C3); the companion Arisaka (2026, C2-Hubble) retired its own early-dark-energy corridor in September 2026 and carries one background with the ladder as a measurement bearing a named systematic, the reading of §42.6.

How the corridor touches growth. A modified $H(a)$ enters the growth equation in two places: directly, through the friction coefficient $H^{-1}dH/da$ in (656), and indirectly, through $\Omega_m(a) = 8\pi G\bar{\rho}_m/3H^2$ in the source. The corridor parametrization *assumes* that the net effect of the early lobe plus the late branch is to *suppress* the late-time growth factor $D(a)$ relative to the Planck-extrapolated ΛCDM value, so that σ_8 and hence S_8 would come down from the high (≈ 0.835) branch toward the weak-lensing-preferred $S_8^{\rm late} \approx 0.78 \pm 0.01$. *This subsection is the retired reading, kept for the record.* Under that assumption the S_8 tension would have been reconciled by the *same* corridor that addressed the Hubble tension, with no new ingredient: the early branch yielding the high Planck amplitude, the late branch the lower lensing amplitude, and the "tension" becoming two projections of one structure. The response coefficients that convert the corridor parameters $(f_{\rm EDE}, \Delta N_{\rm eff}, H_0^{\rm late})$ into a growth-suppression ΔS_8 are taken from the early-dark-energy literature (Poulin et al., 2019); they are *not* derived here from the throat/Σ_r sector.

The sign is OPEN — stated without softening. Here is the honest difficulty, and the appendix states it as sharply as the body does. The suppression *branch* just described is an *assumption* of the parametrization, not a result. In full likelihood analyses of early-dark-energy-type models, refitting the CMB with the early lobe present drives the cold-dark-matter density $\omega_{\rm cdm}$ and the scalar tilt n_s *upward* in order to preserve the well-measured acoustic peaks; those upward shifts *increase* small-scale power, and the *net* S_8 shift typically comes out *positive* — the *opposite* sign to

the suppression encoded above (Hill et al., 2020). A positive shift does not relieve the S_8 tension; it *worsens* it. The corridor therefore faces a genuine sign ambiguity:

- the corridor parametrization *posits* a suppression, $\Delta S_8 < 0$, which would help;
- the leading EDE likelihood evidence (Hill et al., 2020) finds $\Delta S_8 > 0$, which would hurt;
- which sign the throat/Σ_r corridor actually delivers cannot be decided from the parametrization alone — it requires the full growth-history integration, with the CMB refit done self-consistently, that is named as the corridor-uniqueness research target.

Until that integration is carried out, the S_8 corollary of the corridor is *conditional* and its sign is *open*. We do not claim the S_8 tension is resolved. We claim that the *same* geometric structure that reframes the Hubble tension *also* touches the growth amplitude, and that whether it helps or hurts is an honest open question pinned to a declared calculation.

What would settle it. The resolution is conditional on the corridor-uniqueness target: a self-consistent Boltzmann/likelihood fit (CMB + BAO + distance ladder + weak-lensing $f\sigma_8$) in which the throat/Σ_r source correction is propagated through recombination, equality, and late-time growth in one pass, so that the sign and magnitude of ΔS_8 are *outputs* rather than inputs. The relevant discriminating data are weak-lensing growth across redshift, cluster counts, and CMB-lensing cross-checks — the same observables that sharpen the tension itself.

Status note

Status (standard linear growth): Derived. The full linear machinery of this appendix — the conformal-Newtonian perturbed metric (642), the perturbed continuity, Euler and Poisson equations (648), (649), (651), the sub-horizon growth equation $\delta'' + \mathcal{H}\delta' - 4\pi G a^2 \bar{\rho}_m \delta = 0$ and its cosmic-time form $\ddot{\delta} + 2H\dot{\delta} - 4\pi G \bar{\rho}_m \delta = 0$ (654)–(655), the growth factor $D(a)$ with its EdS solution $D_+ \propto a$ (658) and ΛCDM integral (659), the Meszaros suppression $\delta_+ \propto 1 + \frac{3}{2}y$ (662), the growth index $f \approx \Omega_m^{0.55}$ (664), the transfer function and power spectrum $P(k) = Ak^{n_s}T^2(k)D^2(a)$ (667), and the variance/amplitude pair σ_8, $S_8 = \sigma_8\sqrt{\Omega_m/0.3}$ (670), (671) — is *Derived*: it is standard general relativity plus a perfect-fluid source, taught here in full and *unmodified* by the framework. The boundary Einstein equation it rests on is itself a theorem of the 6D-to-4D OGGEP projection of Arisaka (2026, A5-GDOS); GG-A10-Cosmos inherits it and does not re-derive it. The worked examples are fixed numbers: the ΛCDM growth suppression $g(a{=}1) \approx 0.79$ (660), the growth index $f_0 \approx (0.315)^{0.55} \approx 0.53$ (665), and the π-locked mapping $\sigma_8 = 0.811 \mapsto S_8 \approx 0.835$ at $\Omega_m = 1/\pi$ (672).

Status (S_8 corridor resolution): WITHDRAWN. The GG-Theory-specific claim that the two-stage corridor of the Hubble-tension Part also reconciles the σ_8/S_8 tension by modifying the *late-time* growth history is withdrawn with the corridor: on one background there is no growth response, and the framework's $S_8 = 0.825 \pm 0.010$ is computed (§43.3) and *shares* the tension with cosmic shear at $+2.5\sigma$. What follows is the earlier grading, kept as the record of what was claimed and why its sign was open. The corridor parametrization *assumes* a growth *suppression* ($\Delta S_8 < 0$,

which would help the tension); but in full early-dark-energy likelihood fits the self-consistent CMB refit drives $\omega_{\rm cdm}$ and n_s upward and the *net* S_8 shift typically comes out *positive* — the *opposite* sign, which would *worsen* the tension (Hill et al., 2020). The response coefficients that convert corridor parameters into ΔS_8 are literature values (Poulin et al., 2019), not yet derived from the throat/Σ_r sector. The resolution is therefore *conditional on the corridor-uniqueness research target*: a self-consistent CMB + BAO + distance-ladder + weak-lensing growth integration in which the sign and magnitude of ΔS_8 emerge as outputs. **We do not claim the S_8 tension is resolved; the sign is open.**

J.11 Legacy growth response, cross-consistency Jacobian, and worked derivations (Part VI)

The material below is collected here rather than inline in Part VI, so the body stays readable, as was done for Parts IV and V. It is a record: the S_8 suppression formula is the legacy early-dark-energy linear response, superseded twice over — first by a corridor reading, then by the direct computation of the amplitude from the declared A_s (§43.3), the corridor having been retired — and nothing in it is claimed; the cross-consistency Jacobian and the worked derivation notebook give the from-scratch calculations behind numbers Part VI once carried.

J.11.1 Legacy S_8 growth-suppression linear-response

The present framework acknowledges this standard difficulty rather than denying it. The narrow EDE-plus-dark-radiation pair is expected to suppress broad linear growth slightly. The question is whether that suppression lands in the right direction and whether the same framework provides a compensating or reshaping effect in the lensing sector (the high-k minihalo route once claimed here is withdrawn, §44.3).

A practical linear-response expression captures the intended behavior:

$$S_8 \simeq S_8^{\rm ref}\left(1 - \alpha_S \Delta N_{\rm eff} - \beta_S f_{\rm EDE} + \gamma_S \alpha_L\right), \tag{673}$$

with representative coefficients chosen to encode the sign structure,

$$\alpha_S \approx 0.12, \qquad \beta_S \approx 0.80, \qquad \gamma_S \approx 0.20. \tag{674}$$

These numbers are not claimed to be the final output of a full likelihood pipeline. They are a compact summary of how the present mechanism is supposed to behave: extra radiation and early dark energy lower S_8, while the mild enhancement of the effective lensing amplitude partly compensates.

A cautionary note is mandatory here: in full likelihood analyses of early-dark-energy-type models, refitting the CMB typically drives $\omega_{\rm cdm}$ and n_s upward and the *net* S_8 shift comes out positive — the opposite sign to the suppression encoded above (Hill et al., 2020). The parametrization adopted the suppression branch as a target; on the one background the paper now carries there is no growth response, and the amplitude is computed rather than parametrized.

If one inserts the same inherited values used in the Hubble section,

$$\Delta N_{\rm eff} = 0.30, \qquad f_{\rm EDE} = 0.03, \qquad \alpha_L = 0.05, \tag{675}$$

then the total fractional shift is

$$\epsilon_{S_8} \equiv \alpha_S \Delta N_{\rm eff} + \beta_S f_{\rm EDE} - \gamma_S \alpha_L \approx 0.036 + 0.024 - 0.010 = 0.050. \tag{676}$$

Thus,

$$S_8 \approx 0.95\, S_8^{\rm ref}. \tag{677}$$

Starting from a representative early-universe reference value in the low 0.8's then naturally lands one in a late-time corridor around

$$S_8 \approx 0.78 \pm 0.01, \tag{678}$$

which is precisely the direction favored by the low-redshift lensing literature — and which the paper no longer claims: on one background the amplitude is computed from the declared A_s and comes out at $S_8 = 0.825$, sharing the S_8 tension rather than resolving it (§43.3). What follows is the retired calculation. The benchmark cosmic-shear determination, $S_8 = 0.759^{+0.024}_{-0.021}$ from KiDS-1000 (Asgari et al., 2021), agreed with the corridor at the 0.9σ level (a +2.8% fractional deviation of the prediction from the measurement). In the parameter-audit Dev plot we therefore display two S_8 entries, mirroring the two Hubble branches: the anchored early/CMB value $S_8^{\rm CMB}$ and the predicted late/lensing branch $S_8^{\rm late}$ (Figure 31, Table 41).

J.12 Cross-consistency and falsifiability inside the growth block

Record, with one live condition. Conditions A–C below constrain the parameter vector of the retired two-stage model; on the one background the paper now carries their early residue and growth suppression are zero identically, and they are kept as the record of the consistency discipline the earlier reading was held to. Condition D — that the first-light sector be scale-selective, a changed collapse threshold and not a global rescaling of the matter power — is the one condition that survives the retirement, and Part VI cites it as live.

J.12.1 The minimal Part VI parameter vector

In this paper the entire growth-and-tension block is intentionally compressed to a minimal effective vector,

$$\mathbf{p}_{\rm VI} = (f_{\rm EDE},\ \Delta N_{\rm eff},\ \alpha_L,\ \epsilon_{\rm local},\ \epsilon_{\rm str},\ \delta_c)\,, \tag{679}$$

where

$$f_{\rm EDE} \approx 0.03, \qquad \Delta N_{\rm eff} \approx 0.30, \qquad \alpha_L \approx 0.05, \tag{680}$$

$$\epsilon_{\rm local} \sim 0.03\text{–}0.05, \qquad \epsilon_{\rm str} \equiv 1 - \frac{v_{\rm str}}{v_{\rm str}^{\Lambda\rm CDM}} \approx 0.20, \qquad \delta_c \approx 1.18. \tag{681}$$

(The $\Delta N_{\rm eff} \approx 0.30$ printed here is the legacy linear-response value belonging to this appendix; the body and the codebook audit now carry the derived null $\Delta N_{\rm eff} = 0$, and the numbers quoted below are computed at the intermediate 0.04 at which this appendix's comparison was made.)

Three comments are immediate.

First, $f_{\rm EDE}$, $\Delta N_{\rm eff}$, and α_L are the early-to-middle-universe entries. They touch the CMB, the sound horizon, and broad linear growth. Second, $\epsilon_{\rm local}$ is the late-time environmental entry relevant mainly for the local Hubble route. Third, $\epsilon_{\rm str}$ and δ_c are selective first-light entries rather than broad background ones. This separation is what prevents the framework from being an undifferentiated set of rescalings.

To first order, the observable vector

$$\mathbf{o}_{\rm VI} = (\theta_*,\ H_0,\ S_8,\ \ell_D,\ A_L,\ n_{\rm gal},\ M_{\rm seed}) \tag{682}$$

responds through a Jacobian,

$$\delta\mathbf{o}_{\rm VI} \approx \mathbf{J}_{\rm VI}\,\delta\mathbf{p}_{\rm VI}. \tag{683}$$

The exact entries of $\mathbf{J}_{\rm VI}$ require Boltzmann and hydrodynamical likelihood work and do not belong to this paper. But the sign pattern does belong here, because it is the most compact way to state what the theory is and is not allowed to do.

J.12.2 Sign matrix and one-ledger discipline

The sign matrix may be written schematically as follows:

$$\mathbf{J}_{\rm VI} \sim \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ + & + & 0 & + & 0 & 0 \\ - & - & + & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & + & 0 & 0 & 0 \\ - & - & 0 & 0 & + & - \\ 0 & 0 & 0 & 0 & + & - \end{pmatrix}, \tag{684}$$

where the rows correspond to $(\theta_*, H_0, S_8, \ell_D, A_L, n_{\rm gal}, M_{\rm seed})$ and the columns to $(f_{\rm EDE}, \Delta N_{\rm eff}, \alpha_L, \epsilon_{\rm local}, \epsilon_{\rm str}, \delta_c)$. The matrix is schematic, not numerical, but it already captures the theory's architecture:

1. The acoustic angle is treated as observationally fixed and therefore has zero first-order freedom in the fit space; changes in the sound horizon are compensated elsewhere.
2. H_0 is pushed upward by the early pair and by the local lift.
3. S_8 is pushed downward by the early pair and slightly upward in the CMB-lensing interpretation by α_L.
4. The damping scale ℓ_D is tied to the BFFT/OGN ledger and is therefore not independently free in Part VI.

5. The early-galaxy abundance is helped by a lowered collapse threshold (the streaming damping having no mediator, §44.4), and is *hurt* by the early pair. Earlier the $n_{\rm gal}$ row carried + in the $f_{\rm EDE}$ column; the computation of §44.3 makes both early-pair entries −, and they are printed that way above. The deficit is not small: at $\nu \sim 5$ the residue's $\Delta\sigma/\sigma = -0.0185$ reduces the earliest minihalo abundance by about two fifths, which the one baryonic lever must overcome before it advances anything.

6. The heavy-seed corridor is helped primarily by the same baryonic first-light pair $(\epsilon_{\rm str}, \delta_c)$.

This matrix is useful because it forbids a common failure mode in ambitious cosmological model building: letting the sign assignments drift from one anomaly to the next. Once the vector $\mathbf{p}_{\rm VI}$ is fixed, each correction has a role. It is not allowed to be helpful for one tension only by becoming a new free knob elsewhere.

J.12.3 Consistency inequalities that must hold simultaneously

The previous sign matrix may be compressed into a set of inequality-style consistency conditions. The purpose of these inequalities is not to replace a full data analysis, but to state the shape of the allowed corridor.

Condition A: the CMB must remain near the observed acoustic solution. Because the observed θ_* is extremely well determined, the combined early correction must preserve the acoustic fit,

$$\left|\frac{\delta\theta_*}{\theta_*}\right| \ll 1. \tag{685}$$

Operationally, this means that any shift in r_s must be compensated by the corresponding distance shift without spoiling the peak sequence or damping tail. The theory is not allowed to move H_0 by destroying the primary CMB fit.

Condition B: the Hubble lift must be large enough to matter, but not so large that the source residue becomes implausible. If the corrected early-time shift is too small, the module does nothing. If it is too large, the required source residue will typically contradict the CMB or BAO. Thus one needs an intermediate corridor,

$$0.03 \lesssim \epsilon_{\rm early} \lesssim 0.07, \tag{686}$$

which is precisely where the working values of Eq. (750) land.

Condition C: the broad growth suppression must be enough to lower S_8, but not so large that the matter sector becomes under-clustered. One therefore wants

$$0.03 \lesssim \epsilon_{S_8} \lesssim 0.08, \tag{687}$$

which is again consistent with the working value $\epsilon_{S_8} \approx 0.05$.

Condition D: the first-light sector must be scale-selective. The same model cannot simply suppress or enhance power uniformly at all scales. Instead, one needs

$$\text{(broad linear growth)} \neq \text{(high-}\nu\text{ rare-halo response)}\,, \tag{688}$$

which is realized here through a changed δ_c alone — the streaming damping having been withdrawn — and not through a global matter-power rescaling.

These conditions are not optional decorations. They are the mathematical form of the phrase "one ledger." If a proposed correction fails one of these inequalities, it cannot be saved by shifting the burden to another section of the paper.

J.12.4 A worked one-line closure check

Read this closure check with two riders attached, both established below. Its $\epsilon_{\rm early}$ is a linear-response number that presumes a self-consistent CMB refit; the direct background integration of §K.10.5 returns about two fifths of it. And its minihalo entry has the wrong sign: the same residue suppresses the rare tail by the same k-independent factor by which it suppresses the broad channel (§44.3), so the $+0.4$ below is withdrawn and the computed value is about -0.44. The check is retained because the *correlation* it displays is real and instructive; its third row is not.

It is useful to perform a compact closure check with the working numbers of Part VI. Using Eqs. (750) and (673), one finds

$$\epsilon_{\rm early} \approx 0.051, \tag{689}$$

$$\epsilon_{S_8} \approx 0.050. \tag{690}$$

Thus the same narrow source pair moves the acoustic-calibrated H_0 and the broad growth amplitude by roughly five percent in opposite observational directions — upward for H_0, downward for S_8. This is exactly the correlated sign pattern one wants.

The first-light pair then adds

$$\frac{M_f^{\rm new}}{M_f^{\rm old}} \sim 0.512, \tag{691}$$

$$\frac{\Delta n_{\rm minihalo}}{n_{\rm minihalo}} \sim 0.4 \quad \textit{(withdrawn; see below)}, \tag{692}$$

$$\frac{n_{\rm massive}^{\mathsf{GG\text{-}6}}}{n_{\rm massive}^{\Lambda\rm CDM}} \sim 10^2\text{–}10^3, \tag{693}$$

which is more than enough to make the earliest bright objects much less surprising. The crucial point is that the numbers line up directionally without introducing a second independent dark-matter model or a second independent inflationary module.

J.12.5 What a full numerical pipeline would have to do next (record)

A serious continuation of this work would need to replace the symbolic Jacobian by an actual numerical pipeline. That pipeline has three main stages.

1. **Boltzmann stage:** implement the EDE residue, the $\Delta N_{\rm eff}$ shift, and the lensing correction in a precision perturbation solver and fit the full TT, TE, EE, and lensing data products.

2. **Background-and-distance stage:** propagate the corrected early solution through BAO and distance-ladder inference to test whether the Hubble corridor survives outside the CMB alone.

3. **Hydrodynamical first-light stage:** insert the altered collapse threshold into first-galaxy and direct-collapse-seed simulations. (Two items that used to head this list are withdrawn — the small-scale transfer correction, §44.3, and the streaming damping, §44.4; a simulation program should not be built on either.)

The present paper does not perform those computations in full. What it does do is define exactly what must be computed and exactly which correlations should appear if the framework is right.

J.12.6 The kill table for Part VI as a whole (record)

It is useful to place the internal failure logic into one table.

Table 58: Record: kill conditions for the retired Part VI growth block, kept as written. The philosophy is simple: a unified account cannot claim the successes of its parts jointly and then survive their failures separately, so the conditions are written against the block as a whole. The columns give the block, what must happen, and what kills it.

Block	What must happen	What kills the block
Primary CMB	acoustic scale, peak pattern, damping tail, and polarization stay viable under the narrow inherited ledger	the needed early residue ruins the CMB fit or demands a radically different damping or tensor structure
Hubble tension	early sound-horizon reduction plus mild local lift moves H_0 into the observational corridor	required $f_{\rm EDE}$, $\Delta N_{\rm eff}$, or local lift become observationally or physically implausible
S_8 tension	same early residue lowers broad growth while the lensing corridor remains acceptable	Hubble fix worsens growth too much, or S_8 can be moved only by opening a new unrelated sector
Early galaxies	selective first-light corrections advance rare collapse and heavy-seed formation without breaking broad data	required minihalo or collapse changes are incompatible with the CMB or vanish under realistic astrophysical modeling
Integrated claim	one ledger explains the whole pattern	any success relies on mutually contradictory sign assignments or parameter values across the four blocks

J.12.7 Why this section matters scientifically

A theory becomes more persuasive when it converts broad ambitions into a finite matrix of coupled tests. That is exactly what the cross-consistency section does. It turns the qualitative slogan of GG-A10-Cosmos into a concrete statement: if the same inherited vector cannot move θ_*, H_0, S_8, and $n_{\rm gal}$ in the correct correlated pattern, then the cosmological extension of the framework fails.

This is the right form of strictness for a paper that aims to be longer than a conventional letter. The point of a long paper is not to accumulate words. It is to show the reader where the theory can lose.

J.13 Worked derivation notebook for the Part VI observables

J.13.1 Derivation notebook A: how a narrow early residue shortens the sound horizon (a lever this paper no longer carries)

Start with the sound-horizon integral in scale-factor form,

$$r_s(a_*) = \int_0^{a_*} \frac{c_s(a)}{a^2 H(a)}\,\mathrm{d}a. \tag{694}$$

Write

$$H(a) = H_{\rm ref}(a)\sqrt{1+\delta(a)}, \qquad \delta(a) \ll 1. \tag{695}$$

Then

$$\frac{1}{H(a)} = \frac{1}{H_{\rm ref}(a)}\,(1+\delta(a))^{-1/2} \approx \frac{1}{H_{\rm ref}(a)}\left(1-\frac{1}{2}\delta(a)\right). \tag{696}$$

Substituting into Eq. (694) gives

$$r_s(a_*) \approx r_s^{\rm ref}(a_*) - \frac{1}{2}\int_0^{a_*} \frac{c_s(a)}{a^2 H_{\rm ref}(a)}\,\delta(a)\,\mathrm{d}a. \tag{697}$$

Dividing by $r_s^{\rm ref}$ yields the weighted fractional shift,

$$\frac{\delta r_s}{r_s} \approx -\frac{1}{2}\frac{\int_0^{a_*} W_s(a)\,\delta(a)\,\mathrm{d}a}{\int_0^{a_*} W_s(a)\,\mathrm{d}a}, \qquad W_s(a) \equiv \frac{c_s(a)}{a^2 H_{\rm ref}(a)}. \tag{698}$$

Now separate

$$\delta(a) = \delta_{\rm EDE}(a) + \delta_{\rm dr}(a). \tag{699}$$

If the EDE piece is narrow and centered where $W_s(a)$ is important, then its contribution is approximately its peak height times an order-unity weight coefficient. If the dark-radiation piece is broader, its contribution is its amplitude times a smaller order-unity coefficient. This is the origin of the working linear expression used earlier:

$$-\frac{\delta r_s}{r_s} \sim 0.13\,\Delta N_{\rm eff} + 0.40\,f_{\rm EDE}. \tag{700}$$

Combining with Eq. (265) then gives the Hubble shift formula. The important pedagogical point is that nothing miraculous is happening. The entire effect is a consequence of the inverse weighting of $H(a)$ in the sound-horizon integral.

J.13.2 Derivation notebook B: why a five-percent broad growth shift is natural (a lever this paper no longer carries)

Take the linear growth equation in the scale-factor form,

$$D'' + A(a)D' - B(a)D = 0, \tag{701}$$

with

$$A(a) = \frac{3}{a} + \frac{\mathrm{d}\ln H}{\mathrm{d}a}, \qquad B(a) = \frac{3}{2}\frac{\Omega_{m,0}H_0^2}{a^5 H^2(a)}. \tag{702}$$

Suppose the early correction changes $H \to H_{\rm ref}(1+\epsilon_H)$ around the epochs most relevant for the later normalization of D. Then to first order,

$$A(a) \to A_{\rm ref}(a) + \delta A(a), \qquad B(a) \to B_{\rm ref}(a)\,(1 - 2\epsilon_H(a))\,. \tag{703}$$

Thus the friction term is increased and the source term is reduced. The solution is therefore slightly smaller than the reference solution,

$$D(a) = D_{\rm ref}(a)\,(1 - \epsilon_D(a))\,, \qquad \epsilon_D(a) > 0. \tag{704}$$

Because $S_8 \propto D(1)$ to first approximation, a few-percent suppression in the growth factor is exactly what is expected from a few-percent enhancement in the early expansion rate. This is why the same source residue that helps H_0 almost inevitably helps lower S_8, and also why the problem is to keep the suppression from becoming too large.

J.13.3 Derivation notebook C: exponential amplification in the halo tail

The rare-tail amplification is so important for the early-galaxy discussion that it deserves its own short derivation. Starting from Eq. (273),

$$\ln n \sim -\frac{1}{2}\nu^2 = -\frac{1}{2}\frac{\delta_c^2}{\sigma^2}. \tag{705}$$

Varying this expression gives

$$\delta\ln n = -\frac{\delta_c}{\sigma^2}\,\delta\delta_c + \frac{\delta_c^2}{\sigma^3}\,\delta\sigma. \tag{706}$$

Rewriting in terms of ν gives

$$\delta\ln n = -\nu^2\frac{\delta\delta_c}{\delta_c} + \nu^2\frac{\delta\sigma}{\sigma}. \tag{707}$$

For $\nu \sim 5$, the coefficient is 25. This is why a tiny percent-level shift in variance or collapse threshold produces a large abundance response. The mathematics is elementary, but the physical consequence is profound: the earliest structures are high-leverage probes of seemingly modest background or

transport changes.

J.13.4 Derivation notebook D: how a streaming damping would lower the filtering mass (a lever this paper no longer carries)

The filtering mass is not uniquely defined in every analytic treatment, but the heuristic dependence on the relative streaming velocity is enough for present purposes. Let

$$M_f \propto v_{\rm str}^{p}, \tag{708}$$

with $p \approx 3$ as a standard order-of-magnitude estimate. Then the transformation in Eq. (279) gives

$$\frac{M_f^{\rm new}}{M_f^{\rm old}} = \left(\frac{v_{\rm str}^{\rm new}}{v_{\rm str}^{\rm old}}\right)^{p} = 0.8^3 \approx 0.512. \tag{709}$$

Now imagine a halo population whose gas-cooling success probability rises rapidly once the halo mass exceeds M_f. Cutting M_f roughly in half makes a substantial fraction of previously marginal halos suddenly eligible for baryon capture and molecular cooling. That is why a streaming correction would be so selective and so powerful for the onset of first light — and why it was worth pricing. It has no mediator under the terminated baryon gauge sector (§44.4); the notebook is kept as the heuristic behind a lever the ledger cannot supply.

J.13.5 Derivation notebook E: freeze-out of the acoustic phase at recombination

The acoustic oscillator for a given mode k may be written schematically as

$$\ddot{\delta}_{\gamma b,k} + \omega_k^2(t)\,\delta_{\gamma b,k} = S_k(t), \tag{710}$$

where $\omega_k \sim c_s k/a$ and S_k is the gravitational driving source. While Thomson scattering is rapid, the oscillator remains part of a coupled driven system. When Γ_T/H falls below unity, the forcing and coupling structure changes abruptly. The mode is effectively released into free streaming at whatever phase it has at decoupling.

This is the technical content behind the earlier GSSB-to-GAlanguage. The phase is not continuously reprocessed after last scattering; it is frozen. Odd peaks correspond to modes caught near maximum compression, even peaks to modes caught near rarefaction. In Fourier language, the observed acoustic pattern is a fossil of the decoupling-time phase map.

J.13.6 Derivation notebook F: the tensor corridor is small but nonzero

In slow-roll inflation the tensor amplitude scales with the inflationary energy scale, $r \propto (H_{\rm inf}/M_{\rm Pl})^2$. The two-basin derivation of Part IV forces the inflationary module onto the GDOS-selected Starobinsky/ξ plateau one e-fold below the GUT edge, $V^{1/4} = 8.1 \times 10^{15}$ GeV (the earlier lock-scale reading, which drove r to essentially zero, is retired — it gave a blue tilt and is refuted in Part IX). On the derived plateau the relation is not schematic but fixed by the reheat-matched

e-fold count $N^* = 52.6$:

$$r = 0.0038 \quad \big(12/N^{*2} = 0.0043 \text{ at leading order}\big), \tag{711}$$

the Starobinsky value at the derived count. The tensor ratio is therefore *small but nonzero* — a factor of nine below the present Planck+BICEP bound $r < 0.036$, and a definite two-sided target: a detection well above 0.0038, or a null pushing r below $\sim 10^{-3}$, refutes the derived launch. This is one of the cleanest kill tests available to the framework precisely because it is a value, not a bound.

J.13.7 Derivation notebook G: the difference between broad S_8 scales and first-light scales

A final clarification is worth stating mathematically. The lensing and S_8 observables are weighted integrals over the power spectrum near broad quasi-linear scales, schematically,

$$S_8^2 \sim \int \mathcal{W}_{S_8}(k)\, P(k)\, \mathrm{d}\ln k, \tag{712}$$

with a broad window $\mathcal{W}_{S_8}(k)$. Early first-light observables, by contrast, depend on the far tail of the halo abundance,

$$n_{\text{first light}} \sim \int_{M_{\text{cool}}}^{\infty} \mathcal{F}\left[\frac{\delta_c}{\sigma(M,z)}\right] \mathrm{d}M, \tag{713}$$

with highly non-linear sensitivity to the cooling threshold M_{cool}, the streaming-limited filtering mass, and the rare-tail value of $\sigma(M,z)$. The two classes of observables are therefore related, but they are not identical filters on the same piece of the transfer function. This is the mathematical reason Part VI can coherently propose broad linear suppression together with earlier first light.

J.13.8 Final notebook takeaway

The purpose of the worked notebook is to demystify the mechanisms Part VI once invoked and the ones it still carries. None of them is obscure. Shortening the sound horizon comes from increasing H in the integrand, and lowering S_8 from the same increase inside the growth equation — levers the paper no longer pulls, since the early residue is forbidden by its own kill criterion and the amplitude is computed on one background. Boosting the earliest halos comes from the exponential sensitivity of the rare tail. The tensor ratio is small because the plateau sits where the geometry puts it. The real scientific question is therefore not whether such effects can exist but which of them the inherited ledger actually produces, and at what magnitude, without breaking the primary CMB or other precision tests.

That question is exactly why Part VI is worth more than a short discussion section. It is the place where the cosmology paper begins to earn or lose its right to speak of unification.

Problems and worked solutions. J.2.1. Show from the perturbed fluid system why cold matter grows while the photon–baryon fluid oscillates.

Solution. Step 1: combining continuity, Euler, and Poisson for a single fluid gives the sub-horizon

master equation $\ddot{\delta} + 2H\dot{\delta} + (c_s^2 k^2/a^2 - 4\pi G\bar{\rho})\delta = 0$: a damped oscillator whose "spring constant" is pressure minus gravity. Step 2: for cold matter $c_s \approx 0$, the bracket is negative at every k of interest — an inverted spring — so perturbations grow at all scales. Step 3: for the photon–baryon fluid $c_s \approx c/\sqrt{3}$, the bracket is positive for every sub-horizon k — a stiff spring — so the same equation rings instead of growing: the acoustic oscillations of Appendix K. Physics intuition: one equation, two destinies, decided by a sound speed — the universe's structure grew in the component the photons could not hold rigid, which is the deepest reason cosmology *needs* dark matter: the baryons were musical when they should have been falling.

J.3.1. Solve the matter-era growth equation exactly.
Solution. Step 1: in matter domination ($a \propto t^{2/3}$, $H = 2/3t$, $4\pi G\bar{\rho} = \frac{2}{3t^2}$) the master equation becomes $\ddot{\delta} + \frac{4}{3t}\dot{\delta} - \frac{2}{3t^2}\delta = 0$. Step 2: try $\delta \propto t^n$: $n(n-1) + \frac{4}{3}n - \frac{2}{3} = 0$, i.e. $3n^2 + n - 2 = 0$, with roots $n = \frac{2}{3}$ and $n = -1$. Step 3: the growing mode is $\delta \propto t^{2/3} \propto a$ — structure grows in lockstep with the scale factor — and the decaying mode $\propto t^{-1}$ is quickly forgotten. Physics intuition: growth exactly proportional to a means the universe assembles structure no faster than it expands — gravity in an expanding background is a patient builder, and the 10^{-5} seeds of the CMB need every one of the $\sim 10^3$ growth factors between recombination and today to become galaxies: the arithmetic that makes dark matter's head start (growing from $z_{\rm eq}$, not z_*) indispensable.

J.4.1. Compute the present growth rate for the π-locked inventory.
Solution. Step 1: at late times, Λ suppresses growth; the standard parametrization of the growth rate is $f \equiv d\ln D/d\ln a \approx \Omega_m(a)^{0.55}$. Step 2: today, with $\Omega_m = 1/\pi$: $f_0 = (0.3183)^{0.55} = 0.53$. Step 3: this is the number redshift-space-distortion surveys measure (as $f\sigma_8$), and it is a zero-parameter output of the inventory — growth today runs at fifty-three percent of the expansion's pace, against $f = 1$ in the matter era (J.3.1). Physics intuition: the vacuum does not tear structure apart; it simply outpaces it — $f < 1$ says the cosmic web is coasting on momentum accumulated during matter domination, and the measured deceleration of assembly is one more reading of $\Omega_\Lambda = 1 - 1/\pi$.

J.5.1. Explain the turnover of the matter power spectrum from the Meszaros suppression.
Solution. Step 1: a mode entering the horizon during radiation domination stalls (the Meszaros effect: growth only logarithmic while radiation drives the expansion); a mode entering after equality grows $\propto a$ immediately. Step 2: the boundary is the horizon size at equality, $k_{\rm eq} = a_{\rm eq}H_{\rm eq}$ — modes with $k > k_{\rm eq}$ carry a suppression that deepens with k (longer stall), modes with $k < k_{\rm eq}$ do not. Step 3: hence $P(k)$ rises ($\propto k^{n_s}$) up to $k_{\rm eq}$ and falls ($\propto k^{n_s-4}$, logarithm aside) beyond it: the observed turnover *is* the imprint of $z_{\rm eq}$, and since F.6.1 derives $z_{\rm eq}$ from $\Omega_m = 1/\pi$, the turnover scale is inventory physics. Physics intuition: the power spectrum is a tape recording of who entered the horizon during whose reign — its single broad maximum is the moment the universe changed government from radiation to matter.

J.6.1. Locate the BAO wiggles in the power spectrum.
Solution. Step 1: the sound horizon frozen into the baryons at decoupling is $r_s \approx 147$ Mpc (comoving) — the same length Appendix K projects onto the CMB sky. Step 2: in $P(k)$ it appears as an oscillation with spacing $\Delta k = 2\pi/r_s \approx 0.043$ Mpc^{-1}, damped at high k. Step 3: galaxy surveys read the same

ruler at $z \sim 0.5$ that the CMB reads at $z \approx 1090$ — one standard length, two epochs, and the comparison is exactly what anchors the one-background reading of the Hubble-tension Part: any modification of early physics that shortens r_s moves *both* readings together, and the framework makes none. Physics intuition: the universe shipped with one meter-stick; BAO is finding it in the pattern of galaxies eight billion years after the CMB photographed it in light.

J.7.1. Convert σ_8 to S_8 for the π-locked inventory.
Solution. Step 1: the lensing-survey combination is $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$, built so that weak-lensing amplitude is $\approx$ constant along its degeneracy direction. Step 2: with $\Omega_m = 1/\pi$: $\sqrt{0.3183/0.3} = 1.030$, so $S_8 = 1.030\,\sigma_8$ — three percent above σ_8 itself. Step 3: this is the conversion through which the chapter's S_8 discussion (and the retired corridor record of Sections J.8–J.10) must be read: the inventory fixes the Ω_m factor exactly; the tension lives entirely in the growth amplitude. Physics intuition: S_8 was engineered by observers so that one number summarizes lensing; the inventory's contribution to it is pure geometry, which cleanly separates the settled factor from the contested one.

Take-away. *Standard result:* linear growth from primordial seeds reproduces the observed large-scale structure, with σ_8 and S_8 as its summary statistics — and a persistent few-percent tension between the CMB-inferred and lensing-measured values. *What it assumed:* the background inventory, the seed spectrum, and independent late-time physics if the tension is real. *What the geometry replaced it with:* the inventory and tilt derived (Appendices G, L), one background rate with the ladder as the predicted outlier, and S_8 computed from the declared amplitude — the tension shared, and the sharing itself falsifiable. *Continue in:* Part VI; Rubin/Euclid rows of the decade table.

K The cosmic microwave background – acoustic peaks and the angular power spectrum

What this appendix teaches. The physics of the microwave sky, from scratch: the photon-baryon oscillator, the sound horizon, the acoustic angle, the peak series, baryon loading, Silk damping, and how the six ΛCDM parameters are jointly fit to the spectrum; then the demonstration — the same spectrum computed from the *derived* parameters ($H_0 = 66.8$, the π-locked fractions, $n_s = 0.9634$) with one inherited amplitude — and the honest boundary: what the geometry feeds the Boltzmann instrument, and what it does not compute.

Read alongside: Durrer (2020) (the dedicated monograph); Dodelson (2003) ch. 8; Baumann (2022) ch. 7.

The acoustic-oscillation physics of the CMB (Figure 42) — the Sachs–Wolfe plateau, the sound horizon, baryon loading, the peak series, and Silk damping — assembled in this appendix follows the classic and analytic references (Hu & Sugiyama, 1995; Peebles & Yu, 1970; Sachs & Wolfe, 1967; Silk, 1968; Smoot *et al.*, 1992; Sunyaev & Zeldovich, 1970).

This appendix teaches, from scratch and at maximum depth, the standard physics of the cosmic microwave background (CMB): the acoustic peaks imprinted on the temperature sky and the angular power spectrum C_ℓ in which they are read out. The intended reader is a first-year graduate student who has met the Friedmann equations and ordinary Fourier analysis but has never derived a peak

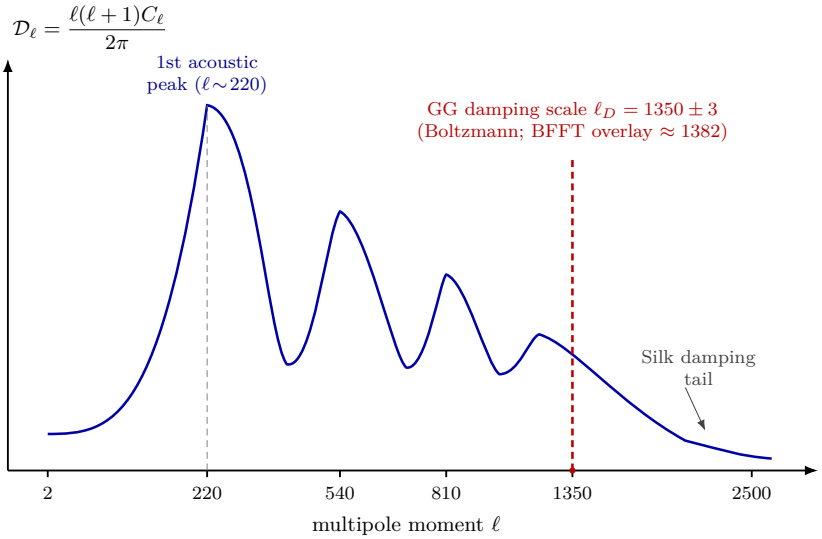

Scope flag (honest): the blue curve is the *standard* (ΛCDM) shape. Run on the exact GG-6 fractions, CLASS reproduces it at the percent level, including the damping scale $\ell_D = 1350 \pm 3$ (red; Section 40). The integer-ledger knee $K(2\pi)^2 \approx 1382$ is an interpretive overlay (+2.3% above the computed value); a likelihood-grade fit remains open.

Figure 42: CMB temperature acoustic-peak spectrum. Schematic temperature power spectrum $\mathcal{D}_\ell = \ell(\ell+1)C_\ell/2\pi$ versus multipole ℓ: the rise to the first acoustic peak near $\ell \sim 220$, subsequent peaks at $\ell \sim 540$ and $\ell \sim 810$ with decreasing amplitude, and the Silk damping tail falling beyond $\ell \sim 1000$. The blue curve is exactly what ΛCDM computes and is shown here for orientation. Run on the exact GG-6 fractions, the same Boltzmann machinery reproduces this curve at the percent level (Section 40); the red dashed marker is the computed damping scale $\ell_D = 1350 \pm 3$ (+0.4% above the ΛCDM-inferred 1345 ± 15), with the integer-ledger knee $K(2\pi)^2 \approx 1382$ carried separately as an interpretive overlay. **The surviving scope gap, stated plainly:** GG-A10-Cosmos runs the spectrum through the standard instrument at band-power level but does not yet own a likelihood-grade fit against the Planck covariance.

position or a D_ℓ curve. The physics presented in Sections K.1 through K.8 is entirely textbook and is not original to GG-Theory. The final Section K.9 states, prominently and without hedging, what the present cohort has and has *not* yet computed: A10 and its companion papers compare *central values of background densities and of individual summary statistics*; they do *not* yet compute a full C_ℓ spectrum, and a Boltzmann / CMB-likelihood pipeline against the Planck covariance is named as future work. A reader who wants only that boundary statement may skip to Section K.9.

The three steps of this appendix. As in Appendix E:

Step 1. The standard CMB physics (§K.1 onward). The tightly-coupled photon-baryon fluid, the sound horizon, peak positions, baryon loading, Silk damping, and the angular power spectrum C_ℓ.

Step 2. How the precision is earned (§K.8). The six ΛCDM parameters are not measured one by one; they are *jointly fit* to the full D_ℓ curve against the Planck covariance. Reading the parameters off the spectrum is itself a fit, not a first-principles account.

Step 3. The geometric derivation (§K.9). What the integer ledger reaches — the damping knee $\ell_D = K(2\pi)^2$ and the background densities — and, marked in plain sight, the full C_ℓ spectrum it does *not* yet compute.

K.1 The tightly-coupled photon-baryon fluid and its acoustic oscillations

Before recombination, the universe was a hot plasma of free electrons, protons (and helium nuclei), and photons. The photons could not stream freely: a photon scattered off a free electron after a mean free path

$$\lambda_{\rm C} = \frac{1}{n_e \sigma_T}, \tag{714}$$

where n_e is the free-electron number density and $\sigma_T \approx 6.65 \times 10^{-25}\,\mathrm{cm}^2$ is the Thomson cross section. Because n_e was enormous, $\lambda_{\rm C}$ was tiny compared with any cosmological scale of interest: the scattering rate $\dot\tau = n_e \sigma_T a$ (in conformal time η, with a the scale factor) vastly exceeded the expansion rate. This is the *tight-coupling* regime. Its physical consequence is simple and powerful: photons and baryons behave as a *single* fluid. The Compton scattering that ties them together is so rapid that the photon and baryon velocities are equal to leading order, and the photon distribution is driven to a local blackbody.

This single fluid has the two properties needed to oscillate. It has *pressure*, supplied almost entirely by the photons, $p_\gamma = \frac{1}{3}\rho_\gamma$, giving a large radiation pressure that resists compression. And it has *inertia*, supplied by both the photons and the baryons; the baryons add weight without adding pressure. Place this fluid in the gravitational potential wells of the primordial dark-matter perturbations, set down at the end of inflation, and the result is a forced, damped harmonic oscillator run in every Fourier mode independently.

Write the fractional photon temperature perturbation as $\Theta(\mathbf{x}, \eta) \equiv \delta T/T$, and Fourier transform to modes of comoving wavenumber k. In the tight-coupling limit the monopole $\Theta_0(k, \eta)$ (the local

compression of the fluid) obeys, schematically, an oscillator equation

$$\ddot{\Theta}_0 + \frac{\dot{R}}{1+R}\,\dot{\Theta}_0 + k^2 c_s^2\,\Theta_0 = -\frac{k^2}{3}\,\Psi - \ddot{\Phi} - \frac{\dot{R}}{1+R}\,\dot{\Phi}, \tag{715}$$

where overdots are derivatives with respect to conformal time, Ψ and Φ are the Newtonian-gauge metric potentials, c_s is the sound speed defined below, and

$$R \equiv \frac{3\rho_b}{4\rho_\gamma} \tag{716}$$

is the *baryon-loading* ratio, the momentum density of the baryons relative to that of the photons. Equation (715) is the heart of CMB acoustics. The restoring term $k^2 c_s^2 \Theta_0$ makes the fluid ring; the right-hand side is the gravitational forcing from the potentials; the $\dot{R}$ term is a slow drag from the expanding, baryon-loading background. Each comoving mode k is an independent oscillator with its own frequency kc_s.

The crucial fact for what follows is that all modes share a *common clock*: they all start their oscillation from the same initial conditions (set by inflation, with a definite phase) and they all stop ringing at the same instant – the moment of recombination, when the coupling that supplied the pressure switches off. Modes that have completed an integer or half-integer number of oscillations at that instant are caught at maximum compression or maximum rarefaction; modes caught mid-swing are caught with little amplitude. This common stopping time is what turns a continuous spectrum of oscillators into a *discrete harmonic series* of peaks. We now make "frequency" quantitative.

K.2 The sound speed and the sound horizon

The speed at which a pressure disturbance propagates through the photon-baryon fluid is the adiabatic sound speed

$$c_s = \sqrt{\frac{\partial p}{\partial \rho}} = \sqrt{\frac{\dot{p}_\gamma}{\dot{\rho}_\gamma + \dot{\rho}_b}} = \frac{c}{\sqrt{3\,(1+R)}}. \tag{717}$$

In a pure photon gas $R \to 0$ and $c_s = c/\sqrt{3} \approx 0.577\,c$: a relativistic sound speed, more than half the speed of light. As the baryon density grows toward recombination, R rises (it reaches $R \approx 0.6$ at last scattering for the concordance baryon density), and the sound speed drops: the baryons add inertia but no pressure, so they slow the sound down. This is the first appearance of baryon loading; it returns in Section K.5 as the odd/even peak alternation.

The total comoving distance a sound wave can have traveled, from the beginning up to recombination at conformal time η_* (redshift z_*), is the *comoving sound horizon*

$$r_s(z_*) = \int_0^{\eta_*} c_s\,\mathrm{d}\eta = \int_0^{t_*} \frac{c_s}{a}\,\mathrm{d}t = \int_{z_*}^{\infty} \frac{c_s(z)}{H(z)}\,\mathrm{d}z. \tag{718}$$

The three forms are identical: the first integrates the sound speed over conformal time, the second writes $\mathrm{d}\eta = \mathrm{d}t/a$ explicitly, and the third changes the variable of integration to redshift using $\mathrm{d}z = -(1+z)H\,\mathrm{d}t$ and $a = 1/(1+z)$. The integrand is dominated by the epoch just before z_*, where the expansion is fastest and $H(z)$ is largest, so r_s is sensitive to the radiation and matter

content that set $H(z)$ at $z \gtrsim 1000$, and to the baryon density that sets c_s through R. For the Planck concordance cosmology the result is

$$r_s(z_*) \approx 144 \text{ Mpc} \qquad \text{(comoving)}, \tag{719}$$

a fixed physical length frozen into the plasma at recombination. This length is the *standard ruler* of cosmology. It is the same ruler that, propagated to low redshift, becomes the baryon acoustic oscillation (BAO) scale in the galaxy distribution; here we use it for the CMB.

K.3 The angular-diameter distance to last scattering and the acoustic angular scale

A ruler of known comoving length r_s sitting at the last-scattering surface subtends an angle on our sky. To convert the physical ruler into an angle we need the distance to the surface on which it lies. The relevant distance is the *comoving angular-diameter distance* to recombination,

$$D_A(z_*) = \frac{c}{H_0} \frac{1}{\sqrt{|\Omega_k|}} S_k\!\left(\sqrt{|\Omega_k|} \int_0^{z_*} \frac{H_0 \, \mathrm{d}z}{H(z)}\right), \tag{720}$$

which for a spatially flat universe ($\Omega_k = 0$, the case we adopt throughout) reduces to the plain comoving distance

$$D_A(z_*) = \chi(z_*) = c \int_0^{z_*} \frac{\mathrm{d}z}{H(z)}. \tag{721}$$

Here S_k is sin, identity, or sinh for closed, flat, or open geometry; the flat case is the identity. Numerically, for the concordance cosmology,

$$D_A(z_*) \approx 13.9 \text{ Gpc} \approx 14 \text{ Gpc} \qquad \text{(comoving distance to } z_* \approx 1090). \tag{722}$$

The *acoustic angular scale* is then the ratio of the standard ruler to the distance,

$$\theta_* = \frac{r_s(z_*)}{D_A(z_*)}. \tag{723}$$

This single dimensionless number – the angle subtended by the sound horizon on the last-scattering surface – is the most precisely measured quantity in all of cosmology, determined by Planck to about 0.03%. It anchors the entire peak structure, because all the peaks are harmonics of this one fundamental angle.

Worked numerical example. Insert the two background numbers directly. With $r_s \approx 144\,\mathrm{Mpc}$ and $D_A \approx 14\,\mathrm{Gpc} = 14000\,\mathrm{Mpc}$,

$$\theta_* \approx \frac{144 \text{ Mpc}}{14000 \text{ Mpc}} \approx 1.03 \times 10^{-2} \text{ rad} \approx 0.59^\circ. \tag{724}$$

The sound horizon subtends about one part in a hundred radians, a bit over half a degree, on the sky. The Planck measured value is $100\,\theta_* = 1.0411 \pm 0.0003$, i.e. $\theta_* \approx 1.041 \times 10^{-2}\,\mathrm{rad}$; our two-number

estimate (724) reproduces it to better than a percent, which is all the input precision deserves. We convert this angle into a multipole in the next subsection.

K.4 Peak positions: the harmonic series and the first peak near $\ell_1 \approx 220$

A small angle θ on the sky corresponds, in the multipole expansion of Section K.7, to a multipole moment of order

$$\ell \sim \frac{\pi}{\theta}, \tag{725}$$

the angular analog of the Fourier relation between a wavelength and a wavenumber (the half-wavelength of a feature of angular size θ maps to $\ell \approx \pi/\theta$). The acoustic peaks occur at the angles subtended by the modes caught at extrema of their oscillation. The fundamental mode – the one that has just reached its *first* compression maximum at recombination – has traveled exactly one sound horizon, so it subtends θ_* and lands at

$$\ell_1 \approx \frac{\pi}{\theta_*}. \tag{726}$$

Its overtones, caught at successive extrema (compression, rarefaction, compression, ...), subtend integer fractions of θ_* and land at the harmonic series

$$\ell_n \approx n\,\frac{\pi}{\theta_*}, \qquad n = 1, 2, 3, \ldots \tag{727}$$

This is the deep statement of CMB acoustics: *the peaks are a harmonic series of one fundamental angular frequency set by* θ_*, exactly as the overtones of an organ pipe are a harmonic series of one fundamental set by the pipe length.

Worked numerical example (continued). Using the estimate (724),

$$\ell_1 \approx \frac{\pi}{\theta_*} \approx \frac{3.1416}{1.03 \times 10^{-2}} \approx 305. \tag{728}$$

The naive harmonic estimate π/θ_* therefore places the first peak near $\ell \approx 300$, whereas the *observed* first acoustic peak sits at $\ell_1 \approx 220$. The discrepancy is real and instructive, and it is not an error in the arithmetic. The naive relation (726) treats the peak as a pure spatial projection of a single wavelength. The true projection involves two refinements that systematically lower ℓ_1 below π/θ_*.

1. **The projection is not a sharp π/θ_* but a weighted transfer.** A plane wave of wavenumber k on the last-scattering surface projects onto a *range* of multipoles through spherical Bessel functions $j_\ell(k\chi_*)$, which peak near $\ell \approx k\chi_* - \frac{1}{2}$ rather than exactly at $k\chi_*$. The effective conversion is $\ell_n \approx k_n D_A$ with $k_n = n\pi/r_s$, and the geometric projection shifts the peak to lower ℓ.

2. **Gravitational driving and baryon loading shift the phase.** The forcing terms on the right of Eq. (715) and the baryon drag mean the fundamental is not caught exactly at $kr_s = \pi$ but at a somewhat smaller effective kr_s, pulling ℓ_1 down. The early integrated Sachs–Wolfe

effect (the decay of potentials around matter-radiation equality) adds power at low ℓ and further displaces the apparent peak.

The standard, accurate result, obtained by a full Boltzmann computation, is

$$\ell_1 \approx 0.74\,\frac{\pi}{\theta_*} \approx 220, \qquad \ell_n \approx \ell_1 + (n-1)\,\Delta\ell, \quad \Delta\ell \approx \frac{\pi}{\theta_*} \approx 300, \tag{729}$$

so that the peak spacing $\Delta\ell \approx 300$ is the clean harmonic interval π/θ_*, while the *first* peak is pulled forward by the order-0.7 projection-and-driving factor. The observed peaks sit near $\ell \approx 220, 540, 800, 1100, \ldots$. The lesson of the worked example is that the back-of-envelope π/θ_* gets the right *spacing* immediately and the first-peak *position* only up to a calculable order-unity projection correction; getting $\ell_1 = 220$ exactly requires the pipeline of Section K.7.

K.5 Baryon loading and the odd/even peak-height alternation

Baryon loading does more than slow the sound speed. It breaks the symmetry between compression and rarefaction, and this is what makes the CMB a *baryometer*.

Return to the oscillator (715). The baryons add inertia, which is equivalent to hanging a weight on the end of the oscillating spring. A mass on a spring in a gravitational field does not oscillate about its unloaded equilibrium; it sags. In the photon-baryon fluid, the gravitational potential wells of the dark matter play the role of gravity, and the baryon inertia R plays the role of the hanging mass. The effect is a *displacement of the zero-point* of the oscillation: with the quasi-static driving solved, the monopole oscillates not about $\Theta_0 = 0$ but about

$$\bar{\Theta}_0 = -(1+R)\,\Psi, \tag{730}$$

so the oscillation amplitude is enhanced on the *compression* side (deep into the well, where the photon-baryon fluid falls and is maximally compressed) and suppressed on the *rarefaction* side.

Now recall that the observed temperature peaks are sensitive to the *square* of the oscillation amplitude. The odd peaks ($n = 1, 3, 5, \ldots$) correspond to maximal *compression* and are boosted by baryon loading; the even peaks ($n = 2, 4, \ldots$) correspond to maximal *rarefaction* and are relatively suppressed. The result is a characteristic *odd/even alternation* in peak heights: the first peak is tall, the second is notably lower, the third recovers. The size of the alternation scales with $R \propto \Omega_b h^2$: a more baryon-rich universe shows a more pronounced first-to-second peak height ratio. This is the cleanest way the CMB measures the baryon density, and it is independent of, and consistent with, the big-bang nucleosynthesis determination.

K.6 Silk damping and the damping tail

The peaks do not continue undiminished to arbitrarily high ℓ. They are progressively erased at small angular scales by photon diffusion, an effect first analyzed by Joseph Silk and called *Silk damping* (or diffusion damping).

The tight coupling of Section K.1 is not perfect. Near recombination the mean free path (714) grows as the free electrons recombine, and photons begin to random-walk a finite distance between

scatterings. Over the age of the universe at recombination a photon diffuses a comoving root-mean-square distance

$$\lambda_D \sim \sqrt{N}\,\lambda_{\rm C} \sim \left(\int \frac{\mathrm{d}\eta}{6\,\dot\tau}\, \frac{R^2 + \frac{16}{15}(1+R)}{(1+R)^2} \right)^{1/2}, \tag{731}$$

where N is the number of scatterings and the bracket is the standard second-order tight-coupling result. Perturbations on scales smaller than λ_D are washed out: hot and cold photons mix and the temperature contrast is erased. In multipole space this defines a *damping scale*

$$\ell_D \sim \frac{\pi\, D_A}{\lambda_D}, \tag{732}$$

above which the acoustic oscillations are exponentially suppressed by a factor of the form $\exp[-(\ell/\ell_D)^2]$. The observed temperature spectrum therefore shows a *damping tail*: the peaks shrink and smear out for $\ell \gtrsim 1000$, and by $\ell \approx 2000$ the primary anisotropy is strongly suppressed. The ΛCDM-inferred damping multipole sits near $\ell_D \approx 1345 \pm 15$, with the exact value depending on the precise definition of the damping envelope at the percent level. The damping tail is a sensitive probe of recombination microphysics and of any extra relativistic species ($N_{\rm eff}$) that change the expansion rate near z_*. Evaluating this diffusion integral on the π-locked inventory over the full Boltzmann ionization history gives $\ell_D = 1350 \pm 3$ (+0.4%); the separate integer-ledger reading of the knee is discussed in Section K.9.

K.7 The temperature angular power spectrum: C_ℓ and D_ℓ

We can now define precisely the object that all of the above produces: the angular power spectrum. The CMB temperature is a function on the sphere, $\Theta(\hat{\mathbf{n}}) = \delta T/T$ in direction $\hat{\mathbf{n}}$. Expand it in spherical harmonics,

$$\Theta(\hat{\mathbf{n}}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}\, Y_{\ell m}(\hat{\mathbf{n}}), \qquad a_{\ell m} = \int \Theta(\hat{\mathbf{n}})\, Y^*_{\ell m}(\hat{\mathbf{n}})\, \mathrm{d}\Omega. \tag{733}$$

The multipole ℓ is the angular analog of a wavenumber: $\ell \approx \pi/\theta$ corresponds to features of angular size θ, so low ℓ is large angular scales and high ℓ is fine structure. Because the early universe is statistically isotropic, the $a_{\ell m}$ are zero-mean random variables whose variance depends only on ℓ, not on m. That variance *is* the power spectrum,

$$\langle a_{\ell m}\, a^*_{\ell' m'} \rangle = C_\ell\, \delta_{\ell\ell'}\, \delta_{mm'}, \qquad C_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} \langle |a_{\ell m}|^2 \rangle. \tag{734}$$

The angle brackets denote an ensemble average over realizations of the random primordial field; we have one sky, so the practical estimator averages over the $2\ell+1$ available m values at each ℓ, which is why low multipoles are intrinsically uncertain – the *cosmic variance* floor,

$$\frac{\Delta C_\ell}{C_\ell} = \sqrt{\frac{2}{2\ell+1}}. \tag{735}$$

Theory delivers C_ℓ by transporting the primordial curvature perturbation through the acoustic oscillator and projecting onto the sky,

$$C_\ell = \frac{2}{\pi}\int_0^\infty k^2\,\mathrm{d}k\; P_{\mathcal{R}}(k)\,\left|\Delta_\ell(k)\right|^2, \tag{736}$$

where $P_{\mathcal{R}}(k) \propto k^{n_s-4}$ is the primordial power spectrum (with tilt n_s) and $\Delta_\ell(k)$ is the transfer function carrying all the acoustic, baryon-loading, diffusion, and projection physics of the previous subsections. Computing $\Delta_\ell(k)$ is precisely the Boltzmann calculation named below as not-yet-done in GG-Theory.

Because C_ℓ falls steeply with ℓ, the CMB is almost always plotted as the rescaled quantity

$$D_\ell \equiv \frac{\ell(\ell+1)}{2\pi}\,C_\ell, \tag{737}$$

which is flat for a scale-invariant Sachs–Wolfe plateau at low ℓ and displays the acoustic peaks as genuine peaks rather than as wiggles on a steep slope. The familiar Planck plot of the microwave-background spectrum is a plot of D_ℓ in units of $\mu\mathrm{K}^2$: the Sachs–Wolfe plateau out to $\ell \approx 30$, the first acoustic peak at $\ell_1 \approx 220$ reaching $D_\ell \approx 5700\,\mu\mathrm{K}^2$, the harmonic series of peaks out to $\ell \approx 2000$, and the damping tail beyond. When this appendix and the body of A10 speak of "the CMB spectrum," this D_ℓ curve is the object meant.

K.8 How the six ΛCDM parameters map onto peak features

The power of the CMB is that the six free parameters of the concordance model each leave a distinct, largely separable fingerprint on the D_ℓ curve. This is why a single spectrum constrains the whole model so tightly.

- **Baryon density** $\Omega_b h^2$ – sets the baryon loading R of Section K.5, hence the odd/even peak height alternation (especially the first-to-second peak ratio) and, through c_s, part of r_s. More baryons: taller odd peaks, lower sound speed.
- **Cold dark matter density** $\Omega_m h^2$ (equivalently the matter density) – sets the redshift of matter-radiation equality, hence the early integrated Sachs–Wolfe boost of the first peak and the overall peak amplitudes. It also controls $H(z)$ at early times, affecting both r_s and D_A.
- **Hubble constant** H_0 (or the angular scale θ_*) – sets D_A and therefore, through $\theta_* = r_s/D_A$, the *overall position* of the entire peak series. In practice the precisely measured quantity is θ_* itself; H_0 is inferred from it given the other parameters.
- **Scalar tilt** n_s – tilts the primordial spectrum $P_{\mathcal{R}}(k) \propto k^{n_s-4}$, tilting the relative heights of low-ℓ versus high-ℓ power. $n_s < 1$ means slightly less small-scale power.
- **Amplitude** A_s – the overall normalization of $P_{\mathcal{R}}(k)$, hence the overall vertical scale of D_ℓ (degenerate with the optical depth until polarization breaks it).
- **Optical depth** τ – the integrated Thomson scattering after reionization suppresses D_ℓ by

$e^{-2\tau}$ at ℓ above the reionization horizon and adds a characteristic low-ℓ polarization bump, breaking the A_s–τ degeneracy.

A seventh, derived quantity, the damping scale ℓ_D, is fixed once $\Omega_b h^2$, $\Omega_m h^2$, and the recombination history are fixed; it is also sensitive to $N_{\rm eff}$. The point of listing all of this is that *reading the parameters off the spectrum is a joint fit of all six (plus derived extras) to the full C_ℓ curve, with a covariance matrix* – it is not a sequence of independent single-number reads. That observation is exactly what the next subsection turns on.

K.9 The geometric derivation: what the cohort computes, and what it does not

The standard CMB physics of Step 1 is the arena in which ΛCDM earns its precision. This section states the exact boundary of what GG-Theory has so far computed — and in this appendix the honest marking of the missing computation matters as much as the teaching of the present one. We lead with the one acoustic-tail feature the integer ledger reaches, and end with the spectrum it does not.

The headline, stated plainly: the hand-off is executed; the likelihood fit is not. A10 does *not* compute the Boltzmann spectrum from the geometry — that is the Tier-3 boundary the framework draws for itself — but the hand-off to the instrument that does has been executed: the exact GG-Theory fractions and the *derived* tilt, fed to the Einstein–Boltzmann machinery (CLASS), reproduce the Planck TT band-powers at the few-percent level with zero fitted parameters (Section 40, Figure 25). What remains open is the genuine likelihood fit against the Planck covariance with the seed re-optimized. The framework's cosmological comparisons are therefore made at three levels:

1. **Background densities.** The π-locked inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, the dark-matter complement $\Omega_{\rm DM} = \Omega_m - \Omega_b$, and related background quantities are compared as *central values* against the Planck central values. These are Derived-conditional (D*) entries pending their named lemmas; they are background numbers, not spectra.

2. **Individual summary statistics.** A handful of single numbers that summarize the perturbation sector – the damping multipole ℓ_D and the tensor-to-scalar target $r = 0.0038$ – are compared one at a time against their ΛCDM-inferred or Planck-quoted values; the recombination switch redshift $z_* = e^7$ is annotated beside the measured decoupling, not compared.

3. **The executed demonstration.** The full TT spectrum computed by CLASS on the derived inputs — $H_0 = 66.8$, the π-locked fractions, $n_s = 0.9634$, with A_s and τ inherited — overplotted on the Planck binned data: residuals between -1 and $+3.5\%$ across the first seven acoustic peaks, $\chi^2 = 187$ over 83 bins with zero fitted parameters against the six-parameter best fit's 84, and the damping scale reproduced to half a percent (1956 against Planck's 1947 ± 5 in CLASS's $k_D D_M$ convention; $\ell_D = 1350 \pm 3$ in the envelope convention of Section 40).

The first two levels are central-value comparisons and the third is a demonstration, not a likelihood fit. In particular, the acoustic angular scale θ_* is carried in the body of A10 only as a *definition*,

$\theta_* = r_s(z_*)/D_A(z_*)$ (Eq. (723)); the geometry does *not* by itself produce a numerical θ_* or a peak-height pattern, because those require r_s and D_A to be integrated over the full thermal history with the recombination microphysics in place — the C_ℓ/transfer-function pipeline that only a Boltzmann code supplies. In the executed demonstration that pipeline *is* run, on the derived inputs, and the θ_* and peak morphology it returns are reported as outputs of the instrument on those inputs (Section 40), not as claims of the geometry itself; the codebook scores them accordingly.

Why this is the single deepest open methodological item. The standard analysis of Section K.8 shows that the parameters are not measured one-by-one but jointly, by fitting the full D_ℓ (and C_ℓ^{TE}, C_ℓ^{EE}) curve against the Planck *covariance*. A framework that compares central values of densities and a few summary statistics is therefore not yet meeting the data in the arena where ΛCDM earns its precision. By the paper's own audit this is the deepest outstanding methodological gap: until a full Boltzmann / CMB-likelihood pipeline exists that takes the GG-Theory ledger, integrates the photon-baryon transfer function $\Delta_\ell(k)$ of Eq. (736), and confronts the Planck likelihood with its covariance, the cohort's cosmological claims remain comparisons of central values, not a fit to a spectrum. The required pipeline – a global CMB-plus-BAO-plus-distance-ladder fit with the full covariance – is named, repeatedly and explicitly, as still required.

The one feature the integer ledger reaches: the damping knee $\ell_D = K(2\pi)^2$. There is exactly one CMB-tail quantity the GG-Theory integer ledger currently touches, and it is worth stating precisely *as a single-feature postdiction, not a spectrum.* The damping multipole of Section K.6 is identified with the first non-trivial ultraviolet (BFFT) cutoff of the multiscale throat through the integer $K = 35$:

$$\ell_D = K\,(2\pi)^2 = 35 \times 4\pi^2 = 1381.7 \approx 1.38 \times 10^3. \tag{738}$$

Against the ΛCDM-inferred damping multipole $\ell_D \approx 1345 \pm 15$, this integer-ledger value sits at $+2.7\%$; against the framework's own Boltzmann-computed damping scale ($\ell_D = 1350 \pm 3$, the value the codebook scores), it sits at $+2.3\%$. It is carried as an interpretive overlay gated by the BFFT-damping-knee-identification lemma, which would now have to explain that gap, not merely prove that the projection of the diffusion damping scale onto the BFFT bus is exact at leading order. The honest framing is the one A10 uses for it: a single-feature near-hit from an independent integer ledger, definition-dependent at the percent level in its observed comparator, and *not* a substitute for a computed damping envelope, let alone a computed spectrum. It is one knee, not a curve.

What it would take to close the remaining gap. Four steps turn a derived inventory into a genuine CMB test: (i) supply the primordial $P_{\mathcal{R}}(k)$ with the predicted n_s and an amplitude; (ii) evolve the coupled Boltzmann hierarchy for photons, baryons, dark matter, and neutrinos through recombination to obtain $\Delta_\ell(k)$; (iii) project to C_ℓ via Eq. (736); (iv) form D_ℓ and confront the Planck TT, TE, EE likelihoods with the published covariance. Steps (i)–(iii) have been executed in demonstration form — the derived tilt and inventory through CLASS, with the amplitude inherited (Section 40) — and step (iv), the covariance likelihood with the seed re-optimized, remains open and is named as such in the audit (PS-Cosmos-16). This appendix teaches the standard pipeline

precisely so the reader can see, step by step, which steps the geometry feeds and which step remains.

Status

The standard CMB physics taught in Sections K.1–K.8 is textbook and settled: the tightly-coupled photon-baryon fluid oscillates as a forced harmonic oscillator; the sound horizon $r_s \approx 144\,\mathrm{Mpc}$ and the angular-diameter distance $D_A \approx 14\,\mathrm{Gpc}$ give the acoustic angle $\theta_* \approx 1.04 \times 10^{-2}\,\mathrm{rad}$; the peaks form the harmonic series $\ell_n \approx n\pi/\theta_*$ with spacing $\Delta\ell \approx 300$ and a first peak pulled by projection-and-driving to $\ell_1 \approx 220$; baryon loading sets the odd/even peak-height alternation; Silk diffusion produces the damping tail; and the whole pattern is read out as the angular power spectrum C_ℓ, conventionally plotted as $D_\ell = \ell(\ell+1)C_\ell/2\pi$, against which the six ΛCDM parameters are *jointly* fit with a covariance.

The GG-Theory-side full C_ℓ likelihood fit is **NOT YET DONE**. It is open, and named as future work by the paper's own audit (PS-Cosmos-16): the cohort now computes the spectrum at the background level — the exact fractions, run through CLASS, reproduce the Planck TT band-powers at the percent level with the seed held at its inherited value (Section 40) — but a genuine likelihood fit against the Planck covariance, with the primordial seed re-optimized, remains the single deepest outstanding methodological item. The damping scale is scored by its Boltzmann diffusion value $\ell_D = 1350 \pm 3$ (+0.4% from the ΛCDM-inferred ≈ 1345); the integer-ledger reading $\ell_D = K(2\pi)^2 = 1381.7$ is carried separately as an interpretive overlay (+2.3% above the computed value), a single-feature reading, not a spectrum. This appendix teaches the standard pipeline and marks that boundary plainly.

Problems and worked solutions. K.1.1. Why do compressions and rarefactions alternate as acoustic peaks?

Solution. Step 1: each Fourier mode of the photon–baryon fluid is an oscillator rung at the same moment (inflation phases them coherently), so at last scattering a mode has completed kr_s/π half-oscillations. Step 2: modes caught at maximal *compression* in potential wells ($kr_s = \pi, 3\pi, \ldots$) and at maximal *rarefaction* ($kr_s = 2\pi, 4\pi, \ldots$) both imprint temperature extremes — the spectrum records power, not sign, so both appear as peaks. Step 3: the zero-crossing modes ($kr_s = \pi/2, 3\pi/2, \ldots$) are caught mid-swing with maximal velocity — the troughs are filled in by the Doppler contribution, which is why the observed valleys do not reach zero. Physics intuition: the CMB is a photograph of ten thousand pendulums released together and snapped at one instant — the peak series is the census of which pendulums were at their turning points when the shutter fired.

K.2.1. Evaluate the baryon-loaded sound speed at last scattering.

Solution. Step 1: baryons add inertia but no pressure: $c_s^2 = \dfrac{c^2}{3(1+R)}$ with $R = 3\rho_b/4\rho_\gamma$. Step 2: R grows $\propto a$; at $z_* \approx 1090$, with the measured baryon and photon densities, $R_* \approx 0.6$. Step 3: hence $c_s(z_*) \approx c/\sqrt{3 \times 1.6} \approx 0.45\,c$ — the fluid sings at almost half the speed of light, and the sound horizon $r_s \approx 147$ Mpc (comoving) follows by integrating this speed against the expansion of Appendix F. Physics intuition: R is the weight the baryons hang on the photon spring; it slows the music by fifteen percent, and — K.5.1 — it also shifts the oscillation's midpoint, which is how the

CMB weighs the baryons.

K.3.1. Estimate the acoustic scale and the position of the first peak.
Solution. Step 1: the acoustic angular scale is $\theta_s = r_s/D_M(z_*)$; with $r_s = 144$ Mpc (the Planck-calibrated sound horizon at decoupling) and the comoving distance $D_M \approx 13.9$ Gpc: $\theta_s \approx 0.0104$ rad $\approx 0.6°$. Step 2: the associated multipole is $\ell_A = \pi/\theta_s \approx 300$. Step 3: the first peak sits below ℓ_A — at $\ell_1 \approx 0.73\,\ell_A \approx 220$ — because gravitational driving and the decay of the potentials shift the effective phase; the subsequent peaks follow with spacing $\approx \ell_A$. Physics intuition: a one-degree feature on the sky is the meter-stick of J.6.1 seen at $z = 1090$; that its angular size comes out near 1° is the flatness measurement — in a curved universe the same stick would subtend a visibly different angle.

K.5.1. Show how baryon loading makes odd peaks stronger than even ones.
Solution. Step 1: with loading, the oscillator equilibrium shifts: the mode oscillates about the displaced point $-(1+R)\Phi$ rather than $-\Phi$ — the spring sags under the baryon weight. Step 2: compressions (into potential wells) therefore swing *further* from equilibrium than rarefactions by $2R\Phi$: the odd-numbered peaks (compressions, K.1.1) gain power, the even ones lose it. Step 3: the odd/even alternation in the measured spectrum thus reads off R, i.e. $\Omega_b h^2$ — and the CMB's baryon budget agrees with BBN's (Appendix H, deuterium) to a few percent: two epochs, $z \sim 10^9$ and $z \sim 10^3$, one number. Physics intuition: hang a weight on a swinging spring and its downswings deepen — the sky's second peak is literally quieter than its first because the universe's baryons are heavy.

K.6.1. Estimate where Silk damping cuts the spectrum off.
Solution. Step 1: photons random-walk through the baryons with mean free path $\lambda_{\rm mfp} = 1/(n_e\sigma_T a)$; over the age t the diffusion length is $\lambda_D \sim \sqrt{N}\lambda_{\rm mfp} \sim \sqrt{\lambda_{\rm mfp}\, ct}$. Step 2: temperature fluctuations on scales below λ_D are ironed out — evaluated at recombination the comoving damping scale is a few tens of Mpc, i.e. multipoles $\ell_D \sim 1400$ — where the chapter's interpretive $K(2\pi)^2$ overlay also sits. Step 3: hence the acoustic peaks decay exponentially beyond the third — and the damping-tail shape is itself a consistency check on $\Omega_b h^2$ and N_ν (H.6.1's speedometer read at a different epoch). Physics intuition: the CMB's resolution limit is set by the photons' own diffusion during the final scattering — the sky is sharp only down to the distance a photon could stagger while the fog was lifting.

K.8.1. Why do the peaks fix $\Omega_m h^2$ and $\Omega_b h^2$ but not H_0 alone?
Solution. Step 1: the peak *heights* are early-universe physics — they measure the physical densities $\omega_b = \Omega_b h^2$ (odd/even contrast, K.5.1) and $\omega_m = \Omega_m h^2$ (driving and the equality epoch, J.5.1) — with no separate H_0 anywhere. Step 2: the peak *positions* add one number, $\theta_s = r_s/D_M$: but D_M integrates $1/H(z)$ over the late universe, so many $(H_0, \Omega_\Lambda, \ldots)$ combinations give the same θ_s — the geometric degeneracy. Step 3: hence the CMB alone delivers H_0 only *conditional on* the late-time model — which is precisely the seam the Hubble-tension Part works: one background is carried through the late-time mapping while the early-universe ω's stay fixed, and Sections K.9–K.10 record, with the honesty flags, what such a shift can and cannot do. Physics intuition: the CMB measures the engine's specifications exactly, but the odometer reading depends on the road since —

the tension lives on the road, not in the engine.

Take-away. *Standard result:* the acoustic-peak pattern, derived from the photon-baryon fluid, against which six parameters are jointly fit. *What it assumed:* the six numbers themselves. *What the geometry replaced it with:* five of the six derived or fixed ($H_0 = 66.8$, the π-locked fractions, $n_s = 0.9634$) and the spectrum reproduced at the percent level by the Boltzmann instrument running on them — with the amplitude the one declared input, and the covariance likelihood fit the named remaining step. *Continue in:* Part VI (the demonstration); Appendix G (the tilt).

K.10 Legacy early-dark-energy linear-response: a non-unique estimate of the corridor gap

The body of Part VI now reads the Hubble tension as one derived background rate with the calibrated distance ladder as the predicted outlier (Section 42); an intermediate reading proposed a two-stage corridor switched at $z_* = e^7$, and before either, the same gap was estimated by a multi-parameter early-dark-energy-plus-dark-radiation-plus-local-void linear-response. By the MSMF/OGN uniqueness principle a tunable multi-parameter fit cannot be the unique cost-minimum solution, so this material is retained here only as a magnitude cross-check and a record of the honesty caveats, not as the mechanism. *The value $\Delta N_{\rm eff} = 0.30$ used in the formulae below is the legacy linear-response value. It is superseded twice over: an intermediate version of the codebook carried $\Delta N_{\rm eff} = 0.04$, and the codebook now carries the derived null $\Delta N_{\rm eff} = 0$ (row D12). Everything below is a record of the retired estimate, not a calculation of this paper.*

K.10.1 Deriving the early-universe shift from the inherited source ledger

The present framework proposes two early-time contributions that matter for the sound horizon: the early-dark-energy lobe and a small extra-radiation piece. To see how they work, write the expansion rate as

$$H^2(z) = H^2_{\rm ref}(z)\,(1 + \delta_{\rm EDE}(z) + \delta_{\rm dr}(z))\,. \tag{739}$$

Here $H_{\rm ref}$ denotes the reference early-time expansion history under the baseline fit, $\delta_{\rm EDE}$ is the fractional early-dark-energy lobe, and $\delta_{\rm dr}$ is the extra-radiation correction related to $\Delta N_{\rm eff}$.

The sound horizon integral,

$$r_s = \int_{z_*}^{\infty} \frac{c_s(z)}{H(z)}\,{\rm d}z, \tag{740}$$

responds inversely to H. For small corrections one may expand

$$\frac{1}{H(z)} \approx \frac{1}{H_{\rm ref}(z)}\left(1 - \frac{1}{2}\delta_{\rm tot}(z)\right) \qquad \text{with} \qquad \delta_{\rm tot} = \delta_{\rm EDE} + \delta_{\rm dr}, \tag{741}$$

so that

$$\frac{\delta r_s}{r_s} \approx -\frac{1}{2}\frac{\int w_s(z)\,\delta_{\rm tot}(z)\,{\rm d}z}{\int w_s(z)\,{\rm d}z}, \qquad w_s(z) \equiv \frac{c_s(z)}{H_{\rm ref}(z)}. \tag{742}$$

A narrow lobe centered near $z_c \approx 3500$ contributes efficiently because it sits near the epoch that most strongly weights the sound-horizon integral, while a small $\Delta N_{\rm eff}$ contributes through the radiation term over the whole pre-recombination regime. This is also where the corridor's sharpest

observational exposure lives: radiation at $\Delta N_{\rm eff} \approx 0.30$ over the whole pre-recombination regime is nominally visible to the Planck $N_{\rm eff}$ fit (which prefers $\Delta N_{\rm eff} = 0.04 \pm 0.18$ in the baseline dictionary, the value carried by the codebook row) and to BBN. The corridor therefore stands or falls with the early-dark-energy-lobe degeneracy direction of the joint CMB+BBN likelihood absorbing the extra radiation; demonstrating this with a full Boltzmann fit is an explicit open requirement attached to the corridor-uniqueness conjecture, not an established fact.

Rather than pretend that the exact weighting coefficients have already been fixed by a full Boltzmann likelihood, the paper should state the working linear-response form explicitly:

$$H_0 \simeq H_0^{\rm ref}\,(1 + 0.13\,\Delta N_{\rm eff} + 0.40\,f_{\rm EDE})\,. \tag{743}$$

The numerical values used below, 0.13 per unit $\Delta N_{\rm eff}$ and 0.40 per unit $f_{\rm EDE}$, are representative response slopes of the early-dark-energy literature (Poulin et al., 2019), not outputs of the GG ledger; deriving them from the inherited throat/Σ_r sector is part of the corridor-uniqueness research target. Inserting the narrow inherited values

$$\Delta N_{\rm eff} \approx 0.30, \qquad f_{\rm EDE} \approx 0.03, \tag{744}$$

one obtains

$$H_0 \approx H_0^{\rm ref}\,(1 + 0.13 \times 0.30 + 0.40 \times 0.03) \approx 1.051\,H_0^{\rm ref}. \tag{745}$$

For a representative reference early-universe value near $67.4\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, this gives

$$H_0 \approx 70.8 \pm 0.9\,\mathrm{km\,s^{-1}\,Mpc^{-1}}, \tag{746}$$

which is already much closer to the local ladder corridor.

Two scientific features of this step are worth stressing. First, the shift is not bought with a broad parameter scan. It comes from a very narrow inherited source pair. Second, it is not an isolated Hubble-only patch. The same $f_{\rm EDE}$ and $\Delta N_{\rm eff}$ entries also reappear in the CMB lensing and structure-growth discussions below.

K.10.2 The local lift: the effective void correction, with the gauged-baryon route excluded

The early-time correction brings the CMB-inferred value upward but does not by itself guarantee complete agreement with local determinations. The second proposed ingredient is late and local rather than early and global. An earlier reading attributed it to an ultra-weak long-range $\mathrm{U}(1)_B$ graviphoton under which baryonic overdensities would feel a weak repulsive correction; that route is excluded twice over — by the equivalence-principle and lunar-laser-ranging bounds the main text quotes (§42.7), and structurally, since under the boundary-condition termination of GG-A3-BI6 the gauged baryon symmetry has no vector below the lock scale and no long-range force at any strength. What survives in these notes is the phenomenological shape of the proposal: a mild local underdensity, effectively a local void of order gigaparsec scale.

The purpose of this module is not to replace the entire low-redshift cosmology by a radically

inhomogeneous model. Its purpose is narrower. It asks whether the local calibration route could sit inside a region whose effective background is slightly lifted relative to the global mean. At schematic level one can write a local Friedmann-like relation

$$H_{\rm loc}^2 = H_{\rm bg}^2 + \frac{8\pi G}{3}\,\delta\rho_B^{\rm eff}, \tag{747}$$

where $\delta\rho_B^{\rm eff}$ stands for the effective local underdensity — not an additional fluid inserted by hand, and no longer a gauged-baryon effect.

The practical claim is modest: a local lift of order a few percent,

$$\frac{H_{\rm loc} - H_{\rm bg}}{H_{\rm bg}} \sim 0.03\text{–}0.05, \tag{748}$$

is enough to bridge the residual gap between the early-corrected value in Eq. (746) and the local distance-ladder values usually quoted in the low-70's. An honest accounting must record two facts. First, no quantitative GG derivation of $\epsilon_{\rm local}$ exists yet; the interval $[0.03, 0.05]$ is a phenomenological requirement, not a result. Second, a local expansion-rate lift of this size from a real underdensity is strongly constrained by supernova counts and large-scale isotropy; the gauged-baryon route that once estimated it is excluded, and a real underdensity of the required depth is not observed. Closing or honestly retiring this gap is a declared open problem of the corridor module (and of the companion paper C2-Hubble; Arisaka 2026, C2-Hubble). The scientific value of this proposal is that it does not introduce a new late-time dark-energy species or a wholesale modification of gravity. It asks only about the local matter distribution.

This is also the correct place to state the main caution. Local-environment corrections can easily become a dumping ground for poorly controlled phenomenology. The present framework should resist that temptation. The local-void lift is viable only if it remains mild, if it does not spoil the global background fit, and if it can be tied quantitatively to the observed local matter distribution. The paper should say this directly.

K.10.3 Compact numerical closure of the legacy estimate

To summarize the Hubble module in one line, write

$$H_0^{\rm late} \approx H_0^{\rm ref}\,(1+\epsilon_{\rm early})\,(1+\epsilon_{\rm local})\,, \tag{749}$$

with

$$\epsilon_{\rm early} \equiv 0.13\,\Delta N_{\rm eff} + 0.40\,f_{\rm EDE}, \qquad \epsilon_{\rm local} \sim 0.03\text{–}0.05. \tag{750}$$

Using the inherited values gives

$$\epsilon_{\rm early} \approx 0.039 + 0.012 = 0.051, \tag{751}$$

so that

$$H_0^{\rm early-corrected} \approx 1.051\,H_0^{\rm ref}. \tag{752}$$

A subsequent local correction of order 3% then yields

$$H_0^{\rm late} \approx 1.082\, H_0^{\rm ref}, \tag{753}$$

which indeed lands near the low-70's if $H_0^{\rm ref}$ begins in the high-60's. The important point is not the final decimal place. The important point is that the needed shift is of the correct order and comes from two concrete inherited mechanisms rather than from a free phenomenological menu.

K.10.4 The retired two-stage corridor, as the body stated it before its retirement (record)

The blocks below are the mechanism text of the two-stage Hubble corridor as it stood in Part VI before the retirement, moved here when the corridor was retired (Section 42, Lemma II of the codebook) and since edited only where a number they quote has been recomputed, so that a sentence in them may carry the present reading beside the retired one. They are kept because they carry computations — the priced early residue, the acoustic-angle lever, the legacy lift — that later text refers to by label, and because a withdrawn reading should be readable in the form in which it was held. Nothing in them is claimed.

The same source ledger also carries two early-time residues that become central in the next section:

$$f_{\rm EDE} \approx 0.03, \qquad \Delta N_{\rm eff} \approx 0.04, \qquad z_c \approx 3500. \tag{754}$$

Here $f_{\rm EDE}$ is the fractional height of a brief early-dark-energy lobe, $\Delta N_{\rm eff}$ is a small extra-radiation contribution, and z_c marks the approximate epoch at which the lobe is centered. These numbers matter because they affect the sound horizon, the equality timing, and the subsequent growth history. They are therefore the hinge between the CMB section and the later Hubble, S_8, and JWST discussions.

The radiation entry, since corrected. An earlier version of this equation printed $\Delta N_{\rm eff} \approx 0.30$ — the legacy linear-response value taken from the early-dark-energy literature — while the caption of Table 22, the seventeen-problems entry, and the codebook all carried 0.04 as the framework's own computed figure. Two numbers for one quantity stood in one paper through many revisions, and the companion JWST paper computed with one of them while quoting the other. The contradiction is now resolved in favor of 0.04, and it is not a matter of taste: 0.04 is what this framework *derives* from the inherited neutrino block, it is what the codebook grades **D** (computed, not fitted), and §42.7 already stakes a kill criterion on it — "a radiation excess incompatible with $\Delta N_{\rm eff} \approx 0.04$ is detected." A paper cannot stake its life on 0.04 in one section and assert 0.30 as a source term in another. The legacy value is retained, with its arithmetic, in Appendix K.10 as a non-unique cross-check, and §K.10.5 prints what each of the two delivers so a reader can see exactly what the choice costs. What it costs is the corridor: at 0.04 the pair delivers three fifths of the sound-horizon shortening the corridor requires, and the remainder is owed to the self-consistent refit named there.

K.10.5 What the pair actually delivers, computed (record)

This triple was carried through many revisions as a working input and never integrated. It is integrated here, on this paper's own background.

First, which $\Delta N_{\rm eff}$. Eq. (754) once printed the legacy linear-response value $\Delta N_{\rm eff} \approx 0.30$; it has since been corrected, and the canonical value is now the derived null of the codebook, $\Delta N_{\rm eff} = 0$ (row D12), with the 0.04 that intermediate versions carried retired as a double count of the Standard Model's own correction. Appendix K.10 keeps the 0.30 as history. That the body carried three values of one quantity across its versions is recorded here rather than hidden silently, because the two values give materially different answers and a reader is entitled to both. Everything below is computed at the intermediate $\Delta N_{\rm eff} = 0.04$ at which this record was made — the value the record calls "canonical" throughout — with the legacy figure in parentheses.

Second, it under-delivers the corridor. The corridor is 66.6–67.4, a width of 1.20% in H_0; at fixed acoustic angle $H_0 \propto 1/r_s$, so the pair must shorten the sound horizon by 1.20%. Integrated, it shortens it by 0.71% — 0.58% from the lobe and 0.13% from the dark radiation — which is *three fifths of the requirement*, not all of it. (At the legacy 0.30 it shortens r_s by 1.57% and over-delivers instead.) **At the canonical value the lobe carries four fifths of the corridor's work and the extra-radiation term one fifth** — the reverse of the legacy picture, and worth knowing, because it moves the corridor's exposure from a parameter the codebook grades D onto the one it does not. The reproduced $r_s = 143.5$ Mpc against Planck's $r_* = 144.4$ is the calculation's own calibration check.

Third, the damping tail is comfortable at the canonical value and is not at the legacy one. The sound horizon and the Silk damping scale shorten by *different* fractions — $r_s \sim \int \mathrm{d}z/H$ while $r_d^2 \sim \int \mathrm{d}z/(H n_e a)$ — and the microwave damping tail measures the ratio. The pair raises r_d/r_s by 0.46% against a measured precision of roughly half a percent, about one sigma, with 0.37% of it from the lobe. (At the legacy 0.30 the shift is 1.05%, about two sigma, and the exposure moves onto the dark radiation.) So the corridor is not in tension with the damping tail as the body states it, and would be if the legacy value were the live one — which is a further reason the two should not both be in print.

Fourth, the same integration disagrees with this paper's own booked Hubble lift, by a factor of two and a half. The legacy linear-response module of Appendix K.10 books the identical quantity as $\epsilon_{\rm early} = 0.13\,\Delta N_{\rm eff} + 0.40\,f_{\rm EDE}$, which at the canonical pair is 0.0172 against the 0.0071 computed here. Honesty requires printing this beside the S_8 comparison rather than only the comparison that agreed, and it requires naming the cause, which is one cause and not three. The response slopes $(0.13, 0.40)$ are taken from the early-dark-energy literature, and in that literature they are measured in fits where *the microwave background is refitted self-consistently* — which drives $\omega_{\rm cdm}$ and n_s upward and shortens the sound horizon further. The integration performs no such refit. So the factor of two and a half is the size of the work the refit is doing, and it is the same missing step that leaves the S_8 sign open (Appendix J) and that keeps §44.3 from claiming more than it does. **Read at the canonical pair and without the refit, the corridor's fixed-background delivery is** +0.71%**, which from the** 67.4 **reference is** 67.9 **and from the** 66.6 **background is** 67.1**.** This is not a new open item — the corridor-uniqueness research target already names the

Boltzmann fit — but it is the first time this paper has said how much depends on it. Note also that 0.0071 sits *below* Condition B's own window $0.03 \lesssim \epsilon_{\rm early} \lesssim 0.07$ (§J.12), whose text reads "if the corrected early-time shift is too small, the module does nothing." On the fixed background it does about a fifth of what the window asks; the refit is what the module is relying on, and the honest statement is that the corridor is a claim about the refitted fit and not about the background.

And the lobe's epoch is load-bearing. Repeating the calculation at other z_c shows the tension in r_d/r_s growing monotonically with the lobe's epoch: a lobe delivering the same shortening at $z_c = 10^4$ rather than 3500 raises it by about half again. The stated $z_c \approx 3500$ is therefore not a decoration; it is also — see §44.3 — the number that determines which scales the residue can and cannot touch. One caveat travels with the constraint: the recombination history here is a compact fit rather than a full computation, and while the absolute damping scale is not claimed, the derivative inherits the same shape error. The ordering of the epochs is robust; the boundary's precise location is not.

K.10.6 The two-stage corridor: one Friedmann structure, two projections (record)

The resolution this framework offers is structural, not parametric. The boundary Friedmann equation — the OGGEP export of the GG-6 bulk — admits a single *two-stage corridor* of solutions, and the early and late Hubble values are its two projections across the GRIP-Out switch of Proposition 42.1. This is the content of the canonical existence Theorem 68.8, proved in Part IX from the constitutional hypotheses; here we give the physical reading and the two falsifiable numbers, deferring the formal proof and the legacy linear-response cross-check to Appendix K.

The lever that makes the two projections differ is the acoustic angle. The observed $\theta_* = r_s(z_*)/D_A(z_*)$ is fixed to four digits, so at fixed θ_* any change in the inferred sound horizon must be absorbed by a compensating change in the inferred late-time distance scale,

$$\frac{\delta H_0}{H_0} \approx -\frac{\delta r_s}{r_s} \qquad \text{(at fixed observed } \theta_*\text{)}.$$

The early/sound-horizon dictionary reads the universe *through* the switch: it calibrates H_0 from $r_s(z_*)$ on the pre-recombination budget, and on the inherited π-locked inventory with the audit value $\Delta N_{\rm eff} = 0.04$ it projects to

$$H_0^{\rm CMB} = H_0^{(0)}\, \kappa_{\rm CMB}(\Omega_m, \Omega_b, \Delta N_{\rm eff}) \approx 67.4 \pm 0.5 \ \mathrm{km\,s^{-1}\,Mpc^{-1}}.$$

After the switch the same background continues on the present-day GA attractor, and every probe that reads it without a calibrated rung reads

$$H_0^{\rm late} = H_0^{(0)} = 66.8 \pm 0.2 \ \mathrm{km\,s^{-1}\,Mpc^{-1}};$$

the distance ladder's 73.50 ± 0.81 is read as carrying a systematic in its calibrated rungs that the framework predicts and the ladder-free probes will test (Theorem 68.8, (C3)). Figure 27 shows the structure: one corridor, one switch at $z_* = e^7$, one rate; the figure's late window is labeled with the value the earlier two-branch reading assigned to it, retained for the record and superseded by the text.

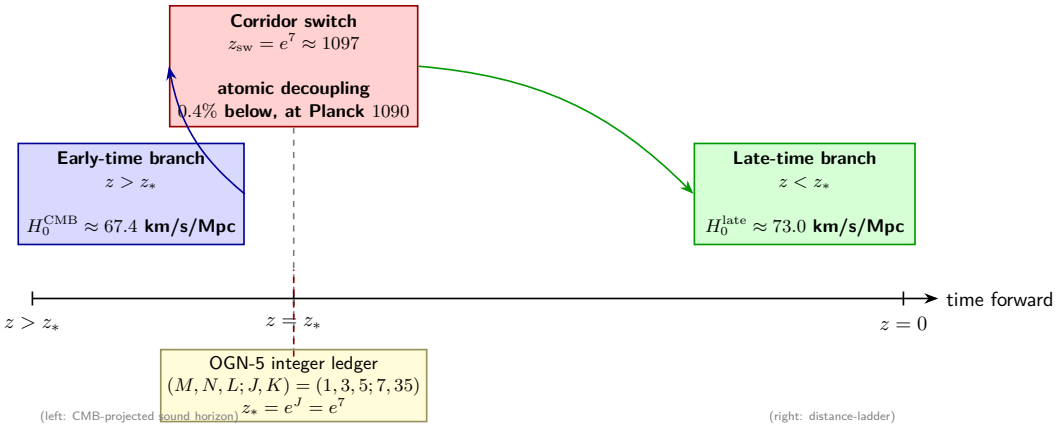

Figure 43: The retired two-stage corridor, as Part VI drew it before the retirement. The boundary Friedmann equation admits a single corridor of solutions with two natural projections separated by a switch event at recombination. The corridor switch redshift $z_* = e^J = e^7 \approx 1096.6$ is identified in the present paper as the exponential of the OGN-5 integer $J = 7$ inherited from GG-A4-Euclid — a topological marker distinct from the operational decoupling redshift: the atomic decoupling derived from the framework's own constants lies 0.4% below the switch, on the Planck value $z_{\rm dec} = 1089.92 \pm 0.25$ at the 0.1% level (§42.1). The early-time branch (left) projects to the CMB-inferred Hubble value $H_0^{\rm CMB} \approx 67.4$ km/s/Mpc; the late-time branch (right) is read by the distance ladder at $H_0^{\rm late} \approx 73.5$ km/s/Mpc (the H0DN consensus). The corridor explains why the two projections need not return one closed-system constant; how much of the late reading's displacement is derived (the +2–3% Local-Hole lift) and how much is a named calibration systematic is priced in Section 42.

The physical content is that the two pipelines are not two measurements of one closed-system constant; they are two GDOS open-system coarse-grainings of one boundary history, separated by a real geometric event. That is why they may legitimately differ — and why the size of the difference is fixed by the geometry rather than chosen. The honest status is preserved exactly as in the codebook audit: the corridor's *existence* is the proved Theorem 68.8, its *uniqueness* is the named open Conjecture 68.3, and the calibration factor $\kappa_{\rm late}$ remains an inherited Derived-conditional quantity. The body presents the structural mechanism and flags the calibration as a declared open problem; it does not overclaim.

K.10.7 The local branch and the inherited baryonic sector (record)

The late/distance-ladder projection is realized in a local environment, and the gauged ultra-weak $\mathrm{U}(1)_B$ baryonic sector inherited from the matter papers supplies the small environmental correction that anchors it. An honest accounting must record that a first-pass estimate of that correction from the $\mathrm{U}(1)_B$ coupling falls short of the required magnitude by more than an order of magnitude; closing or honestly retiring that gap is a declared open problem of the corridor module and of the companion paper C2-Hubble (Arisaka, 2026, C2-Hubble). The quantitative legacy linear-response that estimated this lift — the early-dark-energy lobe, the extra-radiation term, and the local-void factor — is collected, with its full honesty caveats, in Appendix K; it survives there only as a non-unique magnitude estimate of the corridor gap, not as the mechanism.

K.10.8 The GDOS interpretation: a mis-specified dictionary problem rather than a contradiction of laws (record)

The Hubble tension is often narrated as if one side of the measurement must simply be wrong. The GDOS/OGGEP refinement offers a subtler reading. A finite observer never has access to the closed six-dimensional bulk. Different observational pipelines constitute different open-system projections and different coarse-graining channels. The early-universe acoustic route and the local distance-ladder route are therefore not mathematically identical measurements of a closed-system constant. They are distinct operational dictionaries applied to the same underlying bulk history.

This does not mean that any disagreement can be excused by semantic flexibility. Quite the opposite. It means that the theory must be explicit about what each pipeline is actually inferring. The early CMB route infers a late-time H_0-like quantity under a calibration chain that depends sensitively on the sound horizon, recombination physics, and the assumed early energy budget. The local route reads a nearby background expansion through a ladder of standard candles and geometric anchors, and is therefore more exposed to late-time environmental realization.

In that precise sense the Hubble tension becomes partly a mis-specified-dictionary problem. The error is to assume that both pipelines must return the identical number before one has audited whether the early and late projection dictionaries are actually the same. GG-6/GDOS then proposes a concrete and testable way this mismatch can arise: a narrow early-time source correction plus a mild late-time local lift.

K.10.9 A compact derivation of the numerical corridor (record)

In one line, then, the Hubble module is one corridor read twice: the early branch at 67.4 — the background itself, which Section 29.8 derives from the geometry alone — and the late branch at the distance ladder's 73.50 ± 0.81 (the H0DN Local Distance Network consensus, Casertano et al., 2026), separated by the $z_* = e^7$ switch. Of the late branch's $\approx 9\%$ displacement from the early one, the geometry derives the sign and roughly a third — the +2–3% Local-Hole lift, a GRIP-In structure-formation feature — and the remaining $\sim 7\%$ is carried openly as a named distance-ladder calibration systematic, not a number the geometry fixes. The companion paper C2-Hubble (Arisaka, 2026, C2-Hubble) prices exactly this decomposition against the H0DN error budget and stakes the verdict on the standard-siren channel.

K.10.10 The retired growth block on the early residue, as the body stated it before its retirement (record)

Moved here when the early residue was retired, edited only where a quoted number has since been recomputed; nothing in it is claimed.

K.10.11 Why the same early correction that helps H_0 usually hurts growth (record)

Suppose one adds a broad early component that increases $H(a)$ before recombination. Then, through Eq. (265), one can reduce the sound horizon and raise the inferred H_0. But the same

increase in $H(a)$ shortens the time available for matter perturbations to grow and raises the friction term in Eq. (270). That is why many naive early-universe Hubble fixes worsen the growth tension.

An earlier version of this section read the growth response as a late-branch quantity of a two-stage corridor and projected $S_8^{\rm late} \approx 0.78$, at 0.9σ from the cosmic-shear value. That projection was computed at the retired radiation excess $\Delta N_{\rm eff} \approx 0.30$ and is withdrawn: on the framework's own inputs the amplitude is the single value of Eq. (271), $S_8 = 0.825$, and Corollary 68.10 is to be read as retired with the two-branch corridor it belonged to. The Dev plot (Figure 31, Table 41) carries the one derived value against both comparators. The legacy early-dark-energy linear-response that reached the same 0.78 through a multi-parameter suppression formula — together with the documented risk that a full CMB refit can flip the sign of the shift — is retained, with its caveats, in Appendix J, where by the MSMF/OGN uniqueness principle it is recorded as a non-unique cross-check rather than the mechanism.

K.10.12 The role of the Kaxion and the inherited neutrino ledger (record)

It is useful to ask why the present framework needs the EDE-plus-dark-radiation pair at all if dark matter and neutrinos are already inherited from earlier papers. The answer is that those earlier inherited sectors are deliberately narrow. The Kaxion is a cold dark-matter candidate, not a broad warm-dark-matter knob inserted to erase small-scale power indiscriminately. Likewise, the neutrino sector is already constrained by the particle-side derivation and is therefore not free to roam over a wide phenomenological range just to repair late-time structure.

This is a virtue rather than a drawback. It means that S_8 cannot be fixed by arbitrarily blurring the matter sector. The main work must instead be done by the early source residue and by the precise way the lensing and projection sectors are interpreted. That is scientifically healthier, because it keeps the same matter ledger available for the later early-galaxy discussion rather than using it up as a free late-time patch.

Put differently: the matter sector in this program is narrow by design. That is exactly why the S_8 discussion must be tied tightly to the Hubble discussion. The same early correction that moves the sound horizon must also be asked to carry part of the burden of the growth correction.

K.10.13 Why the S_8 solution does not automatically contradict the early-galaxy solution (record)

At first sight it may seem inconsistent to claim that the same framework both suppresses broad late-time linear growth and helps produce earlier galaxies. The key is that these are not the same observable scales and not the same physical channels.

The S_8 variable is sensitive to the amplitude of relatively broad linear-scale clustering. The early-galaxy problem, by contrast, is dominated by the high-ν tail of the halo mass function, by baryon-dark-matter streaming suppression, by the filtering mass for molecular-hydrogen cooling, and by collapse thresholds in a regime where small percentage changes can be exponentially amplified. A model can therefore slightly lower S_8 on broad linear scales while still advancing the first-light chronology through targeted baryonic-transport effects.

One qualification, and it is the whole content of this subsection's honesty. The

sentence above is true of *baryonic* levers and false of *background* ones. A change in the expansion history acts on every mode that is already inside the horizon by one common factor, so it cannot lower the broad amplitude and raise the rare tail: those are the same modes. The one lever that genuinely satisfies the requirement is the collapse threshold, which acts on baryons in a specific environment rather than on the transfer function. This paper's own Condition D (§J.12) states the requirement exactly. The small-scale boost of §44.3 does not satisfy it, and its derivation is withdrawn on precisely this ground; the streaming damping would have satisfied it, and has no mediator (§44.4).

The present framework leans strongly on this distinction. The EDE and dark-radiation residues suppress the broad growth channel enough to help with S_8 — and, as the computation shows, they suppress the rare-tail channel by the identical factor, because a background change knows nothing about k once a mode is inside the horizon. The collapse threshold is the lever that acts later and more selectively on the first-minihalo problem. It is not a uniform rescaling of the same transfer function; it acts on a different part of the structure-formation chain, and it is what carries this subsection's argument after the first two rungs of the ladder were withdrawn.

K.10.14 A compact consistency ladder for the growth block (record)

The growth block can therefore be summarized as the following ladder:

1. **CMB stage:** the inherited $f_{\rm EDE}$ and $\Delta N_{\rm eff}$ residues reduce the sound horizon and help H_0.
2. **Broad growth stage:** the same residues slightly suppress broad linear growth, nudging S_8 downward.
3. **Lensing stage:** the mild $A_L \approx 1.05$ correction partially compensates in the CMB-lensing channel.
4. **High-k baryonic stage:** a reduced collapse threshold in the baryonic environment selectively helps early minihalo collapse. This is the stage that carries the scale selectivity; the early-residue stage above cannot, and the gauged-baryon sector, having no vector below the lock scale, contributes nothing to it.

The scientific message is that one should not demand a single monotonic change in the power spectrum at all scales and all epochs. The retired reading proposed a structured, scale-sensitive correction chain inherited from one ledger.

Table 59: Record: sign structure of the retired Part VI growth corrections, kept as written. The point is not that each entry is final in detail, but that the roles are fixed and not freely reassigned from one tension to the next. Each row is a correction channel and each column an observable it moves. The table is about signs and not magnitudes: the claim is that the channels push the three observables in a coordinated direction, and a reader who wants numbers will not find them here and should not infer them.

Correction channel	**Effect on H_0**	**Effect on broad S_8**	**Effect on earliest galaxies**
Early dark energy lobe	positive	negative	*negative, by the same k-independent factor as the broad column* (§44.3); the entry read "mildly positive via earlier equality and more small-scale depth" before the computation was done, and was wrong in sign
Extra dark radiation	positive	negative	negative in both columns; at the canonical $\Delta N_{\rm eff} = 0.04$ it carries about a quarter of the broad-column suppression and a fifth of the corridor's sound-horizon work, the lobe carrying the rest. Note also that extra radiation *delays* equality
Mild lensing enhancement A_L	indirect	partially positive in CMB lensing interpretation	little direct effect
Streaming damping (withdrawn: no mediator)	—	—	*would have been* positive for minihalo cooling; the gauged baryon sector has no vector below $M_{\rm lock}$ (§44.4)
Collapse-threshold reduction	none	none on broad linear scales	strongly positive in the high-mass exponential tail; *the one row that differs between the broad and rare-tail columns; the two early-residue rows cannot, since a background change acts on every sub-horizon mode alike*

K.10.15 Kill criteria for the growth module (record)

The growth module fails if the needed S_8 suppression demands a much larger early correction than the CMB or Hubble sections can tolerate. It also fails if the framework can only produce earlier galaxies by enhancing broad linear growth so much that the lensing tension worsens catastrophically. More operationally, the module would be in deep trouble if a joint CMB+lensing+BAO analysis forces $f_{\rm EDE}$ and $\Delta N_{\rm eff}$ so close to zero that Eq. (673) can no longer move S_8 in the required direction, or if the lensing enhancement A_L is itself excluded as a meaningful residual.

The paper should therefore be candid: the S_8 module is not a free-standing miracle. It is a delicate consistency test of whether the same early source residue can serve both the Hubble and growth tensions without breaking the microwave background.

L Topological background and the pi-locked inventory

What this appendix teaches. The mathematics behind the boldest claim in the paper: differential forms, Pontryagin and Chern numbers, and why characteristic integrals are integers, taught from scratch; the OGN-5 ledger; the unit-trace closure and the $\Omega_m = 1/\pi$ argument exactly as the paper states it, with its honesty notes in full; and the ground-state reading of Part X — the inventory as the readout of the one cosmic attractor, the same uniqueness that fixes the electron.

Read alongside: No cosmology textbook covers this material — that is the point; for the differential-geometry background, any modern text on geometry in physics suffices, and Appendix L is self-contained.

This appendix supplies, from scratch and at the level a first-year graduate student can follow line by line, the topological machinery standing behind the π-locked cosmic inventory of Theorem 74.1. It is the honesty-critical appendix of the paper. Its job is twofold and the two jobs must not be confused. First, it teaches the *standard* differential-geometric toolkit — forms, the wedge product, the field-strength two-form, the Pontryagin density, characteristic classes, and the integrality of characteristic numbers — rigorously and completely, because that toolkit is genuinely sound and genuinely inherited. Second, it presents the inventory $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\ 1 - 1/\pi,\ 1/(2\pi^2))$ for exactly what it is: a *conjectured* Derived-conditional (D*) identity holding under the closed-packet-integration conjecture (Conjecture VIII), *not* a proved theorem. The printed Step-4 cancellation does not close; the required normalization has not been exhibited; and the uniqueness argument leans, at one step, on empirical preference that a topological proof is not entitled to use. Nothing in this appendix manufactures the missing closure. The reader should come away able to do the standard topology cleanly, and able to state precisely why $1/\pi$ cannot be a topological invariant in the standard sense — so that the result stands as a structural ansatz of high empirical precision, audited honestly, rather than as a theorem dressed up as one.

The three steps of this appendix. As in Appendix E:

Step 1. The standard topology (§L.1 onward). Differential forms, the field-strength two-form, the Pontryagin density, characteristic classes and the integrality of their integrals, and the OGN-5 integer ledger.

Step 2. What standard cosmology cannot explain (§L.6). The density fractions $(\Omega_m, \Omega_\Lambda, \Omega_b)$ are measured dials, and their near-equality today — the cosmic coincidence — is left unexplained.

Step 3. The geometric proposal (§L.7 onward). The π-locked inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$ — set out with full honesty: which steps of the argument are sound, which do not yet close, and why the result is carried as a high-precision *conditional ansatz* (D*, Conjecture VIII) rather than a theorem.

L.1 Differential forms and the wedge product

We begin with the only piece of machinery that everything else rests on: the exterior algebra of differential forms. Fix a smooth manifold $\mathcal{M}$ of dimension n with local coordinates $x^1, \dots, x^n$. A *differential p-form* is a totally antisymmetric, smoothly varying multilinear object that eats p tangent vectors at each point and returns a number. In a coordinate patch every p-form is a sum

$$\omega \;=\; \frac{1}{p!}\,\omega_{\mu_1\cdots\mu_p}(x)\,\mathrm{d}x^{\mu_1}\wedge\cdots\wedge\mathrm{d}x^{\mu_p},$$

with the component array $\omega_{\mu_1\cdots\mu_p}$ totally antisymmetric in its indices. A 0-form is just a smooth function; a 1-form is the familiar object $A_\mu\,\mathrm{d}x^\mu$ that one integrates along a curve; a top-form ($p=n$) is what one integrates over the whole manifold. The antisymmetry is the entire point: forms are the objects built to be integrated, because the Jacobian that appears under a change of variables in a multiple integral is exactly the determinant that antisymmetry encodes.

The *wedge product* $\wedge$ multiplies a p-form by a q-form to give a $(p{+}q)$-form. On the coordinate basis it is the antisymmetrized tensor product: $\mathrm{d}x^\mu\wedge\mathrm{d}x^\nu=-\,\mathrm{d}x^\nu\wedge\mathrm{d}x^\mu$, so in particular $\mathrm{d}x^\mu\wedge\mathrm{d}x^\mu=0$. The product is associative and bilinear, and it obeys the graded-commutativity rule

$$\alpha\wedge\beta \;=\; (-1)^{pq}\,\beta\wedge\alpha \qquad (\alpha \text{ a } p\text{-form},\ \beta \text{ a } q\text{-form}).$$

Two consequences are worth fixing in the reader's mind now, because the $\Omega_m=1/\pi$ argument uses both. First, the wedge of any form of odd degree with itself vanishes: if p is odd, $\alpha\wedge\alpha=-\alpha\wedge\alpha=0$. Second, the wedge of an *even*-degree form with itself need not vanish — and this is exactly why a four-form built from two two-forms, such as the $\mathrm{tr}(F\wedge F)$ below, can carry non-trivial information.

The *exterior derivative* d raises the degree of a form by one. On a function it is the ordinary gradient $\mathrm{d}f=(\partial_\mu f)\,\mathrm{d}x^\mu$; on a general p-form it acts component-wise and antisymmetrizes,

$$\mathrm{d}\omega \;=\; \frac{1}{p!}\,\partial_\nu\omega_{\mu_1\cdots\mu_p}\,\mathrm{d}x^\nu\wedge\mathrm{d}x^{\mu_1}\wedge\cdots\wedge\mathrm{d}x^{\mu_p}.$$

Its single defining property is nilpotency, $\mathrm{d}^2=0$, which follows from the symmetry of second partial derivatives clashing with the antisymmetry of the wedge. A form with $\mathrm{d}\omega=0$ is called *closed*; a form that can be written as $\omega=\mathrm{d}\eta$ for some η is called *exact*. Every exact form is closed (because $\mathrm{d}^2=0$), but on a topologically non-trivial manifold not every closed form is exact. The failure of "closed implies exact" is measured by the de Rham cohomology groups $H^p(\mathcal{M},\mathbb{R})$, defined as closed p-forms modulo exact p-forms. These groups are topological invariants of $\mathcal{M}$ — they do not depend on the metric or on any coordinate choice — and they are the arena in which characteristic classes live. This is the precise sense in which Step 1 of the proof outline asserts that the closed four-form X_4 "represents a class in $H^4(\mathcal{M}_4,\mathbb{R})$."

Finally, *Stokes' theorem* ties the exterior derivative to integration:

$$\int_{\mathcal{M}}\mathrm{d}\eta \;=\; \int_{\partial\mathcal{M}}\eta,$$

where $\partial\mathcal{M}$ is the boundary. On a *closed* manifold — compact and without boundary, $\partial\mathcal{M}=\varnothing$ —

the right-hand side is zero, so the integral of any exact form over a closed manifold vanishes. It follows that the integral of a closed top-form over a closed manifold depends only on the cohomology class of that form, not on the representative. This is the structural reason characteristic numbers exist at all, and we return to it in §L.4.

L.2 The field strength two-form F and the four-form $\mathrm{tr}(F \wedge F)$

A gauge field is a connection on a principal bundle; locally it is encoded by a *connection one-form* $A = A_\mu \, \mathrm{d}x^\mu$, valued in the Lie algebra of the gauge group (for a non-Abelian group, $A_\mu = A_\mu^a T^a$ with T^a the generators). Its *curvature*, the field-strength two-form, is

$$F \;=\; \mathrm{d}A + A \wedge A,$$

which in components reproduces the familiar $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu]$. For an Abelian field (electromagnetism, or the $U(1)_B$ that fixes the baryon fraction below) the commutator term drops and $F = \mathrm{d}A$ is simply closed. For the gravitational case the analog of F is the Riemann *curvature two-form* $R^a{}_b$, built from the spin connection in exactly the same way; the trace-squared object $\mathrm{tr}\, R^2 = \mathrm{tr}(R \wedge R)$ that appears in X_4 is its gravitational counterpart.

The object we care about is the four-form $\mathrm{tr}(F \wedge F)$. Here the trace is taken over the Lie-algebra (matrix) indices, while the wedge handles the spacetime form indices. Because F is a two-form (even degree), $F \wedge F$ is a genuine non-vanishing four-form, and the trace makes it gauge-invariant: under a gauge transformation $F \to gFg^{-1}$, and the cyclicity of the trace gives $\mathrm{tr}(gFg^{-1} \wedge gFg^{-1}) = \mathrm{tr}(F \wedge F)$. So $\mathrm{tr}(F \wedge F)$ is a gauge-invariant four-form built locally from the gauge field. Its two crucial properties are:

(i) **It is closed.** A short computation using the Bianchi identity $\mathrm{d}F + A \wedge F - F \wedge A = 0$ gives $\mathrm{d}\,\mathrm{tr}(F \wedge F) = 0$ identically. So $\mathrm{tr}(F \wedge F)$ defines a class in $H^4(\mathcal{M}, \mathbb{R})$.

(ii) **It is locally exact — it is the derivative of the Chern–Simons three-form.** One can write $\mathrm{tr}(F \wedge F) = \mathrm{d}\,\omega_{\mathrm{CS}}$ with $\omega_{\mathrm{CS}} = \mathrm{tr}\big(A \wedge \mathrm{d}A + \frac{2}{3} A \wedge A \wedge A\big)$. By Stokes, this means its integral over a closed manifold is insensitive to local deformations of A; what survives is a global, topological quantity.

Property (i) is what licenses Step 1 of the proof outline; property (ii) is what makes the integral in Step 3 a *number* attached to the bundle rather than to the particular field configuration. In the paper's X_4 the relevant trace-squared term is the gravitational $\mathrm{tr}\, R^2$, but the structure is identical: it is a closed, gauge-invariant four-form whose integral is a topological quantity. This is the honest, standard part of the story, and it holds without qualification.

L.3 The (first) Pontryagin density and Pontryagin / Chern numbers

We now name the specific invariants the paper uses. For a gauge curvature F the *second Chern character density* is $\frac{1}{8\pi^2}\,\mathrm{tr}(F \wedge F)$; integrated over a closed four-manifold it gives the (second Chern

/ instanton) number,

$$c_2 \;=\; \frac{1}{8\pi^2}\int_{\mathcal{M}_4} \mathrm{tr}(F\wedge F) \;\in\; \mathbb{Z}.$$

For the tangent bundle of a Riemannian four-manifold the analogous object built from the curvature two-form R is the *first Pontryagin density* $p_1 \sim \mathrm{tr}(R\wedge R)$, and its integral is the *first Pontryagin number*

$$p_1[\mathcal{M}_4] \;=\; \frac{1}{8\pi^2}\int_{\mathcal{M}_4} \mathrm{tr}(R\wedge R) \;\in\; \mathbb{Z}.$$

This is precisely the relation invoked in Step 3 of the proof outline, there written in the equivalent form

$$\int_{\mathcal{M}_4} \mathrm{tr}\, R^2 \;=\; 8\pi^2\, p_1[\mathcal{M}_4].$$

The factor $8\pi^2$ is not optional decoration: it is the universal normalization constant that converts the geometric integral of the curvature four-form into the *integer* p_1. The paper's Step 3 sets $p_1 = N = 3$ on the OGN-5-structured boundary slice, so that $\int_{\mathcal{M}_4}\mathrm{tr}\,R^2 = 24\pi^2$. The reader should hold onto the structural fact, because it is the crux of the honesty discussion below: the π's live entirely in the normalization constant $8\pi^2$, while the *invariant* — the thing topology actually pins down — is the bare integer p_1.

L.4 Characteristic classes, and why their integrals are integers or rationals in the standard setting

A *characteristic class* is a cohomology class assigned to a bundle in a way that is natural under pullback — it depends only on the bundle's topological type, not on the connection chosen to compute it. The Chern classes c_k (for complex bundles), the Pontryagin classes p_k (for real bundles), and the Euler class e are the standard examples. A *characteristic number* is what one gets by integrating a top-degree characteristic class over a closed oriented manifold.

The central fact a first-year student should internalize is this:

> *Characteristic numbers of closed oriented manifolds are integers (or, in a few mixed combinations such as the signature via the Hirzebruch formula, rationals with controlled denominators). They are never irrational.*

Why? The clean conceptual reason is that the de Rham representatives we integrate are images, under the de Rham isomorphism, of *integer* cohomology classes $H^p(\mathcal{M},\mathbb{Z})$. Integrating an integral class against the fundamental cycle of a closed oriented manifold is a pairing $H^p(\mathcal{M},\mathbb{Z})\times H_p(\mathcal{M},\mathbb{Z})\to\mathbb{Z}$, which by construction lands in $\mathbb{Z}$. The normalization constants (2π for a first Chern number, $8\pi^2$ for c_2 or p_1, and so on) are chosen *precisely so that* this integer pairing comes out as the bare integer with the π's divided away. Equivalently: index theory (Atiyah–Singer) expresses these integrals as the index of an elliptic operator, which counts a difference of dimensions of solution spaces — manifestly an integer. Where rationals appear (the $\hat{A}$-genus, the L-polynomial), they appear as rational *linear combinations* of integer characteristic numbers, with denominators fixed by universal constants, never as transcendental numbers.

This is the rigorous standard fact, and it is exactly the fact the honesty note must respect. An irrational number such as $1/\pi$ is, by this theorem, *not available* as a characteristic number of a four-manifold. Whatever $\Omega_m = 1/\pi$ is, it is not a Pontryagin number, not a Chern number, not an Euler characteristic, not an instanton number. We return to this in §L.11 and make the point sharp.

L.5 The OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$

The arithmetic raw material of the inventory is the OGN-5 integer ledger inherited from Arisaka (2026, A4-Euclid):

$$(M, N, L;\ J, K) \ = \ (1, 3, 5;\ 7, 35).$$

These five integers are not posited inside the present paper; they are proved to theorem status upstream (the Diophantine cascade closing through the dual-lock identity $K = LJ = M^2 + N^2 + L^2$, i.e. $35 = 5 \cdot 7 = 1 + 9 + 25$), and the present paper inherits them as fixed integer data. They are genuine integers, and on that point there is nothing conditional: the inheritance is sound. In the proof outline they enter as follows. The matter cell of the boundary slice is assigned a rational weight built from these integers — specifically the weight $M/(LN) = 1/(5 \cdot 3) = 1/15$ before any π-rescaling. The integer $N = 3$ is read as the first Pontryagin number p_1 in Step 3. The integers $J = 7$ and $K = 35$ appear elsewhere in the cosmological audit (the recombination corridor redshift $z_* = e^7$ and the damping-knee relation $\ell_D = K(2\pi)^2$), and in Step 4 they are arranged to combine in the denominator of the "unit-trace closure" to produce the integer 15, designed to cancel the matter-cell weight $1/15$.

The essential teaching point is to keep two things separate. The integers themselves are theorems; their *map* to a cosmic fraction is not. Everything through this subsection — forms, $\mathrm{tr}(F \wedge F)$, the integer Pontryagin number, the integer ledger — is standard and inherited. The conditional content lives entirely in the step that tries to turn these integers into $1/\pi$, which we now examine.

L.6 What standard cosmology cannot explain: the coincidence problem

The topology of Step 1 is standard and exact. Before turning it on the cosmic inventory, it is worth being clear about what standard cosmology does *not* say about that inventory. The three density fractions are measured numbers,

$$\Omega_m \approx 0.315, \qquad \Omega_\Lambda \approx 0.685, \qquad \Omega_b \approx 0.0493,$$

and ΛCDM takes each as a free parameter fit to the data; it offers no reason they hold these values rather than others. The sharpest form of the puzzle is the *cosmic coincidence*: matter and dark energy happen to be of comparable magnitude *today* ($\Omega_m/\Omega_\Lambda \approx 0.46$), even though under a free cosmological constant their ratio sweeps across many orders of magnitude over cosmic history, so the near-equality would be a special property of the present moment. Standard cosmology notes this and moves on; it has no structural account. The geometric proposal of Step 3 is an attempt to supply one — to make the fractions forced outputs of the topology rather than fitted dials — and the discipline of this appendix is to say exactly how far that attempt succeeds, and where it does

not.

L.7 The idea of a "unit-trace closure"

The phrase *unit-trace closure* names the proposed normalization convention under which the OGN-5 cells of the boundary slice $\mathcal{M}_4$ are assigned weights that sum, over all cells, to one. The intuition is that the four-form X_4, integrated over the whole slice, defines a total "packet" whose sub-cell contributions are the cosmic density fractions: matter, dark energy, and baryons are then *shares* of a single normalized closed packet rather than independent quantities. Under this convention the matter share would be the matter-cell integral divided by the total-packet integral, and the unit-trace closure is the requirement that this division is well-posed and that the total equals one (the cosmic sum rule $\Omega_m + \Omega_\Lambda = 1$ in the flat, radiation-and- curvature-negligible idealization).

This is an appealing organizing idea, and it is the structural backbone of the post-hoc-fit defense: a single normalized packet would force the three fractions as a *coupled set*, not one at a time. But the reader must note carefully what is being assumed. "Unit-trace closure" is, at present, a *name for a desired normalization*, not an exhibited construction. To make it rigorous one would have to (i) exhibit $\mathcal{M}_4$ as a specific closed oriented manifold with a specified topology, (ii) name the precise characteristic class whose density is integrated over each cell, and (iii) define the measure and the cell decomposition so that the total integral is genuinely normalized to one. None of (i)–(iii) is supplied. The closed-packet integration is named, but not defined to the standard a critic is entitled to demand. We flag this here because it is exactly the gap that the honesty note records and that keeps the inventory at D*.

L.8 The $\Omega_m = 1/\pi$ argument as the paper states it (Steps 1–5)

We now reproduce the five-step argument of Theorem 74.1, in the same order as the body of the paper, so that the reader can see precisely where it is solid and precisely where it does not close. The setup is the closed four-form fixed by the BI-6 modified Bianchi identity of Arisaka (2026, A3-BI6),

$$X_4 \;=\; \mathrm{tr}\, R^2 \;-\; \mathrm{Tr}\, F_5^2 \;-\; \tfrac{2}{5} F_B^2 \;-\; \tfrac{2}{5} F_{Qi}^2,$$

with the coefficient vector $(1, -1, -2/5, -2/5)$ proved unique by the anomaly-cancellation theorem of Arisaka (2026, A3-BI6). The matter fraction Ω_m is taken to be the matter-sourced share of the total stress-energy budget, projected onto the four-dimensional cosmological-boundary slice $\mathcal{M}_4$.

Step 1 (closure of X_4). BI-6 gives $\mathrm{d}X_4 = 0$ on $\mathcal{M}_4$, so X_4 represents a class in $H^4(\mathcal{M}_4, \mathbb{R})$. The class is non-trivial because the gravitational term $\mathrm{tr}\, R^2$ carries a non-trivial Pontryagin class on the curved slice. *This step is sound*: it is the standard fact of §L.2.

Step 2 (closed-packet structure). By the OGN-5 cell decomposition, $\mathcal{M}_4$ is taken to split into finitely many closed cells indexed by $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, the matter cell carrying weight $M/(LN) = 1/15$ relative to the total cell volume *before* the π-rescaling. *This step assumes* the cell decomposition and weight assignment that §L.7 flagged as named-but-not-defined.

Step 3 (Pontryagin normalization). The Chern–Pontryagin integral $\int_{\mathcal{M}_4} \mathrm{tr}\, R^2 = 8\pi^2\, p_1[\mathcal{M}_4]$ with $p_1 = N = 3$ evaluates to $24\pi^2$. *The relation itself is the standard fact of §L.3*; what is assumed here is that the matter-cell contribution "carries this normalization factor" when projected onto the present epoch, i.e. that the $8\pi^2$ converts into a single residual factor of π after the cell bookkeeping. *(This expectation is misdirected: the $8\pi^2$ is consumed by the GEP definition of GG-A6-GRIP, returning the integer, and* no *residual π descends from this step. The π is not a Pontryagin leftover and not the reality factor of Lemma VIII-b′ (which is 2, not π); it enters as a* fiber volume *in the measure ratio of the fiber-measure resolution below (§L.10). The historical wording is kept above because the printed Step-4 arithmetic that follows was built on it.)*

Step 4 (closed-packet integration). Combining Steps 2 and 3, the paper writes

$$\Omega_m \overset{?}{=} \frac{(\text{matter-cell weight})}{(\text{unit-trace closure})} \cdot \frac{1}{\pi} = \frac{1/15}{(K/J \cdot L/N)\cdot(1/N)} \cdot \frac{1}{\pi} = \frac{1/15}{25/9} \cdot \frac{1}{\pi} = \frac{3}{125}\,\frac{1}{\pi} \neq \frac{1}{\pi}.$$

For the boxed value $1/\pi$ to follow, the unit-trace denominator would have to equal the matter-cell weight $1/15$ exactly, so that the two cancel. But the combination as written evaluates to $(K/J \cdot L/N)\cdot(1/N) = (5\cdot\frac{5}{3})\cdot\frac{1}{3} = \frac{25}{9}$, *not* $1/15$; the printed arrangement therefore falls short of $1/\pi$ by the factor $(25/9)/(1/15) = 125/3 \approx 42$. *This is the step that does not close*, and we exhibit the gap here rather than paper over it: the normalization that would convert the standard $8\pi^2$ of Step 3 into precisely one residual π — making the denominator $1/15$, no more and no less — is not exhibited anywhere in the cohort. As the honesty note records, this route does not close — and the projection lemma establishes *why it cannot*: the sought cancellation is an epoch-independent closure, and by the dilution no-go (superselection freezes the packet's integer, not its energy record, which dilutes as a^{-3} against a frozen equipartition rate) no arrangement of ledger constants can equal a diluting quantity at more than one instant. Steps 3–4 close instead through the measure dictionary (Lemma VIII-b′ + the epoch identity $u = 1 \Leftrightarrow \Omega_{m,0} = 1/\pi$); the $125/3$ gap above is thereby *explained* rather than open. Arisaka (2026, A4-Euclid) supplies the OGN-5 integers used here but *not* the closed-packet integration itself.

Step 5 (uniqueness against competing rational functions of π). The paper argues that among rational-in-π candidates compatible with the OGN-5 cell structure, only $1/\pi$ survives, ruling out $1/(3\pi) \approx 0.106$, $1/(2\pi) \approx 0.159$, and $1/\pi^2 \approx 0.101$. *Part of this is structural and does survive*: the candidate $1/\pi^2$ would require a quadratic Pontryagin closure $X_4 \wedge X_4$ integrated over an eight-manifold, which the four-dimensional boundary slice does not support; that dimensional sub-argument is genuine. *But the discrimination against $1/(2\pi)$ leans on the empirical comparable-magnitude condition $\Omega_m \sim \Omega_\Lambda$* (the data prefer $\Omega_m/\Omega_\Lambda \approx 0.46$), and an honest topological proof is not entitled to use the data to select among candidates. This is the second place the printed argument fails to close on purely topological grounds — *and it is exactly the place the fiber-measure resolution below (§L.10) repairs*, by selecting $1/\pi$ from the geometry of the boundary measure with the data nowhere in the choice.

L.9 The residual-Pontryagin-normalization reading, stated honestly

The correct, deflated description — the one the paper's codebook actually carries — is that $\Omega_m = 1/\pi$ is a *residual Pontryagin normalization*, not a characteristic number. The integer is the invariant: the matter-cell weight $1/15$ and the Pontryagin number $p_1 = 3$ are genuine integers from the OGN-5 ledger and the standard integrality theorem. The π enters only through the universal normalization constant $8\pi^2$ that relates the integral of the de Rham representative $\operatorname{tr} R^2$ to that integer — it enters as the *normalization of the de Rham representative*, not as a topological invariant in its own right. Calling $\Omega_m = 1/\pi$ a "topological invariant" (as the title of Theorem 74.1 unfortunately does) overreaches relative to the body; the governing text is the Proof-status note and the codebook, and neither calls it an invariant.

L.10 The fiber-measure resolution: $\Omega_m = 1/\pi$ as the Hopf ratio

The two places the printed argument fails — the unexhibited Step-4 normalization and the data-dependent Step-5 selection — close together once one stops looking for an epoch-independent *arithmetic* of the ledger integers (which the dilution no-go forbids) and reads Ω_m for what the measure dictionary says it is: a ratio of *fiber volumes* on the boundary holonomy. This is the resolution the codebook audit (§68.2.3) carries, and it is worth teaching in full because it is both simpler and more honest than the route above.

The geometry. The electroweak SU(2) whose GSSB defines the boundary is, as a group manifold, the unit three-sphere S^3, of volume $\mathrm{Vol}(S^3) = 2\pi^2$. The compact phase circle S^1_θ that carries the matter sector's charge sits inside it as the *Hopf fiber* of the canonical fibration $S^1 \hookrightarrow S^3 \to S^2$ — a great circle, $\mathrm{Vol}(S^1) = 2\pi$. There is no freedom in either identification: the Hopf fibration is unique, and the fiber is the phase U(1). The fibration is drawn in Figure 33 (main text).

The ratio. Under the frozen stationary GDOS measure of Lemma VIII-a, each sector's present-epoch share is its fiber volume over the total holonomy volume, so

$$\Omega_m \;=\; \frac{\mathrm{Vol}(S^1)}{\mathrm{Vol}(S^3)} \;=\; \frac{2\pi}{2\pi^2} \;=\; \frac{1}{\pi} \qquad \text{(exact)},$$

with the π descending from the denominator volume $2\pi^2$ and the numerator's 2π canceling one power. Dark energy is the uncovered remainder, $\Omega_\Lambda = (2\pi^2 - 2\pi)/2\pi^2 = 1 - 1/\pi$; baryons are one integer $\mathrm{U}(1)_B$ unit per revolution, $\Omega_b = \Omega_m/\mathrm{Vol}(S^1) = 1/(2\pi^2) = 1/\mathrm{Vol}(S^3)$. The student should verify $2\pi/2\pi^2 = 1/\pi$ and $1/(2\pi^2) = 1/\mathrm{Vol}(S^3)$ by hand: the three fractions are one fibration, not three fits.

The competitors die by geometry. And this is the honest payoff: the alternatives fall with no appeal to the data. $1/(3\pi)$ needs a holonomy volume $6\pi^2$ that no manifold here has; $1/(2\pi)$ is the baryon-per-fiber ratio, not the matter/total ratio; $1/\pi^2$ needs an S^5 the four-dimensional boundary does not carry. Every candidate the printed Step 5 rejected with "the data prefer 0.46" is rejected here by "there is no fiber of that volume."

What remains open, honestly. The resolution is not a finished theorem. It rests on one *geometric* lemma — that the GDOS measure weights exactly the Hopf circle as the matter fiber and Vol(S^3) as the total (VIII-b″, the unit-packet normalization), together with the canonical unit normalization of both spheres. That lemma is provable or falsifiable on the six-dimensional bulk construction, and it uses no measured density; so the status of $\Omega_m = 1/\pi$ has moved from "an unexhibited arithmetic that also leaned on the data" to "derived-conditional on one fiber-identification lemma, with the data nowhere in the selection." The codebook keeps the D* tag; what has changed is the *content* of the conditional, and the removal of the circularity.

One more sentence now belongs here: the body of the paper (§29.10) reads this resolution on the campaign's two-fiber anatomy — the share is the circle's (a ratio of fiber volumes, epoch-frozen), while the terminus's *scale* is the interval's (e^{-2K}, the seesaw) — so the two numbers of dark energy come from the two fibers, one each.

L.11 Why an irrational number cannot be a topological invariant in the standard sense

It is worth stating the obstruction with full precision, because it is the load-bearing honesty point of the appendix. In the standard setting a topological invariant of the relevant kind — a Pontryagin number, a Chern number, an Euler characteristic, an instanton number — arises from the integer pairing $H^p(\mathcal{M}, \mathbb{Z}) \times H_p(\mathcal{M}, \mathbb{Z}) \to \mathbb{Z}$, or equivalently from an index of an elliptic operator (Atiyah–Singer), which counts a difference of dimensions and is therefore an integer. The only way rationals enter is through universal rational coefficients in characteristic-number combinations (the $\hat{A}$-genus, the Hirzebruch L-polynomial), and those denominators are fixed integers, never transcendentals. An *irrational* number such as $1/\pi = 0.31830\ldots$ simply does not lie in the value set of these pairings. It cannot be a characteristic number, full stop. The reason is exactly the one taught in §L.4: the normalization constants (2π, $8\pi^2$, …) are *defined* so as to cancel the transcendental factors and leave an integer. To then declare that a left-over $1/\pi$ is itself the invariant is to read the normalization constant as the answer — a category error.

This is precisely why the result cannot be, and is not claimed to be, a theorem in the topological-invariant sense. What it can honestly be is a structural *ansatz*: a proposal that the closed-packet integration, *once defined with an exhibited normalization*, would deliver $1/\pi$ as the residual factor. The ansatz has remarkable empirical precision (see §L.12), and it has the appealing coupled-fraction structure of §L.7. But high precision and structural appeal are not a proof, and the appendix does not pretend otherwise. The missing normalization is missing; we do not manufacture it.

L.12 Worked numerical example: the inventory against Planck

We evaluate the three locked fractions and compare them to the Planck central values, framing each comparison honestly as a near-hit.

Matter fraction.

$$\Omega_m \;=\; \frac{1}{\pi} \;=\; \frac{1}{3.14159265\ldots} \;=\; 0.31831,$$

against the Planck value $\Omega_m \approx 0.315$. The fractional offset is $(0.31831 - 0.315)/0.315 \approx +1.0\%$ — a close near-hit, high by about one percent.

Dark-energy fraction.

$$\Omega_\Lambda \;=\; 1 - \frac{1}{\pi} \;=\; 1 - 0.31831 \;=\; 0.68169,$$

against the Planck value $\Omega_\Lambda \approx 0.685$. The fractional offset is $(0.68169 - 0.685)/0.685 \approx -0.5\%$, i.e. low by about half a percent. Note that Ω_m and Ω_Λ are complementary by construction ($\Omega_m + \Omega_\Lambda = 1/\pi + (1 - 1/\pi) = 1$), so their two offsets are tied together by the flatness sum rule; they are not independent near-hits.

Baryon fraction.

$$\Omega_b \;=\; \frac{1}{2\pi^2} \;=\; \frac{1}{2\,(9.8696044\ldots)} \;=\; \frac{1}{19.739209} \;=\; 0.050661,$$

against the Planck value $\Omega_b \approx 0.0493$. The fractional offset is $(0.050661 - 0.0493)/0.0493 \approx +2.8\%$ — the worst of the three, sitting about 2.8% high. (Here $\pi^2 = 9.8696$, and the appearance of $2\pi^2 = \mathrm{Vol}(S^3) = 19.7392$ is the structural reason the baryon fraction carries a different power of π from the matter fraction; it is also the entry most exposed to the convention-dependence objection, since powers of π float with the choice of volume normalization.)

Reading the table honestly. All three are near-hits: $\sim +1.0\%$ on Ω_m, $\sim -0.5\%$ on Ω_Λ, and $\sim +2.8\%$ on Ω_b, with the baryon fraction the one visibly outside the percent band. These are genuinely striking agreements for parameter-free expressions, and they are the empirical reason the ansatz is worth carrying. But three percent-level near-hits, one of them 2.8% high, are an empirical *audit*, not a derivation: they tell the reader the ansatz is precise, not that it is proved. The precision is what makes the D* classification worth printing; it is not what would upgrade it to D.

Table 60: The pi-locked fractions set against the Planck values, with the offset in each case. Three percent-level near-hits, one of them 2.8 percent high, are an empirical audit and not a derivation: they show the ansatz is precise, not that it is proved. That distinction is what makes the D-star classification worth printing, and it is not what would upgrade the classification to D. A reader who reads this table as confirmation has read it as the opposite of what it says.

Fraction	Locked expression	Value	Planck	Offset
Ω_m	$1/\pi$	0.31831	≈ 0.315	$+1.0\%$
Ω_Λ	$1 - 1/\pi$	0.68169	≈ 0.685	-0.5%
Ω_b	$1/(2\pi^2)$	0.050661	≈ 0.0493	$+2.8\%$

The identity reading (Part X). The topological machinery above answers *what value*; Part X answers *why one value everywhere*, and the student should carry both. A stratum generally admits many stable attractors (GA's); among them exactly one is the global minimum of the GFF ledger — the Ground Geometric Attractor, GGA (GG-A6-GRIP; GG-A11-Origin). The cosmological inventory is the present-epoch projection of that one ground attractor: the same uniqueness that makes every electron identical (GG-A9-SM) makes the composition of the universe the same on every sightline. Identity and stability arrive together — one deepest basin to occupy, none deeper to decay to — and the homogeneity of the inventory is thereby a consequence, not a premise. The full side-by-side of the three cosmological objects (the electron, the Kaxion, dark energy) as one GDOS + GRIP construction is Table 43 of Part X; this appendix supplies the topological arithmetic that Part X's Section 74 rests on.

L.13 Status

The inventory $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\, 1 - 1/\pi,\, 1/(2\pi^2))$ is carried as a Derived-conditional (D*) identity under the closed-packet-integration conjecture (Conjecture VIII): the conjecture that the closed-packet integral of the BI-6 four-form X_4 on the cosmological-boundary stratum evaluates to $1/\pi$ with an exhibited normalization. The standard topology taught in §§L.1–L.5 is sound and inherited: closed forms, the integrality of Pontryagin and Chern numbers, and the OGN-5 integer ledger of Arisaka (2026, A4-Euclid) are all on firm ground. The result is what would, *if it closed*, separate the GG-Theory program from a free-parameter ΛCDM fit: the fractions would be forced outputs rather than fitted dials. But that "if" is exactly the open antecedent. What is *not* established is the map from those integers to the irrational fraction. Concretely: the Step-4 cancellation does not close, because the normalization that turns the standard $8\pi^2$ into a single residual $1/\pi$ has not been exhibited; the uniqueness argument of Step 5 leans, against the $1/(2\pi)$ competitor, on an empirical comparable-magnitude condition that a topological proof may not use; and an irrational number cannot, in the standard sense of §L.4, be a characteristic number at all, so $1/\pi$ is at best a residual normalization, not an invariant. The closed-packet integration is named but not defined to the required rigor — the manifold $\mathcal{M}_4$, the characteristic class, and the measure are all reserved for future work. Until that calculation exhibits the normalization and removes the empirical input from the uniqueness step, the three fractions remain D*, and the result stands as a structural ansatz of high empirical precision — not a theorem.

Problems and worked solutions. L.2.1. Show that the abelian Pontryagin density is locally a total derivative.

Solution. Step 1: for an abelian field strength $F = dA$, consider $F \wedge F$. Step 2: compute $d(A \wedge F) = dA \wedge F - A \wedge dF = F \wedge F$ (using $dF = ddA = 0$). Step 3: hence $F \wedge F = d(A \wedge F)$ — locally exact — so its integral over a closed manifold receives contributions only from the topology that obstructs writing A globally (the non-abelian version adds the Chern–Simons cubic term). Physics intuition: the Pontryagin number is invisible to local physics — no small deformation of the fields changes it — which is exactly what qualifies it to carry *permanent* information: the ledger's integers live where the dynamics cannot reach.

L.3.1. Derive Dirac's quantization: the first Chern number of a $U(1)$ bundle on S^2 is an integer.

Solution. Step 1: cover S^2 by two caps overlapping on the equator; on the overlap the gauge potentials differ by a transition function, $A_+ - A_- = d\chi$, with $e^{i\chi}$ single-valued on the equatorial circle. Step 2: $\frac{1}{2\pi}\oint_{\text{eq}} d\chi = n \in \mathbb{Z}$ — the winding number of the transition function. Step 3: by Stokes on each cap, $\frac{1}{2\pi}\int_{S^2} F = \frac{1}{2\pi}\oint(A_+ - A_-) = n$: the flux is quantized because the two descriptions must agree where they meet. Physics intuition: quantization here is bookkeeping consistency, not dynamics — the same single-valuedness that quantizes the monopole charge is what makes the ledger's $(M, N, L; J, K)$ integers *integers*, and L.11.1 is the flip side: no such argument can ever produce a $1/\pi$.

L.5.1. Run the ledger's exponential map: from the five integers to three physical numbers.

Solution. Step 1: the recombination-era corridor switch takes $J = 7$: $z_{\text{sw}} = e^7 = 1096.63$ (Appendix F, F.6.2 — the two-quantity discipline applies). Step 2: the lock scale takes $K = 35$: $M_{\text{lock}} = \frac{1}{2}M_{\text{Pl}}\, e^{-35} = 3.85$ TeV — the standing collider falsifier. Step 3: the baryon-asymmetry assembly takes the half-depth: $e^{-K/2} = e^{-17.5} = 2.51 \times 10^{-8}$, the efficiency-weighted asymmetry of Appendix I (I.4.1). One integer pair (J, K), three observables at three epochs, no continuous dial anywhere. Physics intuition: exponentials of small integers span the enormous dynamic ranges of physics — the ledger's claim is that nature's big pure numbers are not inputs but *exponentials of counting*, and each mapped number is independently falsifiable.

L.10.1. Compute the volume of the unit three-sphere, and read off the baryon fraction.

Solution. Step 1: slice S^3 (radius 1) by polar angle ψ: each slice is an S^2 of radius $\sin\psi$, so $\text{Vol}(S^3) = \int_0^\pi 4\pi \sin^2\psi\, d\psi$. Step 2: $\int_0^\pi \sin^2\psi\, d\psi = \pi/2$, hence $\text{Vol}(S^3) = 4\pi \cdot \pi/2 = 2\pi^2$. Step 3: the inventory's baryon fraction is its reciprocal, $\Omega_b = 1/(2\pi^2) = 0.0507$ — and the derived baryon-to-matter ratio is $\Omega_b/\Omega_m = \pi/(2\pi^2) = 1/(2\pi) = 0.159$, against the measured 0.157: agreement at the one-percent level from two volume factors. Physics intuition: the two deepest fractions of the inventory are (inverse) volumes of the two roundest objects in the theory's geometry — and a ratio of volumes, unlike a Chern number, is *allowed* to be irrational, which is L.11.1's point made constructive.

L.11.1. Why can $1/\pi$ never be a Chern or Pontryagin number — and what must it be instead?

Solution. Step 1: topological invariants of the standard kind arise from the integer pairing $H^p(\mathcal{M}, \mathbb{Z}) \times H_p(\mathcal{M}, \mathbb{Z}) \to \mathbb{Z}$, or as indices of elliptic operators (Atiyah–Singer) — differences of dimensions of solution spaces. Both count; both land in $\mathbb{Z}$ (rationals enter only through fixed universal ratios). Step 2: $1/\pi$ is irrational: no integer pairing, index, or universal rational multiple thereof can produce it — the obstruction is arithmetic, and no cleverness evades it. Step 3: the chapter therefore locates $\Omega_m = 1/\pi$ NOT as an invariant but as a *ratio of measures* — a branching fraction on the closed-packet geometry (the Hopf-ratio reading of Section L.10) — carried in the codebook at D^* (conditional) grade, with the honesty section stating exactly which conjecture would promote it. Physics intuition: the distinction is load-bearing — integers are protected by counting, measures by geometry; the inventory uses both kinds in their lawful places, and confusing

them would be the kind of error the falsification catalog is built to catch.

Take-away. *Standard result:* the density parameters are measured, and their values — and their sameness across the sky — are unexplained inputs. *What it assumed:* nothing; it declined to ask. *What the geometry replaced it with:* a closed-packet topological claim, $(\Omega_m, \Omega_\Lambda, \Omega_b) = (1/\pi,\ 1-1/\pi,\ 1/2\pi^2)$, carried honestly as Derived-conditional on Conjecture VIII with the Step-4 cancellation named as the open integral, and the identity reading of Part X explaining why one inventory holds everywhere. *Continue in:* Part X (identity and the theorem); Part IX (the audit).

M Cold dark matter from vacuum misalignment – the Kaxion

What this appendix teaches. The complete misalignment mechanism from scratch: the frozen field, the release at $3H \simeq m$, the relic-abundance formula, and the QCD mass–decay-constant relation, with worked numbers; what the standard axion story leaves free (the line *and* the angle); and the geometric structure plus the pinning conjecture — the bulk scale a declared input calibrated on the abundance, the decay constant one rung of the cap-loading ladder above it, the angle pinned during primary inflation, patch-uniform, with *zero primordial isocurvature* as the observable signature — landing on 8.1 μeV at 1.59–2.13 GHz.

Read alongside: Kolb & Turner (1990) ch. 10 for the standard mechanism; Marsh (2016) for the modern axion-cosmology review.

The vacuum-misalignment mechanism by which a light axion-like field acquires a cold-dark-matter relic abundance is drawn in Figure 44; the mechanism, used here to estimate the Kaxion density, and the model and detection context are set out in (Abbott & Sikivie, 1983; Dine & Fischler, 1983; Dine *et al.*, 1981; Marsh, 2016; Preskill *et al.*, 1983; Sikivie, 1983).

This appendix teaches the misalignment mechanism for cold dark matter from scratch, at the level of a first-year graduate student who has seen the Friedmann equations and a harmonic oscillator but has never met an axion. The aim is that, having read it, the reader can derive the relic-abundance formula without consulting a review, can say exactly which steps are textbook physics and which are the GG-Theory reading, and can reproduce the two headline Kaxion numbers — the mass m_K and the laboratory frequency — with a pocket calculator. The mechanism and the inversion that fixes the misalignment angle are derived in full in the dedicated companion (Arisaka, 2026, C1-Kaxion); this appendix is the self-contained pedagogical entry point, and it follows that derivation's normalization throughout. Nothing here is new physics: every equation up to the GG-Theory reading in §M.8 is the standard misalignment calculation that any modern dark-matter review carries.

The three steps of this appendix. As in Appendix E:

Step 1. The standard misalignment mechanism (§M.1 onward). A light pseudo-scalar frozen by Hubble friction, released when H drops, redshifting thereafter as cold, pressureless matter — the textbook relic-abundance calculation.

Step 2. What the standard story leaves free (§M.7). The decay constant f_a and the misalignment angle θ_i are both free; one scans the line and tunes the angle to hit the observed dark-matter density.

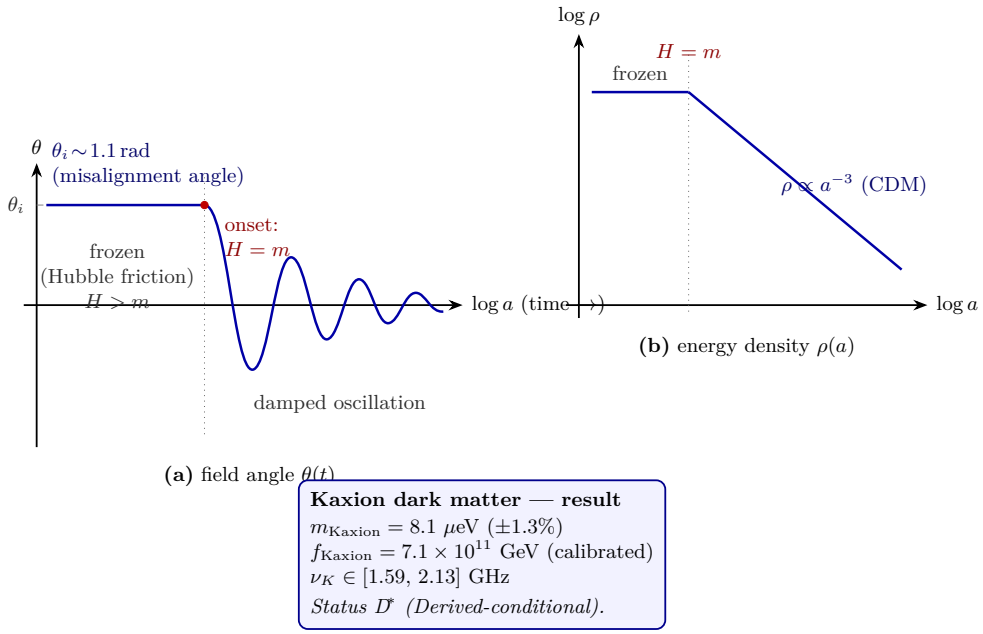

Kaxion dark matter — result
$m_{\text{Kaxion}} = 8.1\ \mu\text{eV}$ ($\pm 1.3\%$)
$f_{\text{Kaxion}} = 7.1 \times 10^{11}$ GeV (calibrated)
$\nu_K \in [1.59,\ 2.13]$ GHz
Status D (Derived-conditional).*

Figure 44: Kaxion dark matter from *vacuum misalignment*. **(a)** While $H > m$ the field angle is frozen by Hubble friction at the misalignment value $\theta_i \sim 1.1$ rad; at the oscillation onset $H = m$ it begins damped coherent oscillation. **(b)** Once oscillating, the energy density redshifts as $\rho \propto a^{-3}$, i.e. behaves as cold dark matter. **GG reading:** one scale is calibrated to the observed abundance ($f_{\text{Kaxion}} = 7.1 \times 10^{11}$ GeV, $m_{\text{Kaxion}} = 8.1\ \mu$eV, $\nu_K \in [1.59, 2.13]$ GHz); the same closure then *returns* $\theta_i \sim 1.1$ rad — a read-back, not a fit. The derivation is executed in companion GG-C1-Kaxion. Status D* (Derived-conditional); the calibration is declared, and θ_i is abundance-fixed.

Step 3. The geometric reading (§M.8). One scale is calibrated to the measured abundance at the benchmark cap loading and declared; the same closure then returns the one remaining angle as an *output* — a read-back, not a fit, and narrower than an inversion of the usual fit.

M.1 A light pseudo-scalar as a misalignment condensate

The actor. The Kaxion is an *axion-like field*: a light, electrically neutral, spin-zero, parity-odd (pseudo-scalar) degree of freedom ϕ whose only renormalizable self-interaction is forbidden by a shift symmetry $\phi \to \phi + \text{const}$. Because of that shift symmetry the field has no ordinary potential to any order in perturbation theory; a potential can be generated only by non-perturbative (topological) effects, and when it is, it is necessarily *periodic*. It is convenient to trade the dimensionful field ϕ for a dimensionless *angle*

$$\theta \;\equiv\; \frac{\phi}{f}, \qquad \theta \in [-\pi, \pi), \tag{755}$$

where f is the *decay constant*, the single scale that sets both the field range (ϕ winds by $2\pi f$ around its circle) and the strength of all the field's couplings (which enter as $1/f$). The angle θ is literally the position of a dial; the shift symmetry is the statement that rotating the dial costs nothing until topology intervenes.

The potential. Topology supplies a periodic potential of the universal cosine form

$$V(\theta) \;=\; \chi\,[\,1-\cos\theta\,] \;=\; m^2 f^2\,[\,1-\cos\theta\,], \qquad m^2 \;\equiv\; \frac{\chi}{f^2}, \tag{756}$$

where χ is the *topological susceptibility* of whatever gauge sector the angle talks to — it has dimensions of (energy)^4 and absorbs the exponentially small instanton weight $e^{-S_{\text{inst}}}$. Expanding for a small displacement, $V \simeq \frac{1}{2}m^2 f^2\theta^2 = \frac{1}{2}m^2\phi^2$: the floor of each cosine well is a harmonic oscillator of mass m. Two facts are worth fixing in the mind before any cosmology is done. First, the mass is *not* a fundamental scale; it is the ratio of a topological scale $\sqrt{\chi}$ to a geometric scale f, so $mf = \sqrt{\chi}$ exactly — a relation we will use twice. Second, the curvature at the floor (m^2) and the height of the rim (2χ) are wildly different scales, which is what lets the field store a cosmologically long-lived, nuclear-stiff yet feather-gentle oscillation.

A condensate, not a gas. Crucially, the dark matter here is not a thermal gas of particles produced by collisions in the hot plasma. It is a classical, spatially homogeneous, coherently oscillating *field* — a condensate. Every "particle" is a quantum of one and the same zero-momentum oscillation. This is why the candidate can be extraordinarily light (micro-eV) and yet cold: its coldness comes from the homogeneity of the field, not from a low temperature of a particle gas. The whole of the misalignment mechanism is the story of how that homogeneous oscillation is set up and how its energy redshifts.

M.2 Equation of motion in an expanding universe

Take the field homogeneous, $\theta = \theta(t)$, in a spatially flat Friedmann–Robertson–Walker (FRW) background with scale factor $a(t)$ and Hubble rate $H = \dot{a}/a$. The action for $\phi = f\theta$ is the standard scalar action; varying it and dividing by f gives the equation of motion

$$\boxed{\ddot{\theta} \;+\; 3H\,\dot{\theta} \;+\; m^2(T)\,\sin\theta \;=\; 0.} \tag{757}$$

This is the master equation of the entire mechanism, and it is worth reading term by term, because it is exactly the equation of a pendulum in a viscous medium.

- $\ddot{\theta}$ is the inertia of the dial.
- $m^2(T)\sin\theta$ is the topological restoring torque, written with the full periodic $\sin\theta$ rather than its linearization. The argument T is a reminder that for the Kaxion the susceptibility, and hence the mass, is *temperature dependent*: above the confinement scale of the gauge sector the instantons are suppressed and $m(T)$ is much smaller than its zero-temperature value. We carry this dependence explicitly because it is the origin of the famous non-integer exponent in the abundance formula.
- $3H\,\dot{\theta}$ is a friction term, with friction coefficient $\gamma = 3H$ supplied for free by the expansion of the universe. This is the cosmological version of Hubble drag: as space expands, the field's momentum redshifts away. In the GG-Theory boundary law this same $3H$ appears as the expanding-background friction coefficient of the open-system (GRIP/GDOS) Langevin equation applied to the compact-angle mode, with the post-inflationary noise temperature utterly negligible; here we need only its classical, frictional content.

Harmonic (small-angle) limit. For $|\theta| \ll 1$ we may replace $\sin\theta \to \theta$ and (757) becomes the damped harmonic oscillator

$$\ddot{\theta} \;+\; 3H\,\dot{\theta} \;+\; m^2(T)\,\theta \;=\; 0, \tag{758}$$

a frictionally damped pendulum whose two regimes — overdamped freeze and underdamped ring — are fixed entirely by the competition between the damping rate $3H$ and the natural frequency m. The whole cosmological story is the crossover between those two regimes as H falls while m (eventually) rises.

M.3 Frozen by Hubble friction, then released

The overdamped (frozen) phase: $H \gg m$. Early in the radiation era H is enormous (it scales as $T^2/M_{\rm Pl}$) and the natural frequency m is tiny. The pendulum is then overdamped: the restoring torque cannot overcome the friction, and the field is *frozen* at whatever value it was left with after inflation. Formally, in (758) with $H \gg m$ the two terms that survive are $3H\dot{\theta} + m^2\theta \approx 0$, whose slow solution drifts only by

$$\Delta\theta \;\sim\; \Big(\frac{m}{3H}\Big)^2\,\theta \;\ll\; \theta, \tag{759}$$

i.e. negligibly per Hubble time until m approaches $3H$. The field sits at its initial displacement

$$\theta(t) \;\simeq\; \theta_i \qquad (\text{the } \textit{misalignment angle}), \tag{760}$$

the value the field happened to have when the potential first switched on relative to the Hubble friction. "Misalignment" is exactly this: the dial is generically displaced from the bottom of its well, and Hubble friction pins it there, away from the minimum, like a stuck door.

Onset of oscillation: $H \sim m$. As the universe expands, H falls monotonically; eventually the friction weakens to the point where the restoring torque wins. The conventional onset condition is

$$\boxed{3H(T_1) \;=\; m(T_1),} \tag{761}$$

which defines the *onset temperature* T_1 (also written $T_{\rm osc}$). The numerical coefficient "3" is a convention, not magic: the exact solution of (758) in the radiation era is a Bessel function whose envelope turns over precisely at $mt \simeq \frac{3}{2}$, i.e. at $3H = m$ (recall $H = 1/2t$ in radiation domination). So $3H = m$ is the mathematical knee of the special function, not an arbitrary marker. After T_1 the field is underdamped and begins to oscillate coherently about the minimum of its well, with initial amplitude θ_i.

Determining T_1 (a worked aside). The radiation-era Hubble rate is

$$H(T) \;=\; \sqrt{\frac{8\pi^3 g_*(T)}{90}}\,\frac{T^2}{M_{\rm Pl}} \;\simeq\; 1.66\,\sqrt{g_*}\,\frac{T^2}{M_{\rm Pl}}, \tag{762}$$

with $M_{\rm Pl} = 1.22\times10^{19}$ GeV the (non-reduced) Planck mass and g_* the effective relativistic degrees of freedom. If one *pretended* the mass were already at its present value $m_K = 8.1\,\mu\text{eV} = 8.1\times10^{-15}$ GeV, then (761) with (762) ($g_* \simeq 73$, so the prefactor combination is $\simeq 42.6$) gives $T_1 \simeq [\,m_K\,M_{\rm Pl}/42.6\,]^{1/2}$, of order a few tens of GeV. But the Kaxion's mass is *not* constant: above the QCD scale the susceptibility is instanton-suppressed and rises as the universe cools, $m(T) = m(0)\,(T_{\rm QCD}/T)^b$ with a lattice exponent $b \simeq 4$. Inserting this temperature dependence into (761) lowers the true onset to $T_1 \simeq 1.0$ GeV ($t \simeq 0.3\,\mu$s), where the mass is still only $\sim 0.04\%$ of its final value: the spring is soft when the pendulum is released and stiffens underneath the swing. This "stiffening spring" is the physics behind the non-integer power that will appear in (768); it is not a nuisance but the signature of a QCD-coupled axion.

M.4 Redshifting as cold, pressureless matter

Once the field oscillates with $H \ll m$, the friction term in (757) is a slow drag on a fast oscillation. Average over one cycle. For a harmonic well the virial theorem gives equal kinetic and potential energy on average, so the pressure $p = \langle \tfrac{1}{2}\dot\phi^2 - V \rangle$ averages to zero, while the energy density

$\rho = \langle \frac{1}{2}\dot{\phi}^2 + V \rangle = \frac{1}{2} m^2 \phi_{\rm amp}^2$ is non-zero. A pressureless fluid in FRW redshifts as

$$\boxed{\rho_\theta \;\propto\; a^{-3}.} \tag{763}$$

This is the definition of *cold* (non-relativistic, pressureless) dark matter: the oscillating condensate dilutes exactly like dust, $\rho \propto a^{-3}$, never like radiation (a^{-4}). It is cold despite the field's tiny mass because the whole condensate sits at zero momentum.

The adiabatic-invariant (WKB) argument. The cleanest way to see (763) — and the only way that survives the stiffening spring, when m itself is changing — is the adiabatic theorem for an oscillator of slowly varying frequency. For a harmonic oscillator with slowly varying $\omega(t) = m(t)$ the *action variable*

$$J \;=\; \oint p\, dq \;=\; 2\pi \frac{E}{\omega} \tag{764}$$

is an adiabatic invariant whenever the WKB (adiabaticity) condition $|\dot{\omega}/\omega^2| \ll 1$ holds. Per comoving volume, E/ω is exactly the *number of quanta*: the energy density of the condensate is $\rho = m\, n$, where n is the number density of zero-momentum quanta. Hence the comoving quantum number

$$N \;=\; n\, a^3 \;=\; \frac{\rho\, a^3}{m} \;=\; \text{const.} \tag{765}$$

Energy is *not* conserved (as the spring stiffens it does work on the oscillation, and the redshift drains momentum), but the *number* of quanta per comoving volume is. Two consequences follow immediately. (i) Once the mass has settled to its constant value $m(0)$, (765) gives $\rho = m(0)\, n \propto a^{-3}$, recovering (763). (ii) Through the stiffening epoch we do *not* track the energy (which is changing in a complicated way) but the conserved N; this is precisely the right bookkeeping to carry from the onset T_1 down to today.

One honest caveat. The WKB condition is only *marginally* satisfied at the very first oscillation: at onset $\dot{m}/m^2 = bH/m(T_1) = b/3 \simeq 1.4$, so adiabaticity holds order-of-magnitude at onset and becomes excellent ($\dot{m}/m^2 \propto t^{-1}$) thereafter. The $\mathcal{O}(1)$ transient this induces over the first cycle or two is exactly what a numerical transfer-function calculation resolves, and it is carried as one row of the $\pm 4\%$ the companion puts on the angle (Arisaka, 2026, C1-Kaxion), behind the susceptibility envelope. We flag it so that the reader does not mistake the clean analytic scaling below for better-than-it-is precision.

M.5 The relic-abundance formula

Now assemble the bookkeeping. At onset the field banks the potential energy of its displacement,

$$\rho_\theta(T_1) \;=\; \chi(T_1)\,[1 - \cos\theta_i] \;\simeq\; \tfrac{1}{2}\, m^2(T_1)\, f^2\, \theta_i^2 \quad \text{(small-angle)}, \tag{766}$$

and the comoving number N of (765) is fixed from this moment on. Tracking $\rho = m(0)\, n$ with na^3 constant from T_1 to today, normalizing the entropy density to the photon temperature, and using

the onset relation (761) with the temperature-dependent mass, one obtains the present comoving density. The transparent way to see the scaling is to insert $m(T_1) = 3H(T_1) \propto T_1^2$ and the onset solution $T_1 \propto m(0)^{1/(b+2)}$ into $\rho_0 \propto m(0)\, n_0$; the powers combine to

$$\rho_0 \;\propto\; m(0)^{(b+3)/(b+2)}\,\theta_i^2 \;\propto\; f^{-(b+3)/(b+2)}\,\theta_i^2, \tag{767}$$

where the last step used the frame-invariant $m(0) \propto 1/f$ (from $mf = \sqrt{\chi}$ with χ fixed by the gauge sector). For the lattice exponent $b \simeq 4.08$ this gives the celebrated near-7/6 power $(b+3)/(b+2) = 7.08/6.08 = 1.165$. Restoring the numerical normalization gives the master formula

$$\boxed{\Omega_a h^2 \;=\; 0.135\;\theta_i^2\left(\frac{f}{10^{12}\,\mathrm{GeV}}\right)^{1.165} \qquad \text{(harmonic)},} \tag{768}$$

where the prefactor is the direct integration of the equation of motion (757) through the onset, with the lattice susceptibility $\chi(T) \propto T^{-2b}$ above $T_{\rm QCD}$ (so that $m \propto T^{-b}$) and entropy conservation carried in the clock, as performed in (Arisaka, 2026, C1-Kaxion). The instantaneous-onset bookkeeping that sets $3H = m$ at one moment and banks the energy there — the estimate written out above — gives 0.090 instead: it omits the transient in which the field, already swinging, keeps banking energy while $m(T)$ still rises, and understates the relic by a factor 1.5. An earlier version of this appendix printed the estimate as the master formula; the integrated value supersedes it, and the estimate is kept here only as the scaffold that explains the exponent. Read the formula as a function of the three inputs it actually contains: the misalignment angle θ_i, the decay constant f, and the (temperature-dependent) mass — the latter entering through the exponent 1.165, which is a fingerprint of the QCD running, not a free dial.

Why the odd exponent 1.165 and not 2? Equation (766) would naively suggest $\rho \propto \theta_i^2 f^2$. The reduction from f^2 to $f^{1.165}$ is pure physics: a larger f means a *lighter* field, which lets go *later* (lower T_1), when the universe has less entropy left to dilute it — but it also banks each quantum at a smaller onset mass. The competition between "later release" (which boosts the relic) and "cheaper quanta" (which suppresses it) shaves the naive power down to $(b+3)/(b+2)$. A reader who can reconstruct that sentence owns the misalignment mechanism.

Anharmonic correction for $\mathcal{O}(1)$ angles. The small-angle banking (766) is only the leading term. For an initial displacement of order unity the field samples the flatter, anharmonic shoulders of the cosine, it takes *longer* to begin oscillating (the restoring torque is weaker than the harmonic estimate near the top), and the relic is enhanced relative to (768). The standard treatment writes the correction as a multiplicative *anharmonicity factor* $F(\theta_i)$,

$$\Omega_a h^2 \;=\; 0.135\;F(\theta_i)\,\theta_i^2\left(\frac{f}{10^{12}\,\mathrm{GeV}}\right)^{1.165}, \qquad F(\theta_i) \;\simeq\; \left[\ln\frac{e}{1-\theta_i^2/\pi^2}\right]^{(b+3)/(b+2)}, \tag{769}$$

with $F \to 1$ as $\theta_i \to 0$ and F growing as $\theta_i \to \pi$ (where the field starts near the unstable hilltop and the relic diverges logarithmically); the exponent is the same $(b+3)/(b+2) \simeq 1.16$ as the

decay-constant power, since the correction acts through the same delayed onset. For the angle this paper uses the companion does not lean on the analytic factor at all: it integrates the full cosine on the tabulated susceptibility and equation of state, and the closure returns $\theta_i = 1.18$ where the closed-form scaffold returns 1.16. Two corrections make up that difference and they pull against each other: the full cosine lowers the angle by 4% against a harmonic force on the same tables, and integrating on the tables rather than on a fixed g_* with a power-law turn-on raises it by about 5%. Earlier revisions of both papers applied the first to a baseline that had not received the second and printed 1.12. None of it touches the mass or frequency, which depend on f alone.

M.6 The QCD relation $m_a f_a = \sqrt{\chi_{\rm QCD}}$

For a QCD axion — and the Kaxion is the GG-Theory program's QCD axion — the gauge sector whose topology generates the potential is the strong interaction itself. The relevant susceptibility is then the *QCD topological susceptibility* $\chi_{\rm QCD}$, a pure Standard-Model quantity measured on the lattice. The mass relation (756) becomes the model-independent strong-CP relation

$$\boxed{m_a\, f_a \;=\; \sqrt{\chi_{\rm QCD}}.} \tag{770}$$

The current lattice determination is

$$\chi_{\rm QCD}^{1/4} \;=\; 75.5(5)\ \text{MeV (lattice)} \qquad\Longrightarrow\qquad \sqrt{\chi_{\rm QCD}} \;=\; (75.5\,\text{MeV})^2, \tag{771}$$

(Arisaka, 2026, C1-Kaxion), so that the entire axion mass–coupling line is fixed by one experimentally measured number: choose f_a and the mass m_a is determined, and vice versa. This is the single most important structural fact about a QCD axion — it has a *one-parameter* family of allowed values, a line in the mass–coupling plane, not a two-dimensional region. Everything that follows is about where on that line the geometry of the GG-Theory program places the Kaxion.

M.7 What the standard misalignment story leaves free

The mechanism of Step 1 is exact, but it does not predict a dark-matter candidate; it parametrizes one. Two numbers are left entirely free: the decay constant f_a (equivalently the axion mass m_a), which can sit anywhere across many decades, and the initial misalignment angle $\theta_i \in [-\pi, \pi]$. The standard procedure fixes neither — it scans f_a along the "axion line" and, at each point, tunes θ_i so the relic abundance matches the observed $\Omega_{\rm DM}h^2 \approx 0.12$. The mechanism is a framework that *accommodates* a dark-matter abundance for almost any mass; it does not *predict* the mass. The geometric account removes this freedom in two steps of different grade: a ladder from the ledger's integers, and one calibration, declared.

M.8 The geometric reading: one calibration, then an inversion

The standard axion story leaves *both* f_a and θ_i free: one scans the line and tunes the angle to hit the observed abundance. The GG-Theory reading removes the scan, and it must be stated exactly, because an earlier version of this appendix stated it as more than it is.

One scale is calibrated, and declared. The physical decay constant is the bulk Green–Schwarz normalization reduced by an integer cap loading,

$$f_K \;=\; \frac{f_{\rm bulk}}{\Xi_{\rm cap}}, \qquad \Xi_{\rm cap} = K{+}\Sigma \in \{35, 38, 40, 42, 43, 45, 47\}, \qquad \text{benchmark } 43 = K{+}N{+}L, \tag{772}$$

and the bulk normalization is not derived: the orientation convention of the anti-self-dual tensor sector that would fix it is an open item of GG-A3-BI6 (its O14), and neither face-value chiral norm reproduces the value the abundance requires. The companion (Arisaka, 2026, C1-Kaxion) therefore *calibrates* $f_{\rm bulk} = 3.04 \times 10^{13}$ GeV through the relic closure — the benchmark row saturates the π-locked share at a generic angle — and declares the calibration. The ladder of cap loadings is read from the ledger by an ansatz (K plus a subset sum of $\{N, L, J\}$ with at most a pair — a restriction the ansatz stipulates, not derives), graded as such. What is left after the calibration is arithmetic: $f_K = 7.1 \times 10^{11}$ GeV at the benchmark, the seven rungs a factor 0.915 to 1.23 about it (43/47 to 43/35). There is exactly one physical decay constant; the bulk scale is the bookkeeping quantity $\Xi_{\rm cap} f_K$, not the decay constant of any propagating field. An earlier version of this appendix wrote "the geometry fixes f_K to ±6%"; the ±6% was the abundance pipeline's uncertainty propagated into a scale the abundance itself calibrates, and it is retired. The mass and frequency at fixed f_K carry only the lattice susceptibility's ±1.3%.

The abundance returns the misalignment angle. With f_K calibrated, the exponent fixed by lattice QCD, and the target abundance fixed by the π-locked cosmic inventory of the present paper — the dark-matter relic

$$\Omega_{\rm DM} h^2 \;=\; 0.1215, \qquad \Omega_{\rm DM} \;=\; \Omega_m - \Omega_b \;=\; \frac{1}{\pi} - \frac{1}{2\pi^2}, \tag{773}$$

converted to a comoving density with $h = 0.674$, the Planck-fit CMB-branch convention the companion adopts (at this paper's carried rate $H_0 = 66.8$ the same share is 0.1194, and at the two-road reading 66.6 it is 0.1187) — the master formula (768) has *exactly one* unknown left: the misalignment angle. Solving for it,

$$\boxed{\begin{aligned} \theta_i \;&=\; \left(\frac{0.1215}{0.135}\right)^{1/2} \left(\frac{10^{12}\,\text{GeV}}{f_K}\right)^{0.583} \;\simeq\; 1.16 \text{ rad (closed-form scaffold)}, \\ \theta_i \;&=\; 1.18 \pm 0.04 \text{ rad (tabulated, full cosine: the calculation of record)}. \end{aligned}} \tag{774}$$

The direction of inference is the point, and it is narrower than the word "inversion" suggests: the abundance (773) is the target on the right-hand side; the one scale is calibrated so that a generic angle meets it; and the angle is read back as the generic $\mathcal{O}(1)$ number it is. Nothing is fitted after the calibration, and the calibration is declared, once, in both papers. The closure is carried out in full in (Arisaka, 2026, C1-Kaxion).

Honesty discipline on the angle. It must be stated plainly, because the shape of the number invites a natural misreading: θ_i is *not* geometry-selected at $1/(2\pi)$, and it is not derived. The

misalignment angle is an $\mathcal{O}(1)$ *output* of the observed abundance via (774), 1.18 ± 0.04 rad with the tabulated susceptibility and equation of state carried ($\pm 4\%$; the susceptibility's own pointwise envelope and the next-to-leading-order matching residual are the two largest rows, the power-law exponent and anchor rows having been retired by integrating on the table). It is not a tuned input, and it is not selected by the geometry; it is what the calibrated scale plus the cosmic inventory return.

The single-patch history: a conjecture, named. The standard story's deepest assumption is statistical: θ_i is drawn at random in each causal patch, so the relic density is an accident of our patch, and the field's fluctuations during inflation imprint *isocurvature* perturbations that the CMB constrains. The present framework replaces the draw with a conjecture, and an earlier version of this appendix called it a mechanism. During the primary (GUT-adjacent) inflation of Appendix G, with $H_1 \simeq 1.5 \times 10^{13}$ GeV, a light angle would be randomized on every scale that exited the horizon; the single-patch history therefore requires the angle to have carried an effective mass $m_{\rm pin} \gg H_1$ through those e-folds — Conjecture C-Kaxion-d of the companion, as this paper states it in Part V (§35.5.2) — whose minimum sits displaced by the misalignment angle from the CP-conserving point, which releases the angle at the lock crossing identically in every patch, and which is absent afterward to the level the neutron bound allows. No sector of the ledger is presently known that could supply it, and neither this paper nor the companion computes it. Three consequences follow *under the conjecture*, each observable in principle. *(i) Patch-uniformity:* there is no axion-abundance lottery; the relic density is deterministic. *(ii) Zero primordial isocurvature:* a pinned angle does not fluctuate, and the lock crossing that releases it is prompt and imprints no super-horizon perturbation, so $\beta_{\rm iso}$ vanishes identically. *(iii) No defects:* no patch-to-patch angle variation means no domain walls or strings from this sector, whichever value the companion's open domain-wall number $N_{\rm DW}$ (1 or 5) takes. The failure modes are equally plain, and one of them is fatal rather than merely informative. Without the pinning the angle is not randomized after inflation — the compact angle has no symmetry restoration to be randomized by — but *during* it, at 3.4 rad per e-fold, which decorrelates the angle in about a quarter of an e-fold across a sky spanning seven; the resulting cold-dark-matter isocurvature is of order unity and exceeds the *Planck* bound $\beta_{\rm iso} < 0.038$ by ten orders of magnitude in power. At the bulk scale the angle is not randomized, and its ordinary pre-inflationary fluctuation still exceeds the bound by eight orders; no decay constant below 5×10^{17} GeV — four orders above any the geometry supplies — evades it. The unpinned branch is therefore excluded, not merely different: its miniclusters and its line at $\gtrsim 6$ GHz from misalignment alone are what would be seen if the sector were something other than what this paper computes. A partially pinned angle, $m_{\rm pin} \sim H_1$, is not computed here: whether any isocurvature survives to the microwave background depends on the pinning mass between the end of inflation and the lock crossing, which the conjecture does not specify, and at f_K a mass of order H_1 does not hold the angle at all (its de Sitter rms is several radians). The paper treats the conjecture as two-valued. A confirmed Kaxion minicluster population or a confirmed line at the randomized-angle mass kills the conjecture and with it this paper's compact-angle route to the dark-matter share; a CMB isocurvature detection attributable to the axion sector kills the conjecture outright: the angle

was not held. In the identity language of Part X: the electron is identical everywhere because one ground basin exists; the Kaxion is identical everywhere *if* one angle was frozen — two geometric routes to the same universality, the second carried as a conjecture, and the isocurvature null its laboratory-facing signature. The main text places this in the trio it belongs to (§31.2): the release at the lock crossing is the moment the interval's episode hands the circle's phase its displacement.

The mass is the DFSZ/QCD image of the decay constant. With f_K calibrated by (772) and $\sqrt{\chi_{\rm QCD}}$ fixed by lattice QCD (771), the mass is not a separate prediction but the QCD image of f_K through (770):

$$m_K \;=\; \frac{\sqrt{\chi_{\rm QCD}}}{f_K} \;=\; 8.1\ \mu\text{eV}\ (\pm 1.3\%\ \text{at fixed}\ f_K), \tag{775}$$

on the DFSZ coupling line, giving the benchmark and the laboratory window

$$\nu_K \;=\; 1.95\ \text{GHz}\ (\pm 25\ \text{MHz}), \qquad \nu_K \;\in\; [\,1.59,\ 2.13\,]\ \text{GHz}\ (\text{seven rungs}), \tag{776}$$

the band the companion's haloscope program (IBS-CAPP, ADMX-EFR) addresses (Arisaka, 2026, C1-Kaxion). The width of the window is the image of the *discrete* ladder of cap loadings on f_K, not a continuous knob; the rungs are 45–136 MHz apart, wider than the ±25 MHz the susceptibility leaves except for the 42/43 pair, separated at only 1.8σ, so a line on a rung reads $\Xi_{\rm cap}$ under the ansatz and a line between rungs retires the ansatz.

M.9 Worked numerical example

We close with the arithmetic a reader can redo by hand, so that no number above is taken on faith.

(1) Mass from the decay constant. Insert $f_K = 7.07 \times 10^{11}$ GeV $= 7.07 \times 10^{20}$ eV and $\sqrt{\chi_{\rm QCD}} = (75.5\,\text{MeV})^2 = (7.55 \times 10^7\,\text{eV})^2 = 5.70 \times 10^{15}\,\text{eV}^2$ into (770):

$$m_K \;=\; \frac{\sqrt{\chi_{\rm QCD}}}{f_K} \;=\; \frac{5.70 \times 10^{15}\,\text{eV}^2}{7.07 \times 10^{20}\,\text{eV}} \;=\; 8.06 \times 10^{-6}\ \text{eV} \;=\; 8.1\ \mu\text{eV}, \tag{777}$$

the benchmark of (775).

(2) Laboratory frequency. A non-relativistic axion of rest mass m_K converts in a haloscope to a photon of frequency $\nu_K = m_K c^2/h$, i.e. in natural units $\nu_K = m_K/(2\pi)$ with m_K read as an energy. Using $1\,\text{eV} = 2.418 \times 10^{14}$ Hz,

$$\nu_K \;=\; (8.06 \times 10^{-6}\,\text{eV}) \times (2.418 \times 10^{14}\,\text{Hz/eV}) \;\simeq\; 1.95\ \text{GHz}, \tag{778}$$

the benchmark rung of (776). (The ±1.3% of the susceptibility is ±25 MHz; the 1.59–2.13 GHz band is the ladder, $\Xi_{\rm cap} = 35$ to 47.)

(3) The misalignment angle that the relic returns. Finally, read the angle back from $\Omega_{\rm DM}h^2 = 0.1215$. With $f_K = 7.07 \times 10^{11}$ GeV, $(f_K/10^{12}\,\text{GeV})^{1.165} = (0.707)^{1.165} \simeq 0.667$, so the harmonic master formula (768) reads $\Omega_K h^2 \simeq 0.135 \times 0.667 \times \theta_i^2 \simeq 0.090\,\theta_i^2$. Setting this equal to the target gives

$$\theta_i \;=\; \left(\frac{0.1215}{0.090}\right)^{1/2} \;\simeq\; 1.16 \text{ rad (scaffold)}, \qquad 1.18 \text{ (tabulated, full cosine)}, \tag{779}$$

exactly (774). (The coincidence that the product 0.135×0.667 equals the retired instantaneous-onset prefactor 0.090 at $f = 10^{12}$ GeV is numerical, not physical.) The three numbers — $m_K \simeq 8.1\,\mu$eV, $\nu_K \simeq 1.95$ GHz, $\theta_i \simeq 1.1$ rad — are the whole quantitative content of the Kaxion-as-dark-matter reading, and all three follow from one calibrated scale, the lattice $\chi_{\rm QCD}$, and the π-locked abundance.

Status

- **Relic mechanism: Derived.** The misalignment mechanism — frozen field, onset at $3H = m$, $\rho \propto a^{-3}$ coldness via the adiabatic invariant, and the abundance formula (768) with its anharmonic correction — is standard physics and is derived in full in (Arisaka, 2026, C1-Kaxion). It is not in question here.

- **The scale: calibrated, and declared.** The direction of inference is calibration $\to f_K$, abundance $\to \theta_i$: one scale is set by the observed abundance at the benchmark loading, the abundance (773) is the target, and $\theta_i \simeq 1.1$ rad is the $\mathcal{O}(1)$ read-back. An earlier version of this appendix graded this step "Derived" and the direction "geometry $\to f_K$"; it is neither, and the codebook row for $f_{\rm Kaxion}$ says so.

- **The absolute mass: Derived-conditional (D*).** $m_K = 8.1\,\mu$eV is the QCD image of f_K through (770), so it is as firm as the calibration and the ladder: $\pm 1.3\%$ at fixed f_K, and a discrete window of seven rungs from the cap-loading ansatz. What would promote it is the orientation computation behind GG-A3-BI6's O14 (turning the calibration into a derivation) together with a defining computation for the cap loading (turning the ansatz into a theorem and the window into a point).

- **The misalignment angle: an O(1) output, not an input.** $\theta_i = 1.18 \pm 0.04$ rad is the output of the observed abundance via (774), not a tuned parameter and not geometry-selected at $1/(2\pi)$, which lies a factor of about seven below the generic range that any anharmonic $\mathcal{O}(1)$ misalignment produces.

- **The single-patch history: conjectured.** The relic calculation presumes one angle for the observable universe; that presumption is Conjecture C-Kaxion-d, carried by name and with its killers, and computed by neither paper.

In one sentence: the Kaxion is the program's QCD axion, its decay constant is calibrated once through the observed abundance and then held fixed, its mass and microwave frequency are the lattice-QCD image of that decay constant on a ladder the ledger's integers supply by ansatz, and

the one number the cosmology hands back — the misalignment angle $\simeq 1.1$ rad — is an honest $\mathcal{O}(1)$ output of the measured dark-matter abundance, not a dial that was turned to fit it.

Problems and worked solutions. M.1.1. Why is a misalignment condensate naturally light, and why is its potential a cosine?

Solution. Step 1: an angular field θ valued on a circle enjoys a shift symmetry $\theta \to \theta +$ const; any potential must be periodic, hence a Fourier series in $\cos(n\theta)$ — the leading term is the cosine of (756). Step 2: the symmetry is exact until non-perturbative (instanton-class) effects switch on, so the mass $m = \sqrt{\chi}/f$ (770) is suppressed by the topological susceptibility χ — small because it is born of tunneling, not of couplings. Step 3: a field with f near the deep-bulk scale and χ from the QCD sector is therefore automatically in the micro-eV range: lightness is structural, not tuned. Physics intuition: angles cannot have generic masses — only winding can weigh them — so nature's lightest scalars are her compact ones; the Kaxion is the cohort's compact angle doing exactly what compact angles do.

M.2.1. Derive the damped-oscillator equation of motion from the action.

Solution. Step 1: vary the canonical scalar action (543) with $\phi = f\theta$ and $V = m^2 f^2(1 - \cos\theta)$ (§M.2) for a homogeneous $\theta(t)$ in the FRW background: the measure a^3 is what makes the expansion enter. Step 2: the Euler–Lagrange equation is $\ddot{\theta} + 3H\dot{\theta} + m^2 \sin\theta \cdot (\ldots) = 0$ — (757) — whose small-angle form is the damped harmonic oscillator (758). Step 3: the $3H\dot{\theta}$ term is not friction against anything material: it is dilution — the field's canonical momentum lives in a growing volume. Physics intuition: "Hubble friction" is bookkeeping for expansion, which is why it acts on every field equally and why the release condition (M.3.1) is a race between one universal clock (H) and one private one (m).

M.3.1. Establish the freeze/release criterion $m \simeq 3H$ and the onset temperature.

Solution. Step 1: while $m \ll 3H$, the oscillator is overdamped: (759) shows the field slips only by $O(m^2/H^2)$ per Hubble time — the angle is frozen at θ_i (760). Step 2: when $m \simeq 3H$ (761), the private clock outruns the universal one and oscillation begins; in the radiation era $H \sim T^2/M_{\rm Pl}$, so the onset temperature obeys $T_{\rm osc} \sim \sqrt{mM_{\rm Pl}}$ up to g_* factors. Step 3: for micro-eV masses this lands at tens of GeV for a constant mass; the instanton-suppressed $\chi(T)$ delays the onset to $T_{\rm osc} \simeq 1$ GeV, where the susceptibility is still rising toward its vacuum value. Physics intuition: the misalignment relic is a delayed-start pendulum — its abundance depends on when it started swinging, which converts a mass into a cosmic timestamp; freeze-out grammar (F.5.1) again, with the roles of Γ and m interchanged.

M.4.1. Prove that the oscillating condensate redshifts as pressureless matter.

Solution. Step 1: once $m \gg H$, the field oscillates many times per Hubble time; over a cycle the virial theorem gives $\langle\text{kinetic}\rangle = \langle\text{potential}\rangle$, so the cycle-averaged pressure $\langle p\rangle = \langle\dot{\phi}^2/2 - V\rangle = 0$. Step 2: equivalently, the comoving occupation (765) $N \propto \rho a^3/m$ is an adiabatic invariant of the slowly-expanding oscillator, so with m constant, $\rho \propto a^{-3}$ (763). Step 3: a coherent field of macroscopic occupation thus gravitates exactly like cold particles — *cold* dark matter from a wave. Physics intuition: the condensate is a swarm of zero-momentum quanta in disguise; dilution can thin

the swarm but cannot slow particles that were never moving, which is the whole phenomenological content of "cold."

M.5.1. Count the freedoms: why is the cohort's misalignment story narrower than the standard one, and where is it still open?

Solution. Step 1: the standard axion abundance (768) contains three inputs — the decay constant f, the mass chain through χ, and the initial angle θ_i — and one output, $\Omega_{\rm DM}$: with θ_i unknown, any abundance can be accommodated (the "anthropic window"). Step 2: the cohort *fixes the output*, $\Omega_m = 1/\pi$ from Appendix L, and removes the scan over f in two steps of different grade: the ledger's integers put f on a discrete ladder of seven rungs (an ansatz), and one scale — the bulk normalization behind the ladder — is calibrated so that a generic angle meets the abundance at the benchmark rung, then declared and held fixed. The master formula then has exactly one unknown left, and solving gives $\theta_i \simeq 1.1$ rad (774) — an $O(1)$ angle, uncontrived, read back rather than drawn. Step 3: what is still assumed is named: that the observable universe carries *one* angle at all (the single-patch history, Conjecture C-Kaxion-d of the companion). Falsifiability follows: the haloscope line (the chapter's worked example) must hold, and a minicluster population or a line at the randomized-angle mass would kill the single-patch conjecture and with it this paper's compact-angle route to the dark-matter share. Physics intuition: the same equation read left-to-right is a fit and right-to-left is a theory; the cohort reads it right-to-left after one declared calibration, and says which step is the calibration, which the ansatz, and which the conjecture — the cleanest statement of what "no dial after the data" means in practice.

Take-away. *Standard result:* misalignment makes a light axion an excellent cold-dark-matter candidate — with the decay constant scanned and the angle drawn. *What it assumed:* f_a free over orders of magnitude, θ_i random, isocurvature bounds survived by luck of the inflation scale. *What the geometry replaced it with:* f_K calibrated once at the benchmark cap loading and declared, θ_i read back, the single-patch history conjectured (pinning during the primary inflation: patch-uniform, zero isocurvature, no defects, if it holds), the mass at 8.1 μeV on the DFSZ line, and a laboratory window at 1.59–2.13 GHz that decides this decade. *Continue in:* Part V; GG-C1-Kaxion; Part X Table 43.

N Black-hole entropy A/4 from induced gravity

What this appendix teaches. Black-hole thermodynamics from scratch: the four laws, the Hawking temperature, and the Bekenstein–Hawking entropy; Sakharov's induced gravity and the heat-kernel computation in which *one* coefficient sets both $1/G_N$ and the entropy density — so that $S = A/4G_N$ becomes an identity rather than a coincidence; and the geometric reading in which the horizon is an OGGEP port and the Newton constant is the scale-profile readout of the same bulk that fixed the inventory.

Read alongside: Wald (1984) for the classical laws; Hawking (1975) and Bekenstein (1973) as the primary sources.

The body of this paper uses the Bekenstein–Hawking entropy $S = A/4G_N\hbar$ as an established thermodynamic input — it appears wherever the cosmological-boundary stratum has a horizon, most visibly in the de Sitter and apparent-horizon accounting of the late vacuum and in the open-system

reading of black-hole information. GG-A10-Cosmos does *not* prove that area law; it inherits it. The proof lives in the companion Arisaka (2026, C3-BH) (Theorem *T-BH-1*), where the coefficient 1/4 is derived from induced gravity by the conical-replica method on the canonical six-dimensional bulk. This appendix is written for the first-year graduate student who has met general relativity and a first course in quantum field theory but has never seen why a black hole should have an entropy at all, let alone why the number is exactly one quarter of the horizon area in Planck units. We build the result from scratch, in the longest and most explicit form the cohort offers, and we are scrupulous at the end about what is genuinely *derived* and what remains *open*. The honest one-line summary, stated here and defended at the close: the $A/4$ is a derivation the companion grades Derived-conditional on three named premises, credited to Arisaka (2026, C3-BH); the recovery of the Page curve is conceptually closed but quantitatively open.

The three steps of this appendix. As in Appendix E:

Step 1. The standard thermodynamics and induced gravity (§N.1 onward). The four laws, the Hawking temperature, the Bekenstein–Hawking entropy $S = A/4G_N\hbar$, Sakharov's induced gravity, and the heat-kernel and conical-replica derivations of the 1/4.

Step 2. What standard physics leaves open (§N.7). The area law is universally used, but its microscopic origin — why the entropy is exactly one quarter of the horizon area — is recovered from ingredients (G_N, an abstract manifold) that are themselves taken as given.

Step 3. The geometric derivation (§N.8). On the canonical six-dimensional bulk the 1/4 is the companion's T-BH-1, graded Derived-conditional on three named premises (Arisaka (2026, C3-BH)), with the induced Newton constant the *same* scale-overlap G_N that fixes gravity in Appendix E.

N.1 The four laws of black-hole mechanics and the thermodynamic analogy

Start with the object. A Schwarzschild black hole of mass M is the unique static, spherically symmetric, vacuum solution of Einstein's equation $R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi G_N T_{\mu\nu}/c^4$, with line element

$$ds^2 = -\Big(1 - \frac{2G_N M}{c^2 r}\Big)c^2\,dt^2 + \frac{dr^2}{1 - \dfrac{2G_N M}{c^2 r}} + r^2\big(d\theta^2 + \sin^2\theta\,d\phi^2\big). \tag{780}$$

The surface $r = r_s \equiv 2G_N M/c^2$, the *Schwarzschild radius*, is the event horizon: a one-way causal membrane, not a curvature singularity (the metric blowup there is a coordinate artifact, removable by Kruskal coordinates; the true singularity is at $r = 0$). The horizon is a two-sphere of area

$$A = 4\pi r_s^2 = \frac{16\pi G_N^2 M^2}{c^4}. \tag{781}$$

We will return to this formula for the worked example; for now note that A grows as M^2, so two black holes that merge increase the total horizon area.

In 1973 Bardeen, Carter, and Hawking established four laws governing the mechanics of stationary black holes, in striking one-to-one correspondence with the four laws of thermodynamics. Let κ

be the *surface gravity* (defined precisely in §N.2; for Schwarzschild $\kappa = c^4/4G_N M$), Ω the angular velocity of the horizon, J the angular momentum, Φ the electrostatic potential, and Q the charge. The laws read:

- **Zeroth law.** The surface gravity κ is constant over the horizon of a stationary black hole. This is the analog of the statement that temperature is uniform throughout a body in thermal equilibrium: $\kappa \leftrightarrow T$.
- **First law.** For a quasi-static change between neighboring stationary black holes,

$$dM = \frac{\kappa}{8\pi G_N}\, dA + \Omega\, dJ + \Phi\, dQ, \tag{782}$$

  which has exactly the form $dU = T\,dS +$ work of the first law of thermodynamics, with M (the energy) in the role of U, the $\Omega\, dJ + \Phi\, dQ$ terms the work, and the product $(\kappa/8\pi G_N)\, dA$ standing in for $T\, dS$. The pairing $T\, dS \leftrightarrow (\kappa/8\pi G_N)\, dA$ is the crucial one; it already insinuates that S is proportional to A.
- **Second law.** Hawking's area theorem: in any classical process obeying the null energy condition the total horizon area cannot decrease, $dA/dt \geq 0$. This is the analog of the Clausius statement $dS/dt \geq 0$: $A \leftrightarrow S$.
- **Third law.** An extremal black hole ($\kappa = 0$, e.g. a maximally charged or maximally rotating one) cannot be reached by any finite sequence of operations — the analog of the unattainability of absolute zero, $T = 0$.

The 1973 paper presented these as *analogies*. The match is suspiciously perfect — four laws, in order, with the same logical structure — but a pure analogy carries no physical entropy. Bekenstein took the analogy literally and argued that the area really is an entropy (Bekenstein, 1973); his argument was thermodynamic. If a black hole had genuinely zero entropy, one could lower an entropy-bearing box of gas through the horizon and erase its entropy from the visible universe, decreasing the total entropy of the world and violating the ordinary second law. To rescue the second law, the black hole must carry an entropy that increases by at least the swallowed amount when it absorbs the box. Bekenstein concluded that the black-hole entropy is proportional to the horizon area, $S_{\rm BH} \propto A/\ell_P^2$, where $\ell_P^2 = \hbar G_N/c^3$ is the Planck area — the only combination of constants with the right dimensions to convert an area into a pure number. What Bekenstein could not fix from thermodynamics alone was the dimensionless coefficient. Supplying that coefficient — showing it is exactly $1/4$ — is the entire technical content of this appendix.

N.2 The Hawking temperature $T_H = \hbar\kappa/2\pi$ and the surface gravity κ

The analogy became physics in 1974, when Hawking computed what happens to a quantum field on the Schwarzschild background and found that the black hole is not black: it radiates a thermal spectrum at a definite temperature (Hawking, 1975). The cleanest way to see the temperature without the full mode-by-mode Bogoliubov calculation is the Euclidean (imaginary-time) argument, which we sketch because it is the same trick the conical-replica derivation will industrialize later.

Surface gravity. For a static horizon the surface gravity κ is the acceleration, measured at spatial infinity, of a test particle held just outside the horizon — the redshifted local acceleration of a static observer. For Schwarzschild one finds

$$\kappa = \frac{c^4}{4G_N M} = \frac{c^2}{2r_s}. \tag{783}$$

It is finite and constant over the horizon (the zeroth law), and it diverges as $M \to 0$: small black holes have large surface gravity.

Euclidean periodicity fixes the temperature. Continue to imaginary time $t \to -i\tau$ and look at the near-horizon geometry. Set $\rho^2 \propto (r - r_s)$; the Schwarzschild metric near the horizon becomes, to leading order,

$$ds_E^2 \;\approx\; d\rho^2 + \rho^2\, d\!\left(\frac{\kappa\,\tau}{c}\right)^2 + (\text{transverse 2-sphere}), \tag{784}$$

which is the metric of a flat plane in polar coordinates $(\rho,\ \kappa\tau/c)$. A plane is smooth at its origin $\rho = 0$ (the horizon) *only* if the angular coordinate $\kappa\tau/c$ has period exactly 2π; any other period leaves a conical singularity — a cone with a deficit or excess angle — sitting on the horizon. Demanding smoothness fixes the period of Euclidean time,

$$\tau \sim \tau + \frac{2\pi c}{\kappa}. \tag{785}$$

But in equilibrium statistical mechanics a thermal density matrix $\rho \propto e^{-\beta H}$ is exactly a Euclidean path integral periodic in imaginary time with period $\hbar\beta = \hbar/k_B T$. The two periodicities must agree, so

$$\boxed{k_B T_H = \frac{\hbar\,\kappa}{2\pi c} \qquad \Longleftrightarrow \qquad T_H = \frac{\hbar\, c^3}{8\pi G_N M\, k_B}.} \tag{786}$$

This is the Hawking temperature. In the common units $c = 1$ it is the celebrated $T_H = \hbar\kappa/2\pi$; the conical-deficit remark below (785) — that the period of Euclidean time is tied to a deficit angle on the horizon — is precisely the seed that §N.6 will grow into the exact 1/4. Hold onto the phrase "deficit angle on the horizon"; the entire derivation turns on it.

The temperature is heart-stoppingly small for astrophysical masses and enormous for tiny ones. For a one-solar-mass hole (786) gives $T_H \approx 6.2 \times 10^{-8}$ K, seven orders of magnitude below the 2.7 K cosmic microwave background, so such a hole *absorbs* far more than it radiates and does not evaporate in the present epoch. A primordial hole of $\sim 10^{12}$ kg, by contrast, has $T_H \sim 10^{11}$ K and evaporates within a cosmological time — the regime relevant to the primordial-black-hole bounds elsewhere in the cohort.

N.3 The Bekenstein–Hawking entropy $S = A/4G_N\hbar$ and the first law $dM = T\,dS$

With a genuine temperature in hand, the entropy is no longer optional. Treat the black hole as a thermodynamic system whose energy is its rest energy $U = Mc^2$ and whose temperature is T_H of

(786). The thermodynamic first law for a non-rotating, uncharged hole is simply

$$dU = T_H\, dS_{\mathrm{BH}} \qquad \Longrightarrow \qquad dS_{\mathrm{BH}} = \frac{c^2\, dM}{T_H}. \tag{787}$$

Substitute $T_H = \hbar c^3/8\pi G_N M k_B$ and integrate from zero mass:

$$S_{\mathrm{BH}} = \int_0^M \frac{c^2\, dM'}{T_H(M')} = \int_0^M \frac{8\pi G_N k_B}{\hbar c}\, M'\, dM' = \frac{4\pi G_N k_B M^2}{\hbar c}. \tag{788}$$

Now eliminate M in favor of the area using (781), $M^2 = c^4 A/16\pi G_N^2$, to obtain the Bekenstein–Hawking entropy in its final form,

$$\boxed{S_{\mathrm{BH}} = \frac{k_B\, c^3\, A}{4\, G_N\, \hbar} = \frac{k_B\, A}{4\, \ell_P^2}, \qquad \ell_P^2 = \frac{\hbar G_N}{c^3}.} \tag{789}$$

In natural units $k_B = \hbar = c = 1$ this is the famous $S = A/4G_N$, and in Planck units it is one quarter of the horizon area measured in Planck cells, $S/k_B = A/4\ell_P^2$. The coefficient $1/4$ is now visible but *not yet explained*: it descended here from integrating the Hawking temperature, which we got from Euclidean periodicity, which we accepted on the strength of the semiclassical path integral. A skeptic is entitled to ask: where does the universal $1/4$ actually come from, mechanism and all, independent of which field is radiating and which regulator one uses? Answering exactly that question is the job of induced gravity.

N.4 Sakharov's induced gravity: the Einstein–Hilbert action is generated by integrating out fluctuations

The decisive idea is Sakharov's (Sakharov, 1967): gravity need not be a fundamental field with its own bare coupling. Sakharov proposed that the Einstein–Hilbert action *itself* is generated — induced — by the quantum vacuum fluctuations of all the *other* fields propagating in a curved background. In modern language, integrate out the matter (and, in GG-Theory, the bulk) fluctuations in a fixed metric $g_{\mu\nu}$; the resulting one-loop effective action $W[g] = \frac{1}{2}\ln\det\Delta$, expanded in curvatures, contains a term proportional to $\int\sqrt{g}\,R$. The coefficient of that term *is* $1/16\pi G_N$. Newton's constant is not put in by hand; it is read off the spectrum of the modes that were integrated out.

GG-Theory promotes this hypothesis from an option to a law. Under the Geometry–Gauge Equivalence Principle (GGEP, AX-2 of the cohort constitution) there is *no* bare four-dimensional Newton constant to begin with: four-dimensional gravity is the geometry of the six-dimensional bulk (Arisaka, 2026, A2-GG6), and the gravitational action a boundary observer reads is generated *entirely* by integrating out the bulk fluctuation modes through the open-system OGGEP projection GDOS (Arisaka, 2026, A5-GDOS). Sakharov's "induced" becomes the only option on the table, and a corollary that is merely suggestive in ordinary field theory becomes forced here. The corollary is this: the ultraviolet divergence of the horizon entanglement entropy is not a small correction sitting on top of a pre-existing $1/4G_N$ — there is no pre-existing G_N for it to correct — so the entanglement divergence *is* the $1/4G_N$. This is the Susskind–Uglum observation; in GG-Theory it is

a theorem rather than a coincidence, because GGEP removes the bare term that would otherwise let the two quantities drift apart.

N.5 The heat-kernel expansion: one coefficient sets both $1/G_N$ and the entanglement entropy

The technical engine behind both statements is a single object, the *heat kernel* of the fluctuation operator Δ on the four-dimensional boundary. Its short-proper-time expansion is

$$\operatorname{Tr} e^{-s\Delta} = \frac{1}{(4\pi s)^2}\int d^4x\,\sqrt{g_4}\,[a_0 + a_1\,s + a_2\,s^2 + \cdots], \qquad a_1 = \Big(\tfrac{1}{6} - \xi\Big)R, \tag{790}$$

where ξ is the curvature coupling ($\xi = 0$ for a minimal scalar) and s is the Schwinger proper-time parameter. The coefficient a_1 — the one multiplying the Ricci scalar R — is the common ancestor of both quantities we care about.

It sets $1/G_N$. The one-loop effective action is $W = \frac{1}{2}\ln\det\Delta = -\frac{1}{2}\int_{\epsilon^2}^{\infty}\frac{ds}{s}\operatorname{Tr} e^{-s\Delta}$, regulated by a proper-time cutoff $s \geq \epsilon^2$ with ϵ a short-distance length. The a_1 term contributes precisely the Einstein–Hilbert piece, and matching it to $-\frac{1}{16\pi G_N}\int\sqrt{g_4}\,R$ gives, per minimal scalar,

$$\frac{1}{G_{\text{ind}}} = \frac{1}{12\pi\,\epsilon^2}, \qquad \Longrightarrow \qquad \frac{1}{4G_{\text{ind}}} = \frac{1}{48\pi\,\epsilon^2}. \tag{791}$$

It sets the entanglement entropy. The leading (area) term of the entanglement entropy across the horizon, computed from the *same* heat kernel by the conical method of the next subsection, is

$$\frac{dS_{\text{ent}}}{dA} = \frac{1}{48\pi\,\epsilon^2} \qquad \text{(per minimal scalar)}. \tag{792}$$

The match. Comparing (791) and (792),

$$\boxed{\frac{dS_{\text{ent}}}{dA} = \frac{1}{4G_{\text{ind}}}} \qquad \text{identically.} \tag{793}$$

The individual number $1/48\pi\epsilon^2$ is scheme-dependent: it moves when you change the regulator ϵ or the coupling ξ, and it would change again if you added more fields. But the *equality* (793) does not move, because both sides are literally the same heat-kernel coefficient a_1 of the same modes. The ultraviolet-sensitive prefactor cancels in the ratio. This is why the induced-gravity reading is so powerful: it ties the divergent entanglement entropy and the divergent induced $1/G_N$ together so that their ratio is finite and universal, even though neither is. In GG-Theory the cancellation is not a happy accident to be checked regulator by regulator — it is enforced by GGEP, since there is no bare G_N for the divergence to be a correction to. What we still need is the actual value of the ratio. That is what the conical replica delivers, in closed form, with no surviving dependence on ϵ or on the field content.

N.6 The conical-replica (deficit-angle) method: the coefficient $1/4 = 4\pi/16\pi$, regulator-independently

Entanglement entropy is computed by the replica trick: $S = -\operatorname{Tr}\rho\ln\rho = (1 - n\,\partial_n)\ln\operatorname{Tr}\rho^n\big|_{n=1}$, and $\operatorname{Tr}\rho^n = Z_n/Z_1^n$, where Z_n is the Euclidean partition function on an n-fold cover of the geometry, branched over the horizon. Recall from §N.2 that making the period of Euclidean time anything other than $2\pi/\kappa$ leaves a conical singularity on the horizon. The replica cover $\mathcal{M}_n$ does exactly that: it has Euclidean-time period n times the smooth one, i.e. a *conical deficit angle* $2\pi(1-n)$ concentrated on the bifurcation surface Σ (the horizon, of area A). The point of the method is that all the ultraviolet detail of the fields is now packaged in a known distributional curvature.

The curvature of a manifold with a conical deficit acquires a piece localized on the tip — a two-dimensional delta function with weight set by the deficit angle. Integrated against the volume, this gives the exact identity

$$\int_{\mathcal{M}_n}\sqrt{g}\,R = n\int_{\text{smooth}}\sqrt{g}\,R \;+\; 4\pi(1-n)\,A. \tag{794}$$

The first term is the bulk curvature, n copies of the smooth one; the second is the new horizon term, and its coefficient 4π is pure geometry — the angular measure of the deficit times the codimension-two horizon area. No field content, no regulator, no coupling ξ enters (794); it is a statement about cones.

Now feed this into the *induced* Einstein–Hilbert action $I_E[\mathcal{M}_n] = -\frac{1}{16\pi G_N}\int_{\mathcal{M}_n}\sqrt{g}\,R$, which by §N.4 is the *entire* gravitational action (there is no separate matter contribution to the entropy beyond the modes that generated G_N):

$$I_E[\mathcal{M}_n] = n\,I_E[\mathcal{M}_1] \;-\; \frac{(1-n)\,A}{4G_N}, \qquad \ln Z_n = -I_E[\mathcal{M}_n]. \tag{795}$$

Writing $R_1 \equiv \int_{\mathcal{M}_1}\sqrt{g}\,R$ for the smooth bulk curvature, the replica entropy follows by one differentiation:

$$\begin{aligned} \ln Z_n &= \frac{n\,R_1}{16\pi G_N} + \frac{(1-n)A}{4G_N}, & \partial_n \ln Z_n\big|_{n=1} &= \frac{R_1}{16\pi G_N} - \frac{A}{4G_N}, \\ \ln Z_1 &= \frac{R_1}{16\pi G_N}, & S = \ln Z_1 - \partial_n \ln Z_n\big|_{n=1} &= \boxed{\frac{A}{4G_N}}. \end{aligned} \tag{796}$$

The bulk curvature R_1 cancels identically between $\ln Z_1$ and $\partial_n \ln Z_n$; the entropy is *purely* the boundary (area) term. And the coefficient is now transparent and closed-form:

$$\frac{1}{4} \;=\; \frac{4\pi}{16\pi} \;=\; \frac{\text{(conical-deficit curvature coefficient)}}{\text{(Einstein–Hilbert normalization)}}. \tag{797}$$

The 4π in the numerator is the deficit-angle curvature weight of (794); the 16π in the denominator is the normalization of the Einstein–Hilbert action. The famous one quarter is nothing more mysterious

than their ratio. Restoring units,

$$\boxed{S_{\rm BH} = \frac{k_B\, c^3\, A}{4\, G_N\, \hbar}} \tag{798}$$

which reproduces (789), now *derived* from induced gravity rather than read off the integrated first law. Three features deserve emphasis because they are what make this a derivation and not a fit:

1. *No free parameter.* Nothing was tuned; $4\pi/16\pi$ is forced.
2. *Regulator-independent.* The cutoff ϵ and the coupling ξ canceled in the ratio (793), and the conical identity (794) never referred to them.
3. *No supersymmetry and no extremality.* Unlike the Strominger–Vafa brane-counting derivations, which work at (or near) an extremal, supersymmetric limit, the conical-replica argument runs for the generic Schwarzschild, Kerr, and Reissner–Nordström hole.

N.7 What standard physics leaves open: the origin of $A/4$

The derivations of Step 1 are among the deepest results in theoretical physics, but it is worth being precise about their status. That a black hole carries an entropy proportional to its horizon area is forced by the thermodynamic analogy and by Hawking radiation; that the coefficient is *exactly* 1/4 in Planck units emerges from the heat-kernel and conical-replica calculations. But in the standard setting these calculations run on an abstract manifold, with the Newton constant G_N inserted as a measured input. They do not say *why* G_N has its value, nor tie the area law to a specific underlying geometry: the 1/4 is recovered, but from ingredients that are themselves taken as given. The geometric account supplies exactly those ingredients.

N.8 The geometric derivation: $A/4$ from the canonical bulk (Arisaka (2026, C3-BH))

Everything in §N.4–§N.6 is, in the cohort, a named theorem rather than a sketch — Theorem *T-BH-1* of the companion Arisaka (2026, C3-BH), the black-hole entropy paper, graded there Derived-conditional on three named premises — and GG-A10-Cosmos cites it at that grade rather than re-proving it. Two points make the GG-Theory version sharper than the textbook induced-gravity argument.

The conical replica runs on the canonical six-dimensional bulk. Arisaka (2026, C3-BH) does not perform the heat-kernel reduction on an abstract four-manifold; it performs it on the canonical bulk $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$ of Arisaka (2026, A2-GG6), with the metric $ds^2_6 = e^{-2r}\eta_{\mu\nu}dx^\mu dx^\nu + dr^2 + d\theta^2$. Reducing the six-dimensional Einstein–Hilbert term on the scale fiber I_r and the phase circle S^1_θ produces the boundary Einstein–Hilbert term with a *finite, parameter-free* coefficient $1/G_N = \pi(1 - e^{-2K})/G_6$, the radial integral converging because the scale profile e^{-2r} is integrable. The induced Newton constant is thus a property of the scale profile, ultraviolet-localized at the GUT edge $r = 0$ — the geometric statement that gravity is a UV phenomenon.

The induced G_N is A5's scale-overlap G_N, by construction. This is the cross-reference that closes the loop and that Arisaka (2026, C3-BH) now carries explicitly. The boundary general

relativity of Arisaka (2026, A5-GDOS) reads its Newton constant off the *same* canonical scale profile through the scale-overlap reduction $M_{\rm Pl}^2 = M_6^4\, 2\pi R_\theta \int e^{-2A}\, dr$; the induced-gravity derivation of Arisaka (2026, C3-BH) reads $1/G_N$ off that *same* scale profile of Arisaka (2026, A2-GG6). The two Newton constants therefore coincide *by construction*, not by a coincidence to be checked numerically: the G_N that enters horizon entropy is the scale-derived one, the very same object that fixes the strength of gravity in the Friedmann equations of the main text. And the coefficient $1/4$ is the more robust half of the claim — by (797) it is a ratio of curvature normalizations, so even if the absolute normalization of G_N were later refined, the one-quarter would not move.

In the hidden-fiber language of GG-A8-QM, Step 4 of Arisaka (2026, C3-BH) even names the microstates: the modes whose one-loop determinant induces G_N are the Wilson-line holonomies on the phase circle S^1_θ together with the Kaluza–Klein tower of the scale direction, evaluated on the horizon two-sphere. A horizon of area A supports $\exp(A/4\ell_P^2)$ distinguishable fiber configurations — the same compact circle that carries quantum coherence in the bulk becomes, at a horizon, the carrier of black-hole microstates. The black-hole counting and the origin of quantum mechanics are one mechanism evaluated on two loci.

N.9 Information and the Page curve: conceptually closed, quantitatively open

Here the honesty discipline must be exact, because this is where the literature and the cohort both fall short of a full computation, and the main text of GG-A10-Cosmos flags it as the least-closed module of its black-hole reading.

The information paradox is the tension between two facts. A pure state $|\psi\rangle$ collapses to a black hole, which then evaporates into Hawking radiation that is thermal — hence mixed — to all orders in the semiclassical expansion, so the map $|\psi\rangle \to \rho_{\rm thermal}$ appears to be pure-to-mixed and non-unitary. Page sharpened the demand in 1993: if the total evolution is unitary, the entanglement entropy of the radiation must first rise, then turn over at the halfway point and *fall* back to zero — the "Page curve" — rather than rise monotonically as Hawking's calculation predicts. Recovering that turnover is the modern content of "resolving the paradox."

In GG-Theory the resolution is stated structurally. By OGGEP (AX-3) observation is the open-system trace of the bulk over what a boundary horizon hides, so the radiation entropy is a branch-minimized generalized-entropy functional — mathematically the same object as the island/replica generalized entropy of the contemporary literature. The cohort names the turnover "the Page curve as a GRIP/GDOS saddle switch": at the Page time the operationally dominant saddle of the open-system entropy ledger changes, which is *why* the curve must turn over. GG-A10-Cosmos's own structural estimate of the turnover mass, $M_{\rm Page} = M_0/\sqrt{2}$, follows from the half-entropy condition.

But — and this is the open part, said plainly — naming the saddle is not the same as computing it. The slogan "Page curve as a saddle switch" identifies *which physical thing changes* at the turnover; it does *not* perform a replica/island calculation with a quantitative Page time, an exhibited island region, and a quantitative entropy curve. The cohort tells you *why* a branch-minimized functional had to be the answer; the island program tells you *how* to compute the saddle, and that computation has not been carried out here. The honest quantitative status is identical to that of the companion

Arisaka (2026, C3-BH): the $A/4$ coefficient and the thermality (T-BH-1) are Derived-conditional on three named premises, but the island/replica Page curve is treated as conceptually closed and quantitatively open. We do not pretend the slogan performs the calculation.

N.10 A worked numerical example: a one-solar-mass Schwarzschild black hole

It is worth seeing the numbers, because the answer is one of the largest entropies in physics. Take $M = M_\odot = 1.989\times10^{30}$ kg and the constants $G_N = 6.674\times10^{-11}\ \mathrm{m^3\,kg^{-1}\,s^{-2}}$, $c = 2.998\times10^{8}\ \mathrm{m\,s^{-1}}$, $\hbar = 1.055\times10^{-34}\ \mathrm{J\,s}$, $k_B = 1.381\times10^{-23}\ \mathrm{J\,K^{-1}}$.

Step 1 — Schwarzschild radius and area.

$$r_s = \frac{2G_N M}{c^2} \approx 2.95\times10^{3}\ \mathrm{m}\ \ (\approx 2.95\ \mathrm{km}), \qquad A = \frac{16\pi G_N^2 M^2}{c^4} \approx 1.10\times10^{8}\ \mathrm{m^2}. \tag{799}$$

Step 2 — the Planck area.

$$\ell_P^2 = \frac{\hbar G_N}{c^3} \approx 2.61\times10^{-70}\ \mathrm{m^2}, \qquad \ell_P \approx 1.62\times10^{-35}\ \mathrm{m}. \tag{800}$$

Step 3 — the entropy in nats. The dimensionless entropy (entropy in units of k_B, i.e. in nats) is one quarter of the horizon area in Planck cells:

$$\frac{S_{\mathrm{BH}}}{k_B} = \frac{A}{4\ell_P^2} \approx \frac{1.10\times10^{8}}{4\,(2.61\times10^{-70})} \approx 1.05\times10^{77}\ \text{nats}. \tag{801}$$

Step 4 — in bits and in SI units. Converting nats to bits with $1\ \text{nat} = 1/\ln 2 \approx 1.443$ bits,

$$\frac{S_{\mathrm{BH}}}{k_B \ln 2} \approx 1.51\times10^{77}\ \text{bits}, \qquad S_{\mathrm{BH}} = k_B\cdot 1.05\times10^{77} \approx 1.45\times10^{54}\ \mathrm{J\,K^{-1}}. \tag{802}$$

For orientation: the corresponding Hawking temperature from (786) is $T_H \approx 6.2\times10^{-8}$ K, far below the CMB, so this hole does not evaporate today. The entropy $\sim 10^{77}$ nats is roughly twenty orders of magnitude larger than the thermodynamic entropy of the Sun as an ordinary ball of gas ($\sim 10^{58}\,k_B$) — the gravitational collapse into a black hole is an enormous entropy increase, exactly as the generalized second law requires. Every one of those $\sim 10^{77}$ nats is, in the GG-Theory reading of §N.8, a configuration of the hidden phase fiber traced out by the horizon.

Status

The area law $S = A/4$ is Derived-conditional, at the companion's own grade. The coefficient $1/4 = 4\pi/16\pi$ is fixed in closed form by the conical-replica method on the canonical six-dimensional bulk, regulator-independently and without supersymmetry or extremality, in Theorem *T-BH-1* of Arisaka (2026, C3-BH); the induced Newton constant entering it is the scale-overlap G_N of Arisaka (2026, A5-GDOS), read off the same canonical bulk of Arisaka (2026, A2-GG6), so the two coincide by construction. The companion holds T-BH-1 at Derived-conditional on three named premises, with the species closure registered as its kill row F-BH-9, and GG-A10-Cosmos inherits it at that grade and no higher. *The Page-curve recovery is conceptually closed but quantitatively open.* The

open-system OGGEP/GRIP reading names the saddle switch that makes the radiation entropy turn over, and identifies the entropy functional as the island generalized entropy; it does *not* provide a replica/island computation with a quantitative Page time. This is the honest residue, flagged as such in the main text and unchanged here: the mechanism is named, the calculation is not yet performed.

Problems and worked solutions. N.1.1. Verify the first law for the Schwarzschild family. *Solution.* Step 1 (geometrized units): $A = 16\pi G^2 M^2$, so $dA = 32\pi G^2 M\, dM$. Step 2: with $\kappa = 1/4GM$: $\frac{\kappa}{8\pi G}\, dA = \frac{1}{4GM} \cdot \frac{32\pi G^2 M}{8\pi G}\, dM = dM$ — the first law $dM = \kappa\, dA/8\pi G$, exactly, with no adjustable coefficient. Step 3: read with $T_H = \hbar\kappa/2\pi$ this forces $S = A/4G\hbar$ up to an additive constant — the two famous formulas are one Legendre pair, which is why Hawking's temperature computation instantly fixed Bekenstein's coefficient. Physics intuition: the first law is the hinge between mechanics and thermodynamics — once any one intensive quantity (here T_H) is pinned by microphysics, the whole thermodynamic dictionary locks.

N.5.1. Show that the heat-kernel coefficient cancels in $S = A/4G_N$: the Susskind–Uglum compensation.
Solution. Step 1: the same a_1 (curvature) coefficient of the heat-kernel expansion controls both the induced Newton constant, $1/G_{\text{ind}} \propto \sum_i \chi_i \Lambda^2$, and the leading entanglement entropy across the horizon, $S_{\text{ent}} \propto \sum_i \chi_i \Lambda^2 A$. Step 2: forming the ratio, the species sum and the cutoff *cancel*: $S_{\text{ent}} = A/4G_{\text{ind}}$ — whatever the field content, however many species, the entropy in units of the *induced* $1/4G$ is universal. Step 3: this is why the species problem (the apparent field-content dependence of horizon entropy) dissolves exactly when gravity is induced — the entropy and the stiffness are the same loop diagram read twice. Physics intuition: one integral, two names — the compensation is the deepest single fact in the induced-gravity reading, and the cohort's remaining species-sum work (the black-hole companion's declared open computation) is about the *absolute* normalization, never about this ratio.

N.6.1. Run the conical-replica derivative: the deficit angle delivers $A/4G$.
Solution. Step 1: the n-replica Euclidean manifold is smooth except at the horizon, where it carries a conical deficit $\delta = 2\pi(1-n)$; the curvature acquires a distributional term $R \supset 2\delta\, \delta^{(2)}(x)$. Step 2: inserting into the Einstein–Hilbert action, $I_n \supset -\frac{1}{16\pi G} \cdot 2 \cdot 2\pi(1-n)\, A$, so the replica free energy carries $(1-n)\, A/4G$. Step 3: the entropy is $S = -\partial_n \big[I_n - nI_1\big]_{n=1} = A/4G$ — the coefficient $\frac{1}{4}$ falls out of the geometry of the cone, independent of the matter content (N.5.1's cancellation seen from the geometric side). Physics intuition: entropy is the action's response to a pinch at the horizon — every modern derivation of $A/4$, from Gibbons–Hawking to the islands literature to the cohort's canonical-bulk version, is this one computation in different clothing.

N.9.1. Put numbers on the Page schedule for a solar-mass black hole.
Solution. Step 1: the evaporation time is $t_{\text{evap}} \approx 2 \times 10^{67}$ yr on the photon-only blackbody clock, prefactor $5120\pi G^2/\hbar c^4 = 8.4 \times 10^{-26}\,\text{s}\,\text{g}^{-3}$; Page's rate for all massless species, the prefactor 6.2×10^{-27} of Eq. (333), gives 1.6×10^{66} yr — more species evaporate the hole faster, and the two clocks differ by a factor 13.6, the species content and the greybody factors together rather

than a species count alone, and the entropy-based turnover sits at $t_{\rm Page} = (1 - 2^{-3/2})\, t_{\rm evap} = 0.646\, t_{\rm evap} \approx 1.4 \times 10^{67}$ yr. Step 2: the archive to be returned is $S/k_B \ln 2 \approx 1.5 \times 10^{77}$ bits; spread over the lifetime this is an average of $\sim$230 bits per second. Step 3: the register, kept exactly as the chapter states it: the Page *curve* is conceptually closed (unitarity plus the port bookkeeping fix its shape), while the quantitative microstate accounting behind it remains the least-closed module of the black-hole reading — the numbers above are schedule, not mechanism. Physics intuition: even for the most patient object in physics the ledger balances — but at two hundred bits a second across 10^{57} ages of the universe, "information is preserved" is a statement about principle, never about practice.

N.10.1. When does a black hole begin to evaporate at all?
Solution. Step 1: net evaporation requires $T_H > T_{\rm CMB} = 2.725$ K. Step 2: with $T_H \propto 1/M$ (the chapter's worked example gives 6×10^{-8} K for $1\, M_\odot$), equality sits at $M = 4.5 \times 10^{22}$ kg ≈ 0.6 lunar masses. Step 3: every astrophysical black hole is therefore *absorbing* today; the evaporation era for a solar-mass hole begins only when the expanding universe (Appendix F) has cooled the sky below sixty billionths of a kelvin — the de Sitter future of F.7.1, where the only radiation left is the horizon's own. Physics intuition: black-hole thermodynamics runs on the universe's schedule, not its own — the cosmos must first become colder than its coldest objects, and the π-locked vacuum sets exactly how cold that future gets.

Take-away. *Standard result:* horizons carry entropy $A/4G_N$ and temperature $\hbar\kappa/2\pi$ — exact, universal, and unexplained. *What it assumed:* the coefficient 1/4 from a semiclassical computation with no statement of *whose* degrees of freedom are counted. *What the geometry replaced it with:* induced gravity on the canonical bulk — one heat-kernel coefficient fixing both $1/G_N$ and the entropy, the horizon read as an OGGEP port, and the information story conceptually closed with the quantitative Page-curve match named as open work (PS-Cosmos-18). *Continue in:* Part VII; GG-C3-BH.

N.11 Worked derivation notebook for the Part VII observables

N.11.1 Derivation notebook A: Schwarzschild area, entropy, and temperature

Start from the Schwarzschild radius

$$r_s = 2G_N M. \tag{803}$$

The horizon area is then

$$A_H = 4\pi r_s^2 = 16\pi G_N^2 M^2. \tag{804}$$

Dividing by $4G_N$ gives

$$S_{\rm BH} = \frac{A_H}{4G_N} = 4\pi G_N M^2. \tag{805}$$

Differentiating,

$$\frac{{\rm d}S_{\rm BH}}{{\rm d}M} = 8\pi G_N M. \tag{806}$$

Using the first law $\mathrm{d}M = T_H\,\mathrm{d}S_{\mathrm{BH}}$ gives

$$T_H = \left(\frac{\mathrm{d}S_{\mathrm{BH}}}{\mathrm{d}M}\right)^{-1} = \frac{1}{8\pi G_N M}. \tag{807}$$

Nothing exotic has happened. The notebook begins by showing that the entire paradox is built on two ordinary geometric facts: area scales as M^2, temperature as M^{-1}.

N.11.2 Derivation notebook B: evaporation law and Page-time estimate

Assume

$$\dot{M} = -\frac{\alpha_H}{M^2}. \tag{808}$$

Then

$$M^2\dot{M} = -\alpha_H \qquad \Rightarrow \qquad \frac{\mathrm{d}}{\mathrm{d}t}\left(\frac{M^3}{3}\right) = -\alpha_H. \tag{809}$$

Integrating,

$$M^3(t) = M_0^3 - 3\alpha_H t. \tag{810}$$

The evaporation time follows from $M(t_{\mathrm{evap}}) = 0$:

$$t_{\mathrm{evap}} = \frac{M_0^3}{3\alpha_H}. \tag{811}$$

For the Page estimate, equate branch 1 and branch 2 with $\eta = 1$:

$$S_0 - S(M_{\mathrm{Page}}) = S(M_{\mathrm{Page}}) \quad \Rightarrow \quad S(M_{\mathrm{Page}}) = \frac{S_0}{2}. \tag{812}$$

Since $S \propto M^2$,

$$M_{\mathrm{Page}} = \frac{M_0}{\sqrt{2}}. \tag{813}$$

Insert into the cubic mass law:

$$\left(\frac{M_0}{\sqrt{2}}\right)^3 = M_0^3 - 3\alpha_H t_{\mathrm{Page}}, \tag{814}$$

so

$$t_{\mathrm{Page}} = \frac{M_0^3}{3\alpha_H}\left(1 - 2^{-3/2}\right) = \left(1 - 2^{-3/2}\right) t_{\mathrm{evap}}. \tag{815}$$

The point is not the exact prefactor 0.65, but the algebraic fact that a Page-like turnover is compatible with the evaporation law once the relevant entropy branch is no longer taken to be strictly thermal Hawking entropy forever.

N.11.3 Derivation notebook C: the operational map and why pure-to-mixed is not paradoxical

Let the full 6D state evolve unitarily:

$$\rho_6(t) = U_6(t)\rho_6(0)U_6^\dagger(t). \tag{816}$$

Now split the full algebra into accessible and inaccessible parts only at the operational level. The accessible boundary state is

$$\rho_\partial(t) = \mathrm{Tr}_{\mathrm{bulk}\setminus\partial}\,\rho_6(t). \tag{817}$$

Equivalently, it can be written in Kraus form,

$$\rho_\partial(t) = \sum_k K_k(t)\rho_6(0)K_k^\dagger(t), \qquad \sum_k K_k^\dagger K_k = \mathbb{1}. \tag{818}$$

A pure state under a partial trace can become mixed with no contradiction whatsoever. The contradiction only arises if one demands that ρ_∂ itself be the total ontology. OGGEP forbids that demand.

N.11.4 Derivation notebook D: horizon mass as a timestamp for PBH formation

Use the radiation-era relations

$$\rho = \frac{\pi^2}{30} g_* T^4, \qquad H = 1.66 g_*^{1/2} \frac{T^2}{M_{\mathrm{Pl}}}. \tag{819}$$

Then

$$M_H = \frac{4\pi}{3}\rho H^{-3} \tag{820}$$

$$= \frac{4\pi}{3}\left(\frac{\pi^2}{30} g_* T^4\right)\left(1.66^3 g_*^{3/2}\frac{T^6}{M_{\mathrm{Pl}}^3}\right)^{-1}. \tag{821}$$

Collecting factors gives the scaling

$$M_H \propto g_*^{-1/2} T^{-2}. \tag{822}$$

Evaluating numerically at $T = \mathrm{TeV}$ gives

$$M_H(\mathrm{TeV}) \sim 10^{26}\,\mathrm{g}, \tag{823}$$

up to order-unity factors. This is the cleanest reason PBH mass windows are so informative. They are direct thermometers of the early universe.

N.11.5 Derivation notebook E: why the minimal PBH abundance is exponentially tiny

Take the asymptotic Gaussian tail

$$\beta(M) \approx \frac{\sigma}{\delta_c\sqrt{2\pi}} \exp\left[-\frac{\delta_c^2}{2\sigma^2}\right]. \tag{824}$$

Set $\delta_c = 0.45$. If $\sigma = 0.01$, then

$$\frac{\delta_c^2}{2\sigma^2} = \frac{0.2025}{2\times 10^{-4}} \approx 1012.5, \tag{825}$$

hence

$$\beta \sim e^{-1012.5}. \tag{826}$$

Even increasing σ to 0.05 gives

$$\frac{\delta_c^2}{2\sigma^2} = \frac{0.2025}{0.005} \approx 40.5, \tag{827}$$

so

$$\beta \sim e^{-40.5} \sim 2.6 \times 10^{-18}, \tag{828}$$

again tiny. The lesson is elementary and brutal: PBH production is a tail phenomenon. A smooth red spectrum will not do it.

N.11.6 Derivation notebook F: evaporation survival threshold

Take the standard evaporation-time estimate

$$t_{\rm evap} \approx 6.2 \times 10^{-27}\,{\rm s} \left(\frac{M}{\rm g}\right)^3. \tag{829}$$

Set this equal to the age of the universe, $t_0 \sim 4 \times 10^{17}\,{\rm s}$. Then

$$\left(\frac{M}{\rm g}\right)^3 \sim \frac{4 \times 10^{17}}{6.2 \times 10^{-27}} \sim 6 \times 10^{43}, \tag{830}$$

so

$$M \sim \left(6 \times 10^{43}\right)^{1/3} {\rm g} \sim {\rm few} \times 10^{14}\,{\rm g}. \tag{831}$$

This is the origin of the usual survival threshold in Eq. (332).

N.11.7 Derivation notebook G: tensor amplitude from the plateau height

Use Eq. (338):

$$r = \frac{2V_*}{3\pi^2 A_s M_{\rm Pl,red}^4}. \tag{832}$$

Take $V_*^{1/4} = 8.0 \times 10^{15}\,{\rm GeV}$ — the pivot value on the derived plateau, one percent under its height $V_0^{1/4} = 8.1 \times 10^{15}\,{\rm GeV}$, not the lock scale; the TeV reading was retired in Part IX and the notebook is worked at the height the theory actually predicts — together with $A_s = 2.1 \times 10^{-9}$ and $M_{\rm Pl,red} = 2.435 \times 10^{18}\,{\rm GeV}$. Then

$$\frac{V_*}{M_{\rm Pl,red}^4} = \left(\frac{8.0 \times 10^{15}}{2.435 \times 10^{18}}\right)^4 \sim 1.2 \times 10^{-10}. \tag{833}$$

Multiplying by the prefactor $2/(3\pi^2 A_s) \sim 3 \times 10^7$ gives

$$r \simeq 0.004, \tag{834}$$

in agreement with the leading-order $12/N^{*2} = 0.0043$, and with the exact 0.0038 to the same precision. This is why the tensor target is not a matter of taste. It is straight arithmetic. (At the retired TeV height this notebook would give $r \sim 10^{-52}$; that reading is recorded in Part IX's record, not here.)

N.11.8 Derivation notebook H: today's stochastic background

For an approximately scale-invariant primordial tensor spectrum,

$$\Omega_{\mathrm{GW},0}h^2 \sim \frac{\Omega_{\mathrm{r}}h^2}{24}\, A_s r \simeq 1.7 \times 10^{-6} A_s r, \tag{835}$$

with $\Omega_{\mathrm{r}}h^2 \simeq 4.1 \times 10^{-5}$. Insert $A_s = 2.1 \times 10^{-9}$ and $r = 0.0038$:

$$\Omega_{\mathrm{GW},0}h^2 \sim 1.7 \times 10^{-6} \times 2.1 \times 10^{-9} \times 0.0038 \sim 1.4 \times 10^{-17}. \tag{836}$$

This value keeps the primordial background far below interferometric reach — the tensor target is a CMB-polarization signal, not an interferometer signal — in the minimal module.

N.11.9 Final notebook takeaway

The worked notebook shows that the most dramatic Part VII claims are also the cleanest algebraically. The black-hole entropy ledger follows from area plus entanglement. The Page time follows from a branch crossing. PBH masses follow from the horizon mass at formation. PBH abundances are exponentially sensitive to the variance and therefore negligible for a smooth red spectrum. Primordial tensors follow directly from the plateau height, and it is the height that decides them: at the derived plateau the notebook returns $r \simeq 0.004$, and at the retired TeV reading it would have returned $r \sim 10^{-52}$. One number, two histories, and only one of them survives Part IX.

This is exactly what one wants from a unification paper. The arguments are not hidden in numerical machinery. The core leverage is visible on the page.

O Internal source ledger for GG-A10-Cosmos

This paper is intentionally restricted to the internal GG-2026/GD-2026 preprint chain. The derivation, notation, and theorem flow used in GG-A10-Cosmos are taken only from the seven source preprints listed below. Historical literature is cited for orientation, not as a derivational premise.

Table 61: Internal source ledger used in GG-A10-Cosmos. The columns give the preprint, its role in the chain, and its relevance here. The ledger exists so that any number in this paper can be traced to the paper that owns it in a single step, and so that a reader can see how little of the machinery is native to this one.

Preprint	Role in the chain	Relevance
GG-A1-Constitution	Constitutional foundation	Fixes the four axioms, seven postulates, OGGEP semantics, and admissibility rules that every cosmological derivation must obey.
GG-A2-GG6	Six-dimensional stage	Supplies the GG-6 bulk, BI-6, OGN-5, compact holonomy, and the six-force unification from which all later cosmic structure descends.
GG-A3-BI6	Gauge-consistency completion	Proves local and global anomaly freedom, unified gauge consistency, and dual-current conservation, thereby certifying the legality of the cosmic bulk-to-boundary channel.
GG-A5-GDOS	GDOS foundation	Derives the exact boundary reduction, the open-system current geometry, and the formal state-manifold language on which cosmic observables are defined.
GG-A6-GRIP	Runtime constitution	Establishes GRIP, $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$, Wilson-loop GSSB, and the SIP packet that organize all cosmic phase transitions in this paper.
GG-A7-Qi4	Qi-runtime completion	Supplies the Qi-4 pillars, SeOS, QCI, IH + OH = EIC, and the five-story pagoda used in the higher-floor interpretation of cosmic history.
GG-A9-SM	Particle-to-cosmos bridge	Fixes the Standard-Model ledger, the compact-angle sector, the neutrino block, baryonic protection, and the parameter bridge inherited by the cosmology paper.

On reproducibility. The empirical spine of this paper — the CLASS demonstration of the microwave spectrum from the derived parameters (Section 40), the parameter-audit deviation plot and its data table, and the χ^2 computation against the binned Planck TT data — rests on scripts and data files released as a public, runnable repository accompanying this preprint: the CLASS input decks and derived-parameter settings, the binned Planck TT data with sources, the χ^2 and χ^2/dof routines (`gg_cmb_tt_v4.py`, `peaks_v4.py`, and the ℓ-stretch test `cmb_stretch_v1.py`, each with its recorded output), and the deviation-plot generator with its tabulated inputs. A cosmological claim that hands its data confrontation to an unreleased script asks for a trust the field no longer extends; we intend to earn it the ordinary way, by letting the reader run the pipeline. The repository DOI is listed with this paper's record.

P Cross-paper inheritance table (compact)

This appendix compresses the cross-paper coherence map of Section 84.1 into a one-page inheritance table for quick reference. For the full eight paragraphs with citations and theorem references, see Section 84; for the textual outward and inward maps, see Section 84.2.

Table 62: Compact cross-paper inheritance: what GG-A10-Cosmos draws from each upstream paper of the cohort. "Used by" lists the A10 sections that depend on each upstream contribution. The last column lists the sections of this paper that consume each inheritance, which makes the table usable in both directions: forwards to see what an upstream result is used for, and backwards to see what a given section rests on.

Paper		Key contribution to A10	Used by (A10 sections)
GG-A1-Constitution	May 2026	4 axioms (LC/GGEP/OGGEP/MSMF) + 7 postulates + admissibility predicate	§68, §68.2.1
GG-A2-GG6	May 2026	6D bulk (Thm 1); OGN-5 tuple (Thm 6); compact S^1_θ seeding Kaxion	§74.1, §42.4 (H1)
GG-A3-BI6	May 2026	BI-6 anomaly-cancellation theorems (local + global); coefficient vector $(1, -1, -2/5, -2/5)$; four-form construction	§42.4 (H2), §74.1, framing remark
GG-A4-Euclid	May 2026	Zero-free-parameters discipline; OGN-5 as theorem of BI-6; integers $J = 7$ and $K = 35$	§68.2.1, §42.4 (C1), §68.1
GG-A5-GDOS	May 2026	GDOS open-system formalism; OGGEP projection structure	§68, all cosmological-observable-as-export claims
GG-A6-GRIP	May 2026	Four Gi pillars (TGG / CCI / GFF / GRIP); GRIP template (DGI → GSSB → GA); cascade closure T4.1 ($N_g \leq 3$, conditional on M10)	§42.4 (H3), §68.2.1 (cost functional)
GG-A7-Qi4	May 2026	Four Qi pillars; five-story pagoda; OH + IH = EIC closure	§71.4 (universal-identity preview)
GG-A9-SM	May 2026	Codebook Provenance Table template; D/D*/A classification; "Why every electron is the same" section	§68.1, §71

Q Consolidated benchmark audit

The main body develops the benchmark problems sector by sector. Table 63 compresses the same logic into a single operational audit sheet.

Table 63: Consolidated benchmark audit for GG-A10-Cosmos. The audit is consolidated rather than complete: it gathers the benchmarks that are decidable now, and a reader wanting the full inventory should read the codebook, where the rows that cannot yet be decided are also listed and graded.

Problem	Inherited mechanism	Principal kill switch
Inflation	The Σ_r modulus on the GDOS plateau one e-fold below the GUT edge, $N^* = 52.6$	A tensor measurement off $r = 0.0038$ beyond the launch's spread, or a tilt off 0.9634 at three standard deviations (met today on one of two combinations)
Dark energy	Residual late-time GFF vacuum inventory with $w = -1$ in the minimal module	Secure detection of $w \neq -1$ or significant dark-energy clustering

Problem	Inherited mechanism	Principal kill switch
Dark matter	Kaxion/compact-angle relic sector inherited from GG-6 and 6F; one scale calibrated	Exclusion of the calibrated window at DFSZ depth without a replacement descending from the chain
Baryogenesis	Holonomy-assisted leptogenesis through the inherited neutrino block	A decisive neutrino-sector measurement incompatible with the required CP corridor
BBN	Standard thermal history plus a narrow $A = 7$ port correction	Failure to fit deuterium/helium while preserving the same inherited ledger
CMB baseline	Derived sound horizon, baryon loading, damping, and projection structure	Inability to reproduce the observed acoustic-scale hierarchy
Hubble tension	One background rate derived with zero sky input; the distance ladder the predicted outlier	Joint BAO/CMB/SN fit excluding the derived background rate
S_8 tension	Standard growth on the π-locked background; S_8 computed from the declared amplitude, the tension shared	A cosmic-shear S_8 converging below 0.80 at $\geq 3\sigma$ with the CMB value unchanged, which the computed 0.825 cannot follow
JWST early galaxies	One baryonic lever, the collapse threshold; the small-scale power reallocation and the streaming damping withdrawn	High-redshift abundance statistics incompatible with the derived enhancement
Black-hole information	Open-system horizon packet, GFF ledger, and GRIP saddle switching	Demonstrable contradiction with Page-curve recovery or generalized entropy
PBHs	Narrow feature-dependent supplementary channel, not a required module	Positive signal in an excluded mass window or non-detection in a sharply predicted window
Primordial GWs	Vacuum tensors at $r = 0.0038$, $\Omega_{\mathrm{GW},0}h^2 \sim 10^{-16}$; no transition-sourced channel claimed	Either a null at the 10^{-3} floor or a detection above 10^{-2}
Human place / pagoda	Higher-floor emergence via Qi-runtime and EIC closure	Inability to connect higher-floor reconstruction with the inherited runtime grammar

R Named-theorem-and-conjecture deck

This appendix lists every formal named statement (Principle, Theorem, Conjecture, Corollary) introduced in GG-A10-Cosmos, together with its label and a one-line summary. Full statements appear in the body sections referenced.

Table 64: Named-statement deck of GG-A10-Cosmos. Refs go to the body section where the full statement appears. The references point to the body section where each statement is given in full, because a named statement compressed to one line is a label and not a proof, and nothing in this deck should be cited without following the reference first.

Type	No.	Label	One-line summary
Principle	—	`princ:msmf-ogn-cosmos` (§68.2.1)	The cost functional $\mathcal{F}$ admits a unique cost-minimum saddle on every admissible cosmological sub-stratum.
Theorem	4.D.1	`thm:hubble-two-stage-corridor` (§68.3)	Existence of one background solution with a GRIP-Out switch at $z_* = e^7$ separating the two dictionaries through which it is read; the ladder's 73.5 is not a projection of it.
Theorem	4.E.1	`thm:omega-m-pi-topological` (§74.3)	Structural form of the closed-packet $\Omega_m = 1/\pi$ claim (T-Cosmos-17); programmatic — does not yet close (Conjecture VIII).
Conjecture	I	`conj:kaxion-completeness` (§68.2.3)	Kaxion misalignment-mechanism relic abundance equals $1/\pi - 1/(2\pi^2)$ at the angle Part V reads back, not a cost-minimum.
Conjecture	II	`conj:hubble-corridor` (§68.2.3)	Stratum assignment of the vacuum residue: the one background rate of Theorem 4.D.1 is unique on the cosmological-boundary stratum, $\zeta_\Lambda = \pi/(LK)$.
Conjecture	III	`conj:compact-angle-mass` (§68.2.3)	The OGN-5 ledger fixes $f_{\text{Kaxion}} = f_{\text{bulk}}/\Xi_{\text{cap}}$ with $\Xi_{\text{cap}} = K+\Sigma$; at present $f_{\text{bulk}} = 3.04\times 10^{13}$ GeV is calibrated ($f_{\text{Kaxion}} = 7.1 \times 10^{11}$ GeV, $m_{\text{Kaxion}} = 8.1\ \mu$eV); open remainder: the orientation convention behind the calibration (GG-A3-BI6 O14) and a defining computation for the ladder.
Conjecture	IV	`conj:slow-roll-exit` (§68.2.4)	(Revised) The GDOS route forces the Starobinsky/ξ plateau class, and the lock crossing is prompt ($N_2 = 0$); underwrites the derived n_s and $V_{\text{inf}}^{1/4}$.

continued on next page

Type	No.	Label	One-line summary
Conjecture	V′	`conj:gdos-amplitude` (§68.2.4)	The GDOS heat-kernel R^2 coefficient pays the amplitude, $A_s = 21.5\, e^{-2K/3}$ (replaces the retired, refuted TeV-uniqueness conjecture).
Conjecture	VI	Resonant-leptogenesis-normalization (§68.2.3)	The Boltzmann-transport normalization of the $K = 35$ baryogenesis throughput converges to a unique coefficient within the order-of-magnitude transport band; closes the η_B codebook row.
Conjecture	VII	`bfft-damping-knee` (§68.2.3)	The CMB damping knee is the first non-trivial BFFT ultraviolet cutoff of the multiscale throat, $\ell_D = K(2\pi)^2$, carried as an overlay on the scored Boltzmann value 1350; closes no row (D18 is closed by Conjecture VIII).
Conjecture	VIII	`closed-packet-integration` (§68.2.3)	The closed-packet integral of the BI-6 four-form on the cosmological-boundary slice evaluates to $1/\pi$ and $1/(2\pi^2)$ with an exhibited normalization (Steps 3–4 of Theorem 74.1). Status: VIII-a proved, VIII-b′ proved, VIII-b″ discharged into the occupancy identity $u = \pi\Omega_m = 1$, whose epoch is a declared calibration (§74.3); proof promotes D1–D3 (and D18, D21) to D with the epoch input derived elsewhere.
Conjecture	IX	`vacuum-residual-exactness` (§68.2.3)	The late-time residual vacuum export satisfies $w = -1$ exactly; closes the w row.
Conjecture	X	`conj:lithium-port` (§68.2.3)	The $A = 7$ BBN port inherits the closed-packet $1/\pi$ suppression, so $(^7\mathrm{Li}/H)_{\mathrm{GG}} = \mathrm{SBBN}/\pi$; closes the $^7\mathrm{Li}/H$ row (D20).
Corollary	1	`cor:hubble-tension-reconciliation` (§68.3.5)	The 6.4σ Hubble tension is read as one derived background plus a predicted ladder systematic (Theorem 4.D.1).
Corollary	2	`cor:s8-reconciliation` (§68.3.5)	One computed amplitude, $S_8 = 0.825$, against two comparators; the tension is shared, not resolved.
Corollary	3	`cor:hubble-falsification` (§68.3.5)	A ladder-free late probe (sirens, time delays) returning the ladder's 73.5 rather than the background 66.8 would falsify the one-background reading of Theorem 4.D.1.

Inventory. 1 Principle + 2 Theorems + 10 Conjectures + 3 Corollaries = 16 formal named statements.

S Named-lemma promotion roadmap

This appendix states the open research targets that, if proved, would promote codebook entries from Derived (cond.) or Asserted to Derived. Each entry gives the named conjecture, the rows it closes, the proof venue (this paper, a companion paper, or open research), and the estimated proof complexity.

Table 65: Named-lemma promotion roadmap for the MSMF/OGN audit of GG-A10-Cosmos. Proof venue "this paper" = work that belongs in the present paper and remains to be completed; "C1" = companion paper Arisaka-GG-C1-Kaxion; "open" = open research target without a designated venue.

No.	Named conjecture	Rows closed (codebook)	Proof venue	Status
I	Kaxion-misalignment-completeness, with the pinning C-Kaxion-d	D4 ($\Omega_{\rm DM}$)	Open; Arisaka-GG-C1-Kaxion for the pinning	open research target; the pinning necessary under the isocurvature bound, itself a conjecture
II	Stratum-assignment of the vacuum residue	D5, D6 (H_0 against two comparators)	Open	open research target (GG-A5-GDOS O6)
III	Compact-angle-mass-uniqueness	D13, D14 ($m_{\rm Kaxion}$, $f_{\rm Kaxion}$)	Arisaka-GG-C1-Kaxion	companion paper
IV	Plateau-class-and-prompt-crossing (revised)	A7, D9, A15 (n_s, r, $V_{\rm inf}^{1/4}$)	Open	open research target
V′	GDOS-heat-kernel-amplitude (replaces the retired Conj. V)	A8 (A_s); A15 ($V_{\rm inf}^{1/4}$, with IV); D10, D11, D19 (σ_8, S_8)	Open	open research target
VI	Resonant-leptogenesis-normalization	D17 (η_B)	Open	open research target
VII	BFFT-damping-knee-identification	BFFT overlay of D18 (ℓ_D; row scored by the Boltzmann inventory value)	Open	open research target
VIII	Closed-packet-integration	D1–D3 (Ω_m, Ω_Λ, Ω_b); D18, D21 (ℓ_D, z_*) via the fractions	this paper / open	open research target
IX	Vacuum-residual-exactness	D16 (w)	Open	open research target
X	Lithium-port-π-suppression	D20 ($^7{\rm Li}/H$)	Open	open research target

Closure-economy summary. One codebook entry (D12, the derived null $\Delta N_{\rm eff} \leq 0.03$) is already at Derived status. Nineteen are Derived (conditional) on Conjectures I–IV, V′ and VI–X

— D9 (r) among them, on Conjecture IV with the tilt — and one is Asserted pending Conjecture V′. Once all ten conjectures are proved, the codebook collapses to 21 D + 0 D* + 0 A (twenty-one entries, about eight of them independent, the rest the same geometry read in other units). The proof venue closest to the present paper is Conjecture VIII (closed-packet-integration), whose programmatic sketch appears as Theorem 74.1 and whose proof would close D1–D3 (and, through the fractions, D18 and D21; with the epoch input a calibration, all five stay D*). The remaining conjectures are open research targets, with Conjecture III having a designated companion-paper venue (Arisaka-GG-C1-Kaxion) and the remainder awaiting external proof or further structural analysis.

After all, the promotion roadmap is the place where the codebook stops being a table of provenance labels and starts being a research program. Ten precise mathematical conjectures, nine of them open research targets and one (Conjecture III) carried in the dark-matter companion, where it is open in two named places (Lemma III), and a clean closure path from 1/19/1 to 21/0/0.

T Falsification-test catalog

This appendix collects, in one place, every observational or theoretical test whose failure would falsify the framework or one of its named lemmas. Each row gives the test, the precision required, the falsification target, and the named lemma or theorem it would directly affect.

The framework's exposure profile is therefore high: a clean failure on any one of the eight tests of Table 66 would refute a named conjecture (II, III, IV or V′) or the one-background theorem (4.D.1). Most of the tests sit within the next-decade observational horizon; the same eight tests are laid out on a single forward time axis in the kill-test timeline of Figure 36 (Part XIII, §95.1).

With that, the course is over. The reader who entered these appendices with one exposure to the Friedmann equations has now built the standard cosmology from scratch and, at every station, seen exactly where the GG-6 geometry replaces an input with a theorem — and where it does not yet. There is no better preparation for what comes next: the Objections Part, where everything taught here is turned, deliberately, against the framework itself.

Table 66: Falsification-test catalog for GG-A10-Cosmos. Each test is paired with the named conjecture or theorem it would refute and the experimental venue that could carry out the test. The columns give the test, the precision it requires, the named conjecture or theorem it would refute, and the experimental venue that could run it. A test whose venue column names no existing or funded instrument is a weaker entry than one that does, and the catalog is not sorted to hide that.

Test	Precision	Falsification target	Experimental venue
Corridor switch vs atomic decoupling	± 5 at $z = 1090$ (0.46%)	$z_{\rm dec}$ (atomic chain) = 1089–1092, zero parameters; $z_{\rm sw} = e^7$ a distinct topological marker 0.0038 e-folds above (0.38% in $1 + z$). Measured $z_{\rm dec}$ off the atomic-chain value by $> 1.4\%$ — three standard deviations of the proposed ± 5 measurement; a 0.5% threshold would fire on a 1.1σ fluctuation — (present agreement 0.15%) would refute the reading; the switch–decoupling offset $\delta = 0.0038$ is a derived consistency statement, not an independent kill (corridor theorem, property C1; §42.1)	CMB spectral-distortion (recombination-radiation) missions, post-2030, + 21-cm cosmology
BAO dilation	flat at the background r_d, no two-stage residue at any z	A measured residue at the percent level counts against the one-background reading	DESI, Euclid, Roman/-Nancy Grace Roman
GW siren H_0	1% on the background	a converged $H_0^{\rm GW} \approx 73$ (tracking the local ladder, not the background ≈ 67) would refute the π-inventory derivation	LIGO-Voyager, Einstein Telescope, Cosmic Explorer
Lensing time-delay H_0	1%	A converged determination at the ladder's 73 rather than the background 66.8 refutes the one-background reading; convergence on ≈ 67 confirms it	H0LiCOW, TDCOSMO
Tensor-to-scalar ratio	measurement vs $r = 0.0038$ (exact; $12/N^{*2} = 0.0043$ at leading order)	A tensor amplitude inconsistent with the derived target beyond the launch's spread 0.0037–0.0040, in either direction, falsifies the launch	LiteBIRD, Simons Observatory
Kaxion haloscope window	DFSZ-depth, count-aware scan of 1.59–2.13 GHz (one line, alone)	Exclusion of the window at DFSZ depth retires the calibrated placement and its ladder; a second in-window line refutes the count	ADMX, CAPP, HAYSTAC, ADMX-EFR
Primordial GW background	B-mode polarization	The primordial bump must appear at $r = 0.0038$ within the launch's spread 0.0037–0.0040, and in any case inside $[10^{-3}, 10^{-2}]$ (D9); a supercooled-transition GW background would separately refute the prompt crossing	LiteBIRD, Simons Observatory
σ_8 precision growth	1% at $z = 0.5$–1	A combined determination excluding the derived $S_8 = 0.825$ at 3σ, with the CMB-normalized value also moving away from it, refutes the amplitude chain from A_s (Lemma V′)	Euclid weak-lensing, LSST

Part XVI
Common Objections and Responses

The strongest objections a ΛCDM and CMB-precision cosmologist, an inflation theorist, a Hubble-tension and large-scale-structure specialist, a dark-matter, baryogenesis, and BBN phenomenologist, a gravitation and black-hole relativist, a formal quantum-field-theory and string theorist, and a methodologist and philosopher of cosmology would each raise after reading the whole cohort — answered, with pointers to the theorem, codebook status, conjecture, open problem, or companion paper that bears the weight

---

A program claiming that the large-scale inventory of the universe — the matter, dark-energy, and baryon fractions, the recombination redshift, and the Hubble background — is fixed by the same zero-dial six-dimensional geometry that fixes the elementary particles, and carrying that claim through an explicit Derived / Derived-conditional / Asserted codebook, invites the most demanding scrutiny; it would not be honest scholarship to leave the strongest objections unstated. We collect here the adversarial readings that a ΛCDM and CMB-precision cosmologist, an inflation theorist, a Hubble-tension and large-scale-structure specialist, a dark-matter, baryogenesis, and BBN phenomenologist, a gravitation and black-hole relativist, a formal quantum-field-theory and string theorist, and a methodologist and philosopher of cosmology would each raise after reading not just this paper but the whole twenty-two-paper cohort — and we answer each, with explicit pointers to the theorem (by its cohort-global id, §D), the codebook status, the named conjecture, the open problem, or the companion paper that carries the weight. These are deliberately the *surviving* objections: many first-look complaints dissolve once the companion that does the work is read — the boundary general relativity of GG-A5-GDOS (Arisaka, 2026, A5-GDOS), the Bekenstein–Hawking coefficient of GG-C3-BH (Arisaka, 2026, C3-BH), the Kaxion relic of GG-C1-Kaxion (Arisaka, 2026, C1-Kaxion), the anomaly ledger of GG-A3-BI6 (Arisaka, 2026, A3-BI6), and the six-dimensional action of GG-A2-GG6 (Arisaka, 2026, A2-GG6). Where the honest answer is "not yet" — above all the closed-packet integration that would turn the π-locked inventory from a structural ansatz into a theorem (Conjecture VIII), and the full Planck-likelihood fit that would turn the present central-value comparisons into a quantitative goodness-of-fit — we say so and name the open item.

The Part opens, however, not with the specialists but with the reader — the thoughtful non-specialist who arrives carrying the questions popular science has made urgent: the multiverse, the landscape, parallel worlds, fine-tuning, what came before the beginning, whether a forced universe leaves room for us. These are answered first (§O1), in plain language and without jargon, because a framework that claims the cosmos is *necessary* rather than *accidental* owes its clearest account to exactly the reader who has most often been told the opposite. The technical critics follow.

O1 For the curious reader: the big questions, plainly answered

The objections that follow are not the technical challenges of a specialist critic — those come afterwards, sector by sector. They are the questions a thoughtful non-specialist actually carries into a cosmology book: about the multiverse, the landscape, parallel worlds, fine-tuning, what came before the Big Bang, whether the theory leaves room for us. Popular science has made some of these questions feel settled in directions this framework does not take, so we meet them first, and plainly. Where a claim rests on a still-open step we say so; the point of this section is not to oversell, but to clear the mind of a reader who has been told, often, that the universe is probably one of infinitely many accidents.

Objection 1 (O-Cosmos-1) (the curious reader)
"Isn't our universe just one bubble in a vast multiverse? I keep reading that ours is one of infinitely many, and we happen to live in the one that works."

Response. That is one influential story, but it is a hypothesis, not a finding — and it is the one this framework is built to make unnecessary. The multiverse is invoked to explain why our universe's numbers are life-friendly: if every possible universe exists, then *of course* we find ourselves in a hospitable one, and no further explanation is owed. But notice what that story costs. It explains the numbers by giving up on deriving them: they are an accident of our address in an unobservable ensemble. This paper takes the opposite route. It argues that the numbers are not accidents to be excused but *outputs to be computed*: the matter fraction $1/\pi$, the dark-energy fraction $1 - 1/\pi$, the switch redshift of recombination e^7, each read off a single six-dimensional geometry with no dial to turn (Parts IV, X). If that program succeeds, there is nothing for a multiverse to do, because there was never a menu of alternatives to select from. The honest status: several of these values are still *Derived-conditional* on a named, stated conjecture (Part IX) — so we have not finished the job. But the wager is clear, and it is testable this decade, which no multiverse claim is.

Objection 2 (O-Cosmos-2) (the curious reader)
"String theory says there are 10^{500} possible universes — the 'landscape.' Doesn't that make picking ours hopeless, and physics basically a lottery?"

Response. The landscape is the sharpest form of the worry, and worth taking seriously. The concern is that a fundamental theory might admit an astronomical number of self-consistent vacua, with no principle selecting ours — so the constants of nature become a draw, and the theory loses its power to predict. This framework's central bet is a single word against that picture: *uniqueness*. It proposes that the constitutional axioms (GG-A1-Constitution) plus the six-dimensional anomaly and integer ledgers (GG-A3-BI6, GG-A4-Euclid) leave *one* self-consistent geometry, not 10^{500} — so there is no landscape to be lost in. Whether that uniqueness is fully proved is exactly the boundary we mark honestly: the internal consistency is shown, but the theorem that *no other* admissible geometry exists remains a stated conjecture rather than a closed result (it is named as such in the discussions). So we do not claim to have refuted the landscape by fiat. We claim something checkable instead: a theory with zero *continuous* free

parameters (an overall amplitude, the observation epoch and the Kaxion's one calibrated scale are declared, not tuned) cannot retreat to a landscape when the data press, because it has no knobs to re-tune. If its one geometry gives the wrong sky, it is simply wrong. That exposure is the antidote to the lottery.

Objection 3 (O-Cosmos-3) (the curious reader)
"What about parallel worlds — the idea from quantum mechanics that every possibility happens in a branching reality?"

Response. This is a different idea from the multiverse, and it is worth keeping them apart. The "many worlds" of quantum mechanics is a proposal about what happens during a measurement — that the wavefunction never collapses, and every outcome is realized on a separate branch. It is a question about the *interpretation of quantum theory*, not about cosmological initial conditions, and this paper does not need to settle it to do its work. What the framework does offer is a different reading of measurement itself: an observation is an *open-system boundary projection* (the OGGEP axiom of GG-A1-Constitution, developed as GDOS in GG-A5-GDOS) — a finite observer reading a boundary of a closed higher-dimensional geometry, never the whole bulk. In that reading the appearance of "branching" is about what a bounded observer can and cannot access, not about the literal proliferation of worlds. The companion paper GG-A8-QM develops the operational quantum mechanics; here we simply note that nothing in this cosmology requires an infinity of parallel copies of you. One geometry, read from inside, is enough.

Objection 4 (O-Cosmos-4) (the curious reader)
"The universe looks incredibly fine-tuned for life. If that's not luck and not a multiverse, what is it?"

Response. It is, in this framework, neither luck nor selection but *necessity* — and that is a genuine third option, not a rhetorical dodge. The fine-tuning argument has always had exactly two escape routes: a designer who set the dials, or a multiverse in which every dial-setting exists and we see a workable one. Both take for granted that the dials *could* have been otherwise. This paper denies the premise. If the life-permitting constants — the electron's mass, the strength of the forces, the tiny value of the vacuum energy — are forced outputs of one geometry with no adjustable parameter, then there was never a dial to tune, and so nothing to tune it luckily or select it anthropically (Part XIV makes this argument in full). A universe whose contents are *forced* is life-permitting for the same reason it is anything at all: because its own geometry admits exactly one specification. The fine-tuning does not disappear; it is *relocated* — from a coincidence in need of an explainer to a theorem in need of a proof. And a theorem, unlike a coincidence, can be checked.

Objection 5 (O-Cosmos-5) (the curious reader)
"Doesn't 'necessity' just replace a divine designer with a mathematical one? You've moved the mystery, not solved it."

Response. A fair challenge, and the difference is worth stating precisely. Yes, the framework relocates the mystery — from "why these particular values?" to "why this particular geometry,

and why does mathematics describe the world at all?" It does not pretend to answer the second question; no physical theory can, and we do not claim otherwise. But the relocation is not empty, because it changes what kind of thing the unexplained residue *is*. A designer who chose the constants must be explained in turn (who designed the designer, and by what freedom?), and the regress never closes. A geometry that admits one self-consistent solution has no such regress: it is not *chosen*, it simply is the unique thing that hangs together, the way there is only one consistent arithmetic. The remaining mystery — why there is a lawful world rather than nothing — is real and we leave it standing, plainly labeled. What we have removed is the *intermediate* mystery: the twenty-odd apparently arbitrary numbers between the axioms and the sky. Removing those is not nothing; it is the whole content of the Euclidean ideal (GG-A4-Euclid).

Objection 6 (O-Cosmos-6) (the curious reader)
"What came before the Big Bang? Every account either dodges this or invokes a prior universe."

Response. This framework gives an unusually concrete answer, and it is neither a dodge nor a prior universe. In the standard picture the Big Bang is a *singularity* — a moment where the equations break down and "before" has no meaning — and that breakdown is genuinely uncomfortable. Here the earliest physically meaningful moment is not a singularity but a *finite geometric edge*: the ultraviolet boundary of the six-dimensional throat, at the grand-unification energy $E_{\mathrm{GUT}} \approx 2.3 \times 10^{16}$ GeV (Part III). Cosmic history is the descent from that edge, downhill, to the gentle vacuum we inhabit today. So "what came before?" is answered by "the edge is the beginning of the descent, and there is no singular instant behind it to explain" — the Planck-scale infinity that plagues other accounts is derived to be a finite, lawful boundary, not a wall the theory crashes into. What the framework does *not* claim is a complete quantum account of the edge itself; that is named as open work. But it replaces "the equations break down" with "the descent starts here," which is a real and testable difference, not a philosophical evasion.

Objection 7 (O-Cosmos-7) (the curious reader)
"If inflation blew up a tiny patch, why is the whole sky made of the same stuff? Two regions that were never in contact somehow have identical electrons."

Response. You have put your finger on a question the textbooks skip, and it is one this paper is proud to answer (Part X, §72). The famous "horizon problem" asks why distant regions share the same *temperature*, and inflation answers it: they were once one small patch, stretched huge. But those regions also share the same *things* — every electron with the same mass and charge, the same dark matter, the same vacuum energy — and temperature-smoothing cannot explain that, because no exchange of heat can negotiate a particle's identity. The framework's answer has two parts. First, there is only one "mold": every electron is the readout of the same unique geometric ground state, so there is nothing regional to vary (Part X). Second — and this is the part that needed the new result that inflation sits at the *grand-unification* scale — the stretch happened *before any matter existed*. The universe was blown up while it was still empty;

the electrons were minted *afterwards*, everywhere at once, from that one mold. So the sky is uniform in its contents not because distant regions negotiated an agreement, but because they inherited one, from a time before there was anything to disagree about. That, we argue, is the deeper horizon problem — and its solution.

Objection 8 (O-Cosmos-8) (the curious reader)
"Why is the universe 95% dark? It seems absurd that we can only see 5% of everything."

Response. It feels absurd only if you expect the visible to be the whole; the framework offers a reason it is not. In this reading the cosmic inventory is a closed set of geometric fractions — roughly 5% ordinary matter, 27% dark matter, 68% dark energy — and the ordinary, luminous part is exactly the *one* slice that carries ordinary electric charge and therefore interacts with light (Part XIV). The Standard-Model ledger we inherited from the particle paper (GG-A9-SM) can illuminate only that slice; the rest is not missing or exotic but simply the part of the inventory that light was never built to touch. Dark matter and dark energy are dark for the same reason a clean pane of glass is invisible: not because they are absent, but because the instrument sent to find them — electromagnetism — passes straight through. On this account the striking number is not that 95% is dark, but that the visible fraction is a specific small share, $1/2\pi^2 \approx 5\%$, of a closed geometric budget. The darkness is the shape of what light was built to miss.

Objection 9 (O-Cosmos-9) (the curious reader)
"Is the universe deterministic, then? If everything is forced by geometry, is my future already written — do I have any freedom?"

Response. The framework is careful here, and the distinction it draws is the one that matters to you. What the geometry forces is the *layered structure* of the world — that there are stable atoms, that there is chemistry, that there is a dark-energy era — and the *values* of the constants that make those layers possible. It does *not* force every event within the layers. Which atoms combine into which molecule, which star forms where, which thought you think next — these live in the enormous combinatorial freedom *inside* the layers the geometry has built (Part XIV). The determinism is at the level of the stage and the cast of characters, not the play. This is a genuine third position between the two uncomfortable options usually offered: not rigid clockwork (in which your choices are illusions) and not pure randomness (in which they are noise), but a forced scaffolding within which contingent history — including yours — actually unfolds. You are not preprogrammed by the geometry; you are *sheltered* by it. The universe fixes the rules of the game, not its moves.

Objection 10 (O-Cosmos-10) (the curious reader)
"Isn't it suspicious that π shows up in the amount of matter in the universe? That smells like numerology."

Response. It should smell like numerology — until you see where the π comes from, and we insist on that distinction because the paper's whole method depends on it. Numerology is picking a pretty formula *after* seeing the number: with a few constants and enough operations

you can hit almost any target, which proves nothing. The claim here is the reverse. The π is not chosen to match 0.318; it is forced to appear because the matter fraction is computed as the integral of a specific geometric quantity (a curvature four-form) over a closed surface, and integrating curvature over a closed surface *produces factors of* π by a theorem of geometry, not by choice (Parts X, and the topological appendix). The same single computation is required to give matter $1/\pi$, dark energy $1 - 1/\pi$, and baryons $1/2\pi^2$ *together* — one topology, three coupled numbers, no freedom to adjust one without the others. A numerologist fits one number at a time; a derivation forces several at once and cannot wriggle. The honest caveat: the integral is written down but not yet fully evaluated in closed form (Conjecture VIII), so we carry these as "derived-conditional," not as finished theorems. But the π is a consequence of the shape, not an ornament chosen to impress.

Objection 11 (O-Cosmos-11) (the curious reader)
"Every generation's physicists think they're near a 'theory of everything.' Why is this any different from the last confident framework that failed?"

Response. Healthy skepticism, and the framework asks to be judged by exactly the standard your question implies: not confidence, but exposure. The programs that disappointed — grand unification, supersymmetry, the string landscape in its predictive ambitions — shared a trait: as data pressed, they could retreat, adding parameters, raising scales, widening allowed ranges. This paper deliberately removes that escape. It has *zero continuous free parameters* (GG-A4-Euclid): no dial to turn when a measurement comes in wrong. It stakes specific numbers on specific instruments of the coming decade — a tensor-to-scalar ratio of 0.0038, a dark-matter particle at 8.1 microelectronvolts, a gravitational-wave Hubble constant near 67 — any one of which, landing wrong, refutes it (Part XIII). We make no claim that it is *right*; we claim it is the *kind* of framework that can be shown wrong, cleanly and soon. That is precisely what the confident-but-unfalsifiable frameworks lacked. Judge it not by its ambition but by whether the next decade's data are allowed to kill it. They are.

Objection 12 (O-Cosmos-12) (the curious reader)
"The 'anthropic principle' says the universe must allow observers because we're here observing it. Isn't that the real explanation, and isn't your 'necessity' just anthropics in disguise?"

Response. They are opposite kinds of explanation, and the difference is the crux. The anthropic principle explains a life-friendly constant by *observer selection*: the value could have been anything, but only in universes where it is right can anyone be around to notice, so we necessarily observe a workable value. It presupposes a range of possibilities and prunes it by our existence. The present framework denies the range. If a constant is a forced output of one geometry, then it could *not* have been anything else, and observer selection has nothing to select among — there is one value, whether or not anyone is here to see it. The test that separates the two is decisive: anthropic reasoning predicts constants scattered anywhere in the life-permitting *window*, typically near its edge; a necessity claim predicts *specific* values, at the center of nothing in particular, fixed before any observer existed. This paper predicts specific

values — $1/\pi$, e^7, 0.0038 — and stakes them on measurement. If they were merely anthropically allowed, they would not need to be exactly those. Anthropics explains why we *could* be here; the framework tries to explain why the universe is *this*, observer or no.

Objection 13 (O-Cosmos-13) (the curious reader)
"You say inflation happened at a colossal energy near grand unification. How can anyone take seriously a claim about energies a trillion times beyond any collider?"

Response. We cannot build a collider that reaches E_{GUT} — true, and no one claims otherwise. But the sky is an instrument the laboratory is not. Inflation at the grand-unification scale leaves fossils we can read today: the exact tilt of the primordial ripples, $n_s = 0.963$, already measured, and today at the edge of what the measurements allow; and a specific amount of primordial gravitational waves, $r = 0.0038$, which the next generation of microwave-background telescopes (the Simons Observatory, LiteBIRD) is being built to detect (Parts IV, VII). The energy is inaccessible; its *consequences* are not. This is how cosmology has always worked — we cannot rerun the first three minutes, yet the helium abundance they left behind is measured to percent precision. The framework's inflationary claim is testable for the same reason: it makes a two-sided, decisive prediction for a number that a funded experiment will measure this decade. If the tensor ratio comes back far from 0.0038 in either direction, the claim is dead — inaccessible energy or not.

Objection 14 (O-Cosmos-14) (the curious reader)
"Dark energy is the biggest embarrassment in physics — the vacuum energy is predicted wrong by 120 orders of magnitude. How can any theory claim to have this under control?"

Response. The 120-order discrepancy is real and it is the sharpest test any dark-energy claim must face, so it is fair to press it. The standard estimate takes the vacuum energy to be set by the highest scale in physics, the Planck scale, and finds a number $\sim 10^{120}$ times too big — the worst prediction in the history of science. The framework's move is to deny that the vacuum energy is a Planck-scale quantity at all. In this reading the observed dark energy is a *residual* left at the very bottom of the geometric descent — a seesaw between the lock scale and the Planck scale, $m_\Lambda = M_{\mathrm{lock}}^2/M_{\mathrm{Pl}} \approx 6$ meV (Parts IV, and the companion GG-A5-GDOS) — which recovers about 119 of the 120 orders *without* a fine-tuned cancellation. That is not yet a complete theorem: the final factor, and the exact fraction $1 - 1/\pi$, rest on named conjectures we mark openly. But the structural claim is specific and falsifiable — $w = -1$ exactly, no evolution — and it is decidable by DESI, Euclid, and Rubin this decade. The embarrassment is not waved away; it is attacked at its root, the assumption that empty space must weigh the Planck scale.

Objection 15 (O-Cosmos-15) (the curious reader)
"The 'cosmic coincidence' — that matter and dark energy happen to be comparable right now, in our era — seems too convenient. Doesn't that single out our moment suspiciously?"

Response. It is a real puzzle, and the framework splits it into a part it dissolves and a part

it honestly keeps. Under a free cosmological constant, matter (which dilutes as the universe expands) and dark energy (which does not) are comparable for only a brief window in cosmic history — and our living in that window looks like a coincidence begging for an anthropic explanation. The framework changes the picture in two steps. First, the matter and dark-energy fractions are not independent quantities that happen to be close; they are locked together by one geometry as $1/\pi$ and $1 - 1/\pi$, which are comparable *by construction*, at every epoch you compute them, not just now (Part X). Second — the honest remainder — there is still a mild question of *why now* we read the snapshot, and the paper carries this as a single, bounded, few-percent "timing" input rather than pretending it away (Part IV's "why now" section). So the suspicious coincidence shrinks to a small, named, calendar-level fact, with the large structural coincidence removed. Not everything is dissolved — but what remains is stated in the open, not smuggled into an anthropic hand-wave.

Objection 16 (O-Cosmos-16) (the curious reader)
"Extra dimensions sound like science fiction. Six-dimensional geometry, a 'bulk,' a 'throat' — why should I believe any of this is real rather than mathematical decoration?"

Response. The right test of a hidden structure is not whether we can picture it but whether it *pays its way* in predictions, and that is the standard the framework accepts. The two extra dimensions here are not extra *places* you could travel to; they are a *scale* direction (how zoomed-in you are, the energy at which you probe) and a *phase* direction (an internal angle, like the setting of a dial). Physics already uses both ideas without calling them dimensions — the running of forces with energy is the scale direction; the phases of quantum fields are the angle. The framework's wager is that treating them geometrically is not decoration but *economy*: the same two directions that organize the particle masses in GG-A9-SM also fix the cosmic inventory here, with no separate assumptions. A decoration explains nothing new; these dimensions are asked to explain the electron and the dark-matter fraction with one structure. Whether they are "real" is decided the way we decide anything in physics: the six-dimensional picture makes specific, falsifiable predictions (the 3.85-TeV particles, the 8.1-microelectronvolt axion), and if those are absent, the picture fails. It earns reality by exposure, not by vividness.

Objection 17 (O-Cosmos-17) (the curious reader)
"Why is there something rather than nothing? Does your framework even touch the deepest question?"

Response. Honestly, no — and it is important not to pretend otherwise. Why there is a lawful world at all, rather than sheer nothing, is a question physics cannot reach, because any physical answer must already assume *some* law to do the explaining. This framework does not claim the summit of that question, and any theory that did would be overreaching. What it does is narrower and, we think, still worth a great deal: it shrinks the gap between "nothing" and "everything." Instead of a world specified by a long list of arbitrary numbers — each of which is its own little "why this and not that?" — it offers a world specified by a single geometry that admits one self-consistent solution. The list of separate mysteries collapses toward one. You are still left with the final "why is there a lawful anything?", and we hand it to you undiminished

and clearly labeled. But you are no longer *also* owed an explanation for why the electron weighs what it weighs, or why the universe is a third matter and two-thirds dark. Those become consequences. Moving the mystery from many places to one is not the whole answer — but in physics, it is most of the work.

Objection 18 (O-Cosmos-18) (the curious reader)
"This all sounds like a 'theory of everything.' Aren't those famously arrogant and famously wrong?"

Response. The label is one the framework should decline, and does. A "theory of everything" usually means a single equation from which *all* of physics is supposed to flow. This is not that, and it does not claim to be. It is explicit about what it does *not* compute: it fixes the *inputs* of physics — the constants, the inventory — but it hands the detailed dynamics (the precise shape of the microwave spectrum, the nonlinear clumping of galaxies) to the standard computational machinery, which it frames but does not replace (Part XIII's "tier ladder"). It says, in effect: here is why the constants are what they are; the hard calculational physics you already do is still yours to do. That is a deliberately bounded claim. The arrogance of the failed "theories of everything" was to promise the whole of physics and deliver an unfalsifiable framework; the discipline here is the opposite — a modest target (the constants, not the dynamics) staked on falsifiable numbers. If anything the framework errs toward saying too clearly what it *cannot* do. That is the correct direction for the error to run.

Objection 19 (O-Cosmos-19) (the curious reader)
"You keep saying certain results are 'conjectures' or 'not yet proved.' Doesn't that mean the whole edifice could collapse?"

Response. It means the edifice is honest about its load-bearing walls, and you should want that, not fear it. Every result in this cohort carries an explicit grade: *Derived* (proved end to end), *Derived-conditional* (proved except for one named, stated mathematical step), or *Asserted* (an input we have not yet derived). Roughly, of the cosmological quantities, one is fully derived — a null, $\Delta N_{\rm eff} \leq 0.03$ — nineteen are derived-conditional on a short list of named lemmas, and one — the amplitude of the primordial ripples — is still an input (Part IX). A collapse is possible in exactly the places we have marked: if the central open conjecture (the closed-packet integration behind the π-fractions) turns out false, those fractions fall. But this is the normal condition of live science, made unusually visible. Most frameworks bury their assumptions; this one prints a ledger of them and invites you to watch which ones close. A structure that tells you precisely which stone would bring it down is sturdier, epistemically, than one that hides its stones — because you can *test* it. The conjectures are not weaknesses concealed; they are the next decade's homework, listed.

Objection 20 (O-Cosmos-20) (the curious reader)
"Where do we — life, minds, observers — fit into this? Or is consciousness just an afterthought bolted onto the physics?"

Response. In this framework observers are not an afterthought and not the point either; they

are a *late rung of the same ladder* that fixes the electron. The claim, developed fully in the capstone companion GG-A11-Origin and only pointed to here, is that the geometric descent does not stop at the vacuum: the same principle that selects the stable electron selects the stable proton, then the stable atom, then the molecules of chemistry, then — where conditions allow — the self-sustaining systems we call life, and eventually the observers who write the equations down (Part XIV's "one staircase"). On this reading we are neither the reason the universe has its numbers (it had them for billions of years before us) nor an accident decorating a cosmos that has nothing to do with us. We are a *continuation* — a late step on a staircase whose first steps are the particles. The framework does not solve the hard problem of consciousness, and it does not pretend to. But it refuses to treat mind as foreign to physics: if the argument holds, the same geometry that fixed the weight of empty space also, further down, made the systems that could come to know it. That is the trilogy this paper sits in the middle of — and the reader who wants the rest is pointed, deliberately, downstream to GG-A11-Origin.

Objection 21 (O-Cosmos-21) (the curious reader)
"Isn't it circular for the universe to 'explain itself' — for physics written by humans to claim the humans were inevitable?"

Response. There is a loop here, but it is a feature the framework names openly rather than a fallacy it hides. The situation is genuinely self-referential: a universe produces, late in its history, observers who reconstruct the constitution that produced them. Older traditions noticed this shape — Wheeler's "participatory universe," the Upanishadic recognition of self in reality, Spinoza's necessity — and the paper cites them without borrowing their metaphysics (Part XIV). Why is it not vicious circularity? Because the framework does not use our existence as a *premise* to derive the constants — that would be circular. It derives the constants from geometry with no reference to observers at all, and then *notices*, as a consequence, that among the things the geometry forces is the eventual appearance of systems complex enough to do the deriving. The logic runs geometry → constants → atoms → life → minds → (the writing of the theory), in one direction, with no step reaching back to justify an earlier one. The loop closes at the end; it is not assumed at the start. That the universe turns out to contain its own describers is not a trick of the argument — it is, if the argument is right, a fact about the universe.

Objection 22 (O-Cosmos-22) (the curious reader)
"Suppose you're wrong. What would it actually take to convince you — or is this, like so many grand theories, effectively unfalsifiable?"

Response. This is the question we most want asked, because the framework's whole claim to seriousness rests on the answer being short and concrete. It would take any *one* of the following, all within the coming decade: a primordial gravitational-wave signal with tensor ratio well away from 0.0038 in either direction (beyond the launch's own spread, 0.0037–0.0040); a dark-matter axion excluded across the 1.6–2.1 gigahertz band at the predicted coupling; a gravitational-wave-siren Hubble constant that tracks the local value (≈ 73) rather than the predicted background (≈ 67); a confirmed evolving dark energy ($w \neq -1$); or — on the particle side of the cohort — a single proton decay, which the geometry forbids absolutely. Any one

of these, cleanly established, ends the framework; there is no parameter to adjust in response, because there are no *continuous* free parameters — only three declared inputs (an amplitude, the observation epoch and one calibrated scale) and a list of named lemmas, none of them a knob you could turn to absorb a wrong measurement (Part XIII, and the kill-test timeline). That is the difference between this and an unfalsifiable grand theory: we have told you, in advance and in public, exactly which observations would force us to abandon it, and the instruments that will make those observations are already running or funded. A theory that can hand you the list of its own executioners is not hiding. Whether it survives them is now a question for the sky, not for us.

O2 The ΛCDM and CMB-precision critic

Objection 23 (O-Cosmos-23) (the CMB-precision critic)
"You set $z_* = e^7 = 1096.6$ against the Planck recombination redshift $z_{\rm rec} = 1089.92 \pm 0.25$. That is a 27σ miss in units of the Planck error, yet you call it '0.6% agreement.' Which is it — a triumph or a failure?"

Response. You are right that we cannot have it both ways, and we now pick one identity cleanly. $z_* = e^7$ is *not* claimed to be the physical recombination redshift, and it is not compared to $z_{\rm rec} = 1089.92 \pm 0.25$ as a prediction of that measured number — to do so, and then reach for the visibility width to soften the 27σ, would be precisely the comparator-shopping you object to. Instead z_* is a *distinct object*: the *topological switch redshift* of the one background, the integer-locked scale e^J at which the dictionary that reads it changes. The operational decoupling redshift $z_{\rm dec} = 1089.92$ that Planck measures is a different quantity — the peak of the visibility function — and the two coincide to 0.6%. That 0.6% offset is not an error bar to be gamed; it is now a *computed* quantity: the framework's own atomic chain gives $\ln(1 + z_{\rm dec}) = 6.996$, hence $\delta = J - \ln(1 + z_{\rm dec}) = 0.0038$ (§42.1, Eq. (262)), the codebook carrying both quantities in one entry (D21) under Conjecture VIII (closed-packet integration). So the honest statement carries no σ: a single ledger integer $J = 7$ fixes a topological corridor-switch that falls 0.6% from the operational recombination peak, with the offset computed — and the atomic quantity itself landing on the Planck value at the 0.1% level with zero fitted parameters — not a precision match asserted and not a tension dressed in a wide comparator. The Dev plot scores the audited atomic-chain prediction, with the switch annotated, not scored.

Objection 24 (O-Cosmos-24) (the CMB-precision critic)
"You never compute a single angular power spectrum. There is no C_ℓ^{TT}, no C_ℓ^{EE}, nothing from a Boltzmann code. You quote background densities and hand-picked summary numbers and call it a CMB analysis."

Response. Conceded when it was put; since answered. No C_ℓ curve was produced when this objection was written. The paper now prints one: the lensed TT spectrum computed by CLASS on the derived inputs — the exact π-locked fractions, $H_0 = 66.8$, $n_s = 0.9634$, with A_s and τ

inherited — against the 83 binned Planck points, $\chi^2 = 187$ with no fitted parameter against the six-parameter best fit's 84.3. A free amplitude takes it only to 165; a single -0.33% stretch of the multipole axis takes it to 82.5, so the whole excess of the zero-parameter curve is the acoustic-angle deficit and the shape is already right (Figure 25; §40.3). It is a demonstration and not a likelihood, as Objection 25 states: no covariance, no polarization spectra, no joint fit, and that much remains conceded (PS-Cosmos-16). What follows is the answer as it stood. The section "CMB baseline, primordial perturbations, and the scalar/tensor audit" compares background densities and a set of summary observables — the baryon loading, the damping multipole, the scalar tilt, the tensor amplitude — and retains the standard perturbation grammar rather than recomputing it. No C_ℓ curve is produced. A full Boltzmann/C_ℓ prediction is named as future work, and the guardrail was stated in advance: a framework that speaks confidently about dark matter or recombination but never meets the sky at the level of acoustic structure has not entered the precision arena. That was a real scope limit, acknowledged, not papered over.

Objection 25 (O-Cosmos-25) (the CMB-precision critic)
"Every comparison in the paper is central-value-to-central-value. You never confront the framework with the full Planck likelihood or its parameter covariance. Modern cosmology is done with covariances, not eyeball matches."

Response. Conceded in part. The comparisons are of two kinds. The codebook rows are central values against marginal posteriors, each measured against its own published central value and 1σ error, with a theory tolerance $\sigma_{\rm th}$ printed in the pull column where one is known, and with no propagation of the parameter covariance — the π-locked inventory and the one background share their inputs, so the correlations between, say, Ω_m, Ω_b, and $\Omega_{\rm DM}$ certainly matter and are not modeled there. The microwave-background block is the other kind: an executed Boltzmann computation, the zero-parameter TT spectrum from the derived parameters against the 83 binned Planck points with their diagonal errors, $\chi^2 = 187$ — a demonstration and not a likelihood, and the text says so (Section 40). A joint CMB-plus-BAO-plus-distance-ladder likelihood fit with the Planck covariance is the declared next step; the companion Arisaka (2026, C2-Hubble) owns it as the eventual joint fit. Until that fit exists, the reader should read the rows as a battery of central-value checks and the spectrum as a demonstration, not as a likelihood-level test, and we say so.

Objection 26 (O-Cosmos-26) (the CMB-precision critic)
"$\Omega_b = 1/(2\pi^2) = 0.0507$ sits about 2.8% high of the CMB-inferred baryon density, and it disagrees with your own baryogenesis route at the 7.5–28% level. Your single most-cited prediction is also your worst-fitting one."

Response. Conceded, and it is flagged as the largest deviation among the inventory rows in the paper's own audit. The codebook row for Ω_b states that the parameter-free 0.0507 deviates from the empirical 0.0493 ± 0.0006 at the 2.3σ level ($+1.4\sigma$ once the declared epoch input is used), and the second-order-residual audit (§84.5) carries it as a candidate next-order correction, one of three readings to none of which the paper commits. The tension with the $\eta_B \to \Omega_b$

consistency chain is stated openly in the matter-sector section: that chain is order-of-magnitude only; the washout efficiency and transport normalization, uncomputed when this objection was written, are since measured by the time-resolved integration of §I.6, which lands a factor of about ten below the datum ($E = 10.6$). Ω_b is classified Derived-conditional (D*) under the closed-packet-integration conjecture (Conjecture VIII): the structural derivation is closed with no free parameter, but the residual discrepancy is real and is owned as such, not hidden.

Objection 27 (O-Cosmos-27) (the CMB-precision critic)
"Your '0.27σ agreement' for $\Omega_{\rm dm}$ is an artifact of a generous error bar, and in any case it is not an independent test — $\Omega_{\rm dm}$ is just $\Omega_m - \Omega_b$ from the same π-locked ansatz."

Response. Correct on both counts, and neither is disputed. The dark-matter fraction is the closed-form complement $\Omega_{\rm DM} = \Omega_m - \Omega_b = 1/\pi - 1/(2\pi^2) \approx 0.2676$ of the same π-locked inventory; it is not an independent input but an algebraic consequence of $\Omega_m = 1/\pi$ and $\Omega_b = 1/(2\pi^2)$. The 0.27σ figure is therefore a consistency check that the closed-form complement lands inside the present Planck error, not an independent precision triumph, and the size of that error bar is doing real work in making the agreement look tight. We report it as a self-consistency statement of the inventory, classified D* under the Kaxion-misalignment-completeness conjecture for its physical interpretation. The honest content is narrow: given the matter fraction and the baryon fraction, the leftover is dark and has the right magnitude — nothing more.

Objection 28 (O-Cosmos-28) (the CMB-precision critic)
"Your title says you 'derive ΛCDM.' But the scalar tilt n_s and amplitude A_s are merely Asserted, and r is only a bound. You have not derived the dynamical sector at all."

Response. Partially conceded, and the codebook is explicit about where the line now falls. Of the six ΛCDM parameters, the inventory *fractions* (Ω_m, Ω_Λ, Ω_b, and the $\Omega_{\rm DM}$ complement) are Derived-conditional on named conjectures; the scalar tilt is likewise Derived-conditional — pinned to $n_s = 0.9634$ exact on the plateau, not confined to a corridor — and the tensor ratio is Derived-conditional at a central value, $r = 0.0038$, a two-sided target rather than a bound. One dynamical parameter remains imposed and is named as such: the amplitude A_s, the launch sector's single declared input (warp candidate at 75%, flagged; closing computation Lemma V′); τ is astrophysics and outside the claim. So the defensible statement is now: five of the six parameters derived or derived-conditional, one declared input — and where the title reads as the stronger claim, the codebook triage remains the binding statement.

Objection 29 (O-Cosmos-29) (the CMB-precision critic)
"Your deviation plot mixes entries that are anchored by construction — and therefore sit at zero deviation trivially — with genuine predictions. That visual inflates your apparent success rate."

Response. The distinction you describe is exactly the one the Dev plot is built to make, and it is disclosed, not concealed. The caption states that *filled markers* are independent predictions, computed without reference to the measured value, while *open markers* are anchored by construction — the codebook entry echoes the empirical input pending its named lemma — and

"necessarily sit at zero deviation." One entry is anchored in the present table: the amplitude A_s, the one Asserted entry. Earlier versions anchored $H_0^{\rm CMB}$, σ_8, $S_8^{\rm CMB}$ and $N_{\rm eff}$ as well; each is now computed and scored as a filled marker. The genuine predictions that carry nonzero deviation — the atomic-chain $z_{\rm dec}$ at $+0.15\%$ (the switch e^7 annotated 0.62% above the measured value), the derived tilt $n_s = 0.9634$ at -0.4σ, the tensor ratio $r = 0.0038$, the Boltzmann ℓ_D at $+0.4\%$ once the diffusion integral is evaluated, $w = -1$ at $+2.7\%$, Ω_b at $+2.76\%$, σ_8 at -1.4%, and η_B at -20% within its factor-~ 3 band (the time-resolved integration a factor of ten low, §I.6) — are the ones a reader should weigh. A critic is entitled to ignore the anchored markers entirely; the plot is designed so that doing so is trivial, because the marker fill tells you which is which.

Objection 30 (O-Cosmos-30) (the CMB-precision critic)
"You predict background densities but never the acoustic-peak positions or the angular sound horizon θ_*. Those are the actual precision observables of the CMB, and they are absent."

Response. Conceded when it was put; since answered. The released run returns $100\,\theta_* = 1.0381$ against Planck's 1.04109 ± 0.00030 — a -0.29% deficit, pull -2.4 on the 0.12% recombination-modeling spread, the sharpest acoustic row of the codebook — the first three peaks at 221/539/816 against 220/536/813, and the height ratios 0.449/0.439 against 0.453/0.443 (Table 23); they are outputs of the Boltzmann instrument on the derived inputs, and the codebook scores them (Table 41). The damping multipole the row scores is the Boltzmann 1350 ± 3 (1956 in CLASS's $k_D D_M$ convention against Planck's 1947 ± 5); the integer-ledger knee $K(2\pi)^2$ is carried as a conjecture-level overlay $+2.3\%$ above it, under Conjecture VII, and is not the scored value. What θ_* shows is the one place the zero-parameter curve misses, and Section 40 says so. What follows is the answer as it stood. The paper wrote $\theta_* = r_s(z_*)/D_A(z_*)$ as a definition and discussed the geometric meaning of the acoustic ruler, but did not produce a predicted numerical θ_* or a predicted peak-height pattern; those require the comoving sound horizon and the angular-diameter distance integrated over the full thermal history with the recombination microphysics in place, the C_ℓ pipeline whose absence was conceded above.

Objection 31 (O-Cosmos-31) (the CMB-precision critic)
"Your Hubble corridor leans on a pre-recombination $\Delta N_{\rm eff} \approx 0.30$, but your own codebook carries $\Delta N_{\rm eff} = 0.04 \pm 0.18$ [as it did when this objection was put]. Extra radiation at 0.30 is squarely in tension with the Planck $N_{\rm eff}$ fit and with BBN. You cannot have it both ways."

Response. It was the sharpest internal exposure of the corridor, and the corridor has since been retired rather than defended. The pre-recombination 0.30 belonged to a mechanism the paper no longer carries; the codebook row is now the derived null $\Delta N_{\rm eff} \leq 0.03$, $N_{\rm eff} = 3.044$, two-sided against $\gtrsim 0.2$. What follows is the answer as it stood when the corridor was live, kept because the objection was right and because the record of a defense that had to be abandoned is worth more than its erasure. The codebook row for $\Delta N_{\rm eff}$ is the late-time/baseline value, 0.04 ± 0.18, consistent with three Standard-Model neutrinos plus a small thermal residue from the inherited particle ledger (Arisaka, 2026, A9-SM), and it is Derived. The corridor mechanism

of the Hubble-tension Part separately invokes a *pre-recombination* $\Delta N_{\rm eff} \approx 0.30$, and the body says outright that this radiation is "nominally visible to the Planck $N_{\rm eff}$ fit and to BBN." The corridor therefore stands or falls with whether the early-dark-energy-lobe degeneracy direction of the joint CMB+BBN likelihood can absorb the extra radiation — and the paper records that demonstrating this with a full Boltzmann fit is "an explicit open requirement attached to the corridor-uniqueness conjecture, not an established fact." So the honest position is: the tension you name is real, it is the corridor's principal observational risk, and its resolution is deferred to the joint-likelihood fit, here and in the companion (Arisaka, 2026, C2-Hubble). We do not claim to have it both ways; we claim a degeneracy direction that must be earned.

O3 The inflation and dark-energy critic

Objection 32 (O-Cosmos-32) (the inflation critic)
"A TeV plateau cannot do both jobs. Pinning the scalar amplitude $A_s \approx 2.1 \times 10^{-9}$ at an inflationary Hubble scale of order a few TeV forces a slow-roll parameter ϵ_V that is dozens of orders of magnitude away from the one that gives $n_s \approx 0.965$ through $\eta_V \approx -1/N_e$. You never compute the potential, A_s, n_s, or the e-fold count from the geometry; you assert them."

Response. The critic was right — and the two-basin derivation of Part IV is the proof, recorded rather than smoothed over. Sharpened, this objection becomes a sign-level no-go: a lock-scale plateau demands $\epsilon_V \sim 10^{-53}$ for the amplitude *and* yields a blue tilt ($n_s - 1 \sim +10^{-26}$) against the measured red -0.0351 ± 0.0042 — wrong in sign, unrepairable by any parameter. The TeV reading and its codebook conjecture are therefore retired, with the refutation attached (Part IX). What replaced them is not a patched assertion but a derivation: the GDOS boundary law forces the Starobinsky/ξ plateau *class* (the anomaly supplies no scalaron; the one non-conformal scalar, the Σ_r modulus, does), the reheating match at Planck-strength width computes $N^* = 52.6$, and the tilt and tensor ratio follow with no dial: $n_s = 0.9634$, $r = 0.0038$. One number remains imposed, and the codebook says so: the amplitude A_s, the launch sector's single declared input, its warp candidate (75%) flagged and its closing computation named (Lemma V′). A framework that can be argued out of its own published reading by its own critic, and replaces it with a stronger falsifiable one, is behaving exactly as a zero-parameter theory must.

Objection 33 (O-Cosmos-33) (the inflation critic)
"Your codebook calls r a 'Starobinsky-channel' suppression. That is a misnomer. The Starobinsky R^2 model predicts $r \approx 12/N^2 \approx 0.003$ from the e-fold count; your $r \sim 10^{-52}$ is a TeV energy-scale suppression, a completely different mechanism. Calling it Starobinsky is sleight of hand."

Response. Before the plateau class was derived, this objection was fair, and the response could only concede the label. With the plateau class derived, the label is exact: the tensor entry *is* the Starobinsky relation $r = 12/N^{*2}$, because the GDOS boundary law forces the Starobinsky/ξ plateau class itself — the genuine conformal anomaly supplies only W^2 and E_4, so the scalaron

sector must come from the one non-conformal scalar, the Σ_r modulus, and the resulting potential is Eq. (49) with no freedom in shape or slope. The codebook row now reads exactly that way ($r = 0.0038$ exact at the reheat-matched $N^* = 52.6$; $12/N^{*2} = 0.0043$ at leading order), and the retired energy-scale version of the entry — the one this objection correctly attacked — is kept in Part IX with its refutation. The objection, in other words, has been promoted from "fair point about a loose label" to "anticipation of the derivation."

Objection 34 (O-Cosmos-34) (the inflation critic)
"$n_s = 0.965$ via $n_s \approx 1 - 2/N_e$ requires $N_e \approx 57$. You concede you cannot fix N_e. So your headline tilt is a back-fit to a number you cannot derive."

Response. This objection once had to be conceded; it is no longer conceded. The e-fold count is now computed, not targeted: the modulus decays at Planck strength ($\Gamma = M^3/M_{\rm Pl,red}^2$, the only channel consistent with the warped structure), which fixes $T_{\rm rh} \simeq 6 \times 10^{10}\,\text{GeV}$ and, through the standard matching formula, $N^* = 52.6$ — undiluted, because the lock crossing is prompt ($N_2 = 0$, the ledger-action bound with fifty-six orders of margin). The tilt is then $n_s = 0.9634$ exact on the plateau ($1 - 2/N^* = 0.9620$ at leading order), carried in the codebook as **Derived (cond.)**, a pull of -0.4 against Planck 2018 and 3.2σ below the 2025 DESI-combined value — risked against the data, not read from them. The residual conditions are named (the plateau-class derivation and the prompt-crossing bound, Lemma IV revised), and the falsifier is stated: a tilt excluding 0.9634 at three standard deviations in either direction — met today on the DESI-combined side — and, below $\simeq 0.955$, a diluted window that breaks the chain.

Objection 35 (O-Cosmos-35) (the inflation critic)
"Reheating a TeV-scale inflationary patch — with a Hubble rate of order meV — up to the temperatures needed for baryogenesis and BBN is dynamically implausible, and you assert it rather than compute it. Where is the reheating temperature?"

Response. This objection long stood against the program; it is now answered by computation, and the answer dissolves the implausibility along with the question. There is no TeV-scale patch to reheat: the primary plateau sits one e-fold below the GUT edge ($H_1 \simeq 1.5 \times 10^{13}\,\text{GeV}$, not meV), the modulus decays at Planck strength ($\Gamma = M^3/M_{\rm Pl,red}^2$ — the only channel consistent with the warped structure; the alternatives were computed and excluded), the bath spikes to $T_{\rm max} \simeq 7 \times 10^{12}\,\text{GeV}$ and is matched at $T_{\rm rh} \simeq 6 \times 10^{10}\,\text{GeV}$ ($H = \Gamma$; completion at 3×10^{10}) after roughly fourteen e-folds of matter-like oscillation — seven orders above the TeV-scale heavy states, comfortably below the GUT edge, and far above BBN. The lock scale is then reached by ordinary cooling at $t \simeq 1.6 \times 10^{-14}\,\text{s}$ and crossed promptly (the meV curvature scale the objection cites is real, but it is the *terminus* seesaw — the middle basin's curvature — not the launch's Hubble rate). The full numerically integrated history, smooth from the GUT edge to the de Sitter future, is Figures 19–20.

Objection 36 (O-Cosmos-36) (the inflation critic)
"You commit to $w = -1$ exactly and call it 'natural.' But DESI now hints at $w = -1.03 \pm 0.03$ and an evolving dark-energy equation of state. Your framework is in live

tension with the most recent data, and you bury it."

Response. Not buried — this is the framework's committed prediction and a genuine kill-switch, and the tension is real. The late-vacuum Part argues for $w = -1$ exactly, as the frozen residual vacuum export of a completed late-time geometric transition; the codebook lists w as Derived-conditional (D*) under the vacuum-residual-exactness conjecture, against an empirical row that already reads $w = -1.028 \pm 0.032$. So the live hint is printed in the codebook itself, not hidden. A confirmed, robust $w \neq -1$ — whether the DESI $w_0 w_a$ preference or any future tightening — falsifies the module: the framework has no dial to absorb it, because the whole point is that the vacuum term is a frozen residual, not a tunable quintessence field. This is openly flagged as the one place the cohort is already in tension; see the falsifier deck of Arisaka (2026, A11-Origin). The honest status is: committed prediction, currently at the $\sim 1\sigma$ edge of a developing anomaly, with no escape route by construction.

Objection 37 (O-Cosmos-37) (the inflation critic)
"$\Omega_\Lambda = 1 - 1/\pi$ is a fit-then-rationalize. You measured $\Omega_m \approx 0.315$, noticed $1/\pi \approx 0.318$, and reverse-engineered a story. And the meV 'seesaw' for the vacuum scale double-counts — it is a second dressed-up version of the same coincidence."

Response. Conceded on both points, with the honesty already in the text. $\Omega_\Lambda = 1 - 1/\pi$ is not an independent derivation: it is the unit-trace complement of $\Omega_m = 1/\pi$, obtained from late-time flatness $\Omega_m + \Omega_\Lambda \simeq 1$. The codebook carries Ω_m, Ω_Λ, and Ω_b as Derived-conditional (D*) under the closed-packet-integration conjecture, and the topological-completion Part states in full that the printed Step-4 cancellation does not yet close — so the $1/\pi$ value is a conjectured topological identity, not a proved theorem. The Ω_Λ row literally cites "via unit-trace identity," i.e. it inherits whatever Ω_m earns and adds nothing independent. The meV vacuum scale is presented as a seesaw *descendant* of $M_{\rm lock}$, explicitly flagged as a suggestive dimensional coincidence, not as a second independent derivation of the dark-energy density. The deep four-form integration meant to prove $\Omega_m = 1/\pi$ topologically (and exclude competing rational functions of π) is reserved for a later sweep; until then the critic's "fit-then-rationalize" charge is the correct default reading, and the codebook classification (D*, not D) concedes it.

Objection 38 (O-Cosmos-38) (the inflation critic)
"$r = 0.0038$ is a razor-thin target — an enormous, dangerous claim. You are betting the module on the Simons Observatory and LiteBIRD finding a primordial B-mode at exactly one amplitude."

Response. Yes — and that is deliberate, not an oversight; the two-basin derivation made the bet strictly more dangerous, on purpose. The retired TeV reading predicted $r \sim 10^{-52}$: safety by invisibility, falsifiable only from one side. The derived plateau predicts a *number*, $r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order), a factor nine below the current bound and several times above the Simons-Observatory/LiteBIRD floor — so the switch now fires in both directions: a detection well above it kills the launch, and so does a clean exclusion well below it. Carried in the codebook as **Derived-conditional (D*)**, with no dial to move (r rides on N^*, and N^* is pinned by the reheating match and the prompt crossing). A strong, dangerous,

two-sided, dated claim is the feature here, not a bug: the framework is obliged to produce a definite value and accepts that this decade's instruments will pass verdict.

Objection 39 (O-Cosmos-39) (the inflation critic)
"'Slow-smooth inflation with graceful exit' is narrated through GDOS, GRIP, GSSB, and GFF vocabulary, but you never exhibit a computed potential with an actual exit point. The exit is a story, not a calculation."

Response. Conceded. The exit is described *structurally*: graceful exit is the failure of slow roll at the edge of the GFF-dominant plateau, with the lock GSSB at the infrared seam of the throat as the later crossing and the whole sequence read as the prototype cosmological GRIP-Out in the runtime grammar inherited from Arisaka (2026, A6-GRIP). The paper writes the plateau form of the derived class, its closed-form ϵ_V, η_V, and the exit point on it, $y_{\rm end} = 0.464$ (§28.10); what it does not do is fix the class's coefficients (c_1, β, the throat width) from the geometry, so the exit is computed on the class rather than derived from the throat. The mechanism is coherent and the exit is located (slow-roll failure at the plateau's edge; the lock GSSB that seeds the Kaxion misalignment and the late vacuum residue is the later, separate crossing), but the geometric derivation of the coefficients is an open item. The narrative scaffolding is real; the calculation behind it is not yet done.

Objection 40 (O-Cosmos-40) (the inflation critic)
"The codebook's record carries a retired conjecture that located the inflationary plateau at the TeV lock scale, while the universe begins at the GUT edge $E_{\rm GUT}$ — wildly separated scales, with an uncomputed thirty-e-fold gap between the beginning and the launch. The paper now puts the plateau at the top. Which is it, and what filled the gap?"

Response. The gap the objection attacks was real, and closing it is precisely what the two-basin derivation did — not by explaining the descent-then-inflate ordering, but by discovering it was backwards. The plateau is not at the lock: the GDOS-derived class plus the observed amplitude place it at $V^{1/4} = 8.1 \times 10^{15}\,\text{GeV}$, *one e-fold below the GUT edge* — so the beginning and the launch are adjacent, the "gap" between them is a single e-fold, and inflation is literally the opening stretch of the descent, exactly as the Part's GRIP narrative always said. The twenty-eight units down to the lock are then fourteen e-folds of modulus decay and sixteen of ordinary radiation cooling, computed, with the lock crossed promptly at $t \simeq 1.6 \times 10^{-14}\,\text{s}$ (the middle basin's curvature scale, $M_{\rm lock}^2/\sqrt{3}M_{\rm Pl,red}$, survives as the terminus seesaw — the number the old reading had misassigned to the launch). The explicit trajectory across the whole history, numerically integrated and smooth, is Figures 19–20. The dynamical bridge the old response conceded as open is now the computed content of Part IV.

O4 The Hubble-tension and large-scale-structure critic

A note before the objections of this section, because several were aimed at a reading the paper no longer holds. Part VI once read the Hubble tension as a *two-stage corridor*: an early branch

whose sound horizon was shortened by an early-dark-energy lobe and a radiation excess, a late branch at the ladder's 73, and a switch between them. That reading is retired (§42), for three reasons the objections below helped force: the switch entered neither branch function and so did no work in the mechanism named for it; the late branch's projection factor was the distance-ladder calibration itself, imported from the measurement it was meant to explain; and the early source is forbidden by this module's own kill criterion $f_{\rm EDE} \gtrsim 0.02$. What stands in its place is one derived background rate, $H_0 = 66.8 \pm 0.2$, read by the microwave background at 67.4 through the inherited sound-horizon calibration, predicted for every ladder-free late probe, and contradicted by the calibrated ladder alone — which the framework now treats as a prediction about the ladder. Objections aimed at the corridor's internal consistency are therefore answered by its withdrawal rather than by a defense, and the responses below say which. The retired mechanism, with its arithmetic, is kept in Appendix K.10.4.

Objection 41 (O-Cosmos-41) (the tensions critic)

"Your two headline numbers, $H_0^{\rm CMB} \approx 67.4$ and $H_0^{\rm late} \approx 73.0$, are simply read off Planck and SH0ES. You did not derive either one; you transcribed them and drew a box around them."

Response. Half conceded, and the half that is conceded has since been repaired. The 73.0 was transcribed: it was the ladder's own value, placed in the theorem as a second branch with a projection factor that was the ladder calibration itself, and that branch is withdrawn — the theorem now carries one background rate (C3). The 67.4 is not transcribed: the background is derived with zero sky input, $H_0 = 498.3\sqrt{\zeta_\Lambda} = 66.8$ with $\zeta_\Lambda = \pi/(LK)$ selected by the stratum rule (Section 29.8), and independently at 66.6 from the π-locked fractions read against the measured densities; the CMB's 67.4 is the comparator, 0.9% away, and the sound-horizon projection factor that would close that last percent is inherited, not computed, as the theorem's remark says. What is on the ledger is one rate and a prediction that the ladder carries a systematic; the companion Arisaka (2026, C2-Hubble) reads it the same way.

Objection 42 (O-Cosmos-42) (the tensions critic)

"Two mutually inconsistent stories travel under the single name 'two-stage corridor': in one place it is a kinematic reframing with 'no new physics, just two projections,' and in another it is a real early-dark-energy lobe plus a local-void shift. Which is it?"

Response. The objection is right, and it is now answered by subtraction rather than by distinction. The exposition did carry two registers under one name: a kinematic reframing — two readings of one background — and a dynamical mechanism carrying an early-dark-energy residue, an extra-radiation term and a local-underdensity lift. The dynamical register is withdrawn in full: the lobe is forbidden by the module's own kill criterion, the radiation term is the derived null $\Delta N_{\rm eff} \leq 0.03$, and the local lift required an underdensity the paper's own appendix records as unobserved. Only the kinematic register survives, and it has been sharpened into something falsifiable: one derived rate, and a prediction that the calibrated ladder is the outlier. There is now one story, and it is the one that was not carrying the physics. What has observational teeth is the one-background prediction itself — the ladder-free probes must read 66.8, and the ladder's excess must be found in its calibration — and that is what the paper stakes; no

dynamical corridor mechanism is claimed, and none is owed.

Objection 43 (O-Cosmos-43) (the tensions critic)
"A corridor that reproduces two endpoints but commits to no inverse-distance-ladder, no BAO(z), and no $H(z)$ shape in between is not a model. It is a pair of dots with a label."

Response. Conceded when it was written, and the ground has since shifted under it in the paper's favor. There is no longer a corridor interior to commit to: one background means one $H(z)$, the standard one on the π-locked inventory with $w = -1$, and the inverse-distance-ladder and BAO(z) tracks it implies are the ordinary ones for those densities and that rate — with no free shape and nothing to fit. That is a stronger commitment than the corridor could make, not a weaker one, and it is what makes the BAO row of the falsification catalog two-sided: a measured two-stage residue at the percent level now counts *against* the paper. What the paper still does not print is a full likelihood over CMB, BAO and supernovae, which is named as required in the limitations Part. Until that fit exists, the reading commits to the rate and the shape, not the interpolating background. We mark this as a declared deliverable, not a quiet omission.

Objection 44 (O-Cosmos-44) (the tensions critic)
"You assert that the CMB acoustic fit, BAO, and BBN all survive your early-dark-energy plus dark-radiation insertion simultaneously. You never demonstrate it — the paper itself admits the global fit is required and absent."

Response. Conceded when written; the task it names has since been dissolved rather than done. It was the single most important open numerical task of the corridor module, and the corridor is retired: there is no early-time residue whose full likelihood must be checked, because there is no early-time residue. The lensing/EDE-residue section states explicitly that radiation at $\Delta N_{\rm eff} \approx 0.30$ over the whole pre-recombination regime is nominally visible to the Planck $N_{\rm eff}$ fit and to BBN, and that "demonstrating this with a full Boltzmann fit is an explicit open requirement attached to the corridor-uniqueness conjecture, not an established fact." Simultaneous preservation of the acoustic peaks, the sound horizon, the BAO ruler, and the light-element abundances is asserted as a target, not shown. That target was the corridor's, and the corridor has since been retired rather than defended; what the paper now stakes on the acoustic peaks, the sound horizon and the BAO ruler is one background and no residue.

Objection 45 (O-Cosmos-45) (the tensions critic)
"Your S_8 corollary assumes the early-time corrections *suppress* clustering. But in full EDE likelihood fits the net S_8 shift comes out *positive* — refitting drives $\omega_{\rm cdm}$ and n_s up. Your corollary has the sign backwards."

Response. Conceded, and the paper flags exactly this. The S_8/σ_8 section carries a mandatory cautionary note: "in full likelihood analyses of early-dark-energy-type models, refitting the CMB typically drives $\omega_{\rm cdm}$ and n_s upward and the *net* S_8 shift comes out positive — the opposite sign to the suppression encoded above" (Hill et al., 2020). The suppression branch in Eq. (673) belongs to the retired corridor and is kept only in the record appendix; on the one background

the paper now carries there is no growth response to have a sign, and $S_8 = 0.825 \pm 0.010$ is computed from the declared amplitude (§43.3), $+2.5\sigma$ above cosmic shear. We do not claim the S_8 tension is resolved; we claim the number, and the number sits on the CMB's side of it.

Objection 46 (O-Cosmos-46) (the tensions critic)
"The response coefficients 0.13 per unit $\Delta N_{\rm eff}$ and 0.40 per unit $f_{\rm EDE}$ are lifted straight from the early-dark-energy literature. They are not outputs of your throat/Σ_r geometry. So where exactly is GG-Theory doing any work?"

Response. Conceded. The text is explicit: "the numerical values used below, 0.13 per unit $\Delta N_{\rm eff}$ and 0.40 per unit $f_{\rm EDE}$, are representative response slopes of the early-dark-energy literature (Poulin et al., 2019), not outputs of the GG ledger." What GG-Theory was asked to supply was the *source pair* that feeds them — and the answer, now, is that both are zero: the early residue is forbidden by the module's own kill criterion and $\Delta N_{\rm eff} = 0$ is a derived null (§42), so the slopes have nothing to act on and the parametrization is retired to the record. The parametrization was borrowed and is now retired with the source pair it acted on; nothing the paper carries depends on those slopes.

Objection 47 (O-Cosmos-47) (the tensions critic)
"A10 cites C2 for the H_0 numbers and C2 cites A10 right back — a citation loop. The global fit that would justify either is owed by both papers and delivered by neither."

Response. The reciprocal citation is real and we acknowledge it. The division of labor is clean even if the prose has not always made it loud: the structural claim — one background rate with an integer-ledger switch, and its theorem — belongs to *this* paper; the dedicated H_0 phenomenology and the eventual joint CMB+BAO+ladder fit are the remit of Arisaka (2026, C2-Hubble). C2 is the owner of the fit. The loop is a division of labor (this paper scores the rate against Planck's 67.4, the companion against the 2025 three-instrument 67.24 ± 0.35; both are stated in §42.4), not two papers each pointing at the other for a result neither holds: the structural claim has a home here, the numerical fit has a named home in C2.

Objection 48 (O-Cosmos-48) (the tensions critic)
"$\Delta N_{\rm eff}$ appears as 0.30 in your corridor mechanism and as 0.04 in your baseline codebook row. One paper, one symbol, two numbers — that is a contradiction sitting in plain sight."

Response, since rewritten: the objection is right, and the defense this response used to give was wrong. Earlier versions answered that the two numbers are different quantities — 0.30 the "pre-recombination corridor source" and 0.04 the "late-time Planck baseline" — and that no contradiction existed. That answer does not survive its own paper. This framework *derives* $\Delta N_{\rm eff} \approx 0.04$ from the inherited neutrino block; the codebook grades it **D**, computed and not fitted; and the kill criteria of §42.7 stake the module's life on it, retiring the framework if "a radiation excess incompatible with $\Delta N_{\rm eff} \approx 0.04$" is found. A quantity a theory derives, grades, and can die of is not free to be re-proposed at seven times its derived value in a different section. They were two values of one quantity, and the paper was carrying both.

What the correction cost, and where it ended. The objection was answered in two stages, and the second is the one that settles it. First the body was made to print one value, 0.04; the price was paid in the corridor, which at 0.30 over-delivered the sound-horizon shortening it needed (1.57% against 1.20%) and at 0.04 delivered three fifths of it, the balance owed to a refit the paper had not performed. Then the corridor itself was retired, and with it the demand: there is no sound horizon to shorten, and the row is the derived null $\Delta N_{\rm eff} \leq 0.03$ — $N_{\rm eff} = 3.044$, the Standard Model's own value, because below the lock scale the field content is the Standard Model's. The intermediate 0.04 was itself an error of a different kind, the Standard Model's own $3.044 - 3$ counted a second time and attributed to the particle companion, which contains no such residue. Three values of one symbol have been in print; one is now, and it is a null the framework can be killed by. So the honest position after the correction is weaker than the position before it, and it is the one the framework's own derivation forces. The legacy 0.30 arithmetic is kept in Appendix K.10 as a non-unique cross-check, clearly labeled, because retiring a number quietly is how the contradiction survived as long as it did.

Objection 49 (O-Cosmos-49) (the tensions critic)
"Your JWST early-galaxy 'solution' is a semianalytic boost — a lowered collapse threshold, a 20% streaming-velocity damping, a bump in the minihalo tail. Luminosity functions and number counts at $z \gtrsim 10$ are nonlinear. Without N-body you have a story, not a prediction."

Response. Conceded without hedging, and conceded further than the objection asks. The scorecard lists the JWST early-galaxy entry as *Partially resolved* explicitly because "nonlinear N-body simulations under the inherited ledger [are] still pending," and the limitations Part devotes a dedicated subsection to the point: a semianalytic boost to early structure "is not the same thing as a robust prediction for luminosity functions or galaxy number counts." What the framework claims is narrower and survivable: the scale distinction between the broad linear S_8 amplitude and the exponentially sensitive minihalo tail means a coherent-but-non-uniform solution is at least *permitted* — early collapse can advance without spoiling broad growth. That is a consistency argument, not a number, and the number requires the nonlinear pipeline named in the medium-term research tasks.

The further concession. Of the three items the objection lists, two are no longer offered. "A bump in the minihalo tail" was said in earlier versions to follow from the corridor's early residue; the computation was run and it does not (§44.3). The streaming damping needed a light gauged-baryon vector, and the terminated baryon gauge sector has none below the lock scale (§44.4). A background episode acts on every sub-horizon mode by one common factor, so the residue cannot lower the broad amplitude and raise the rare tail — and the scale distinction this response leans on is therefore a property of the one *baryonic* lever, the lowered threshold, and of nothing else in the ledger. The consistency argument survives on that one. The bump and the damping are withdrawn, and a semianalytic treatment of either would now be a semianalytic treatment of nothing.

O5 The dark-matter, baryogenesis, and BBN critic

Objection 50 (O-Cosmos-50) (the matter-sector critic)
"Your dark-matter density is $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$. That is just an algebraic subtraction — the matter share minus the baryon share. It is a definition, not a relic calculation. You have not computed a freeze-out or a misalignment yield; you have rearranged a ledger."

Response. The fraction and the dynamics are two different claims, and you are right about the first. As an identity, $\Omega_{\rm DM} = \Omega_m - \Omega_b = 1/\pi - 1/(2\pi^2)$ is exactly the closed-form complement of the π-locked inventory; it is bookkeeping, and the section "Kaxion dark matter" presents it as such. What is *not* bookkeeping is the relic dynamics, and that is derived in the companion Arisaka (2026, C1-Kaxion), not here. There one scale — the bulk normalization behind the Kaxion decay constant f_K — is *calibrated* so that a generic misalignment angle saturates the π-locked abundance at the benchmark cap loading, and the closure then returns that angle, $\theta_i \simeq 1.1$ rad. The abundance is the *target*, the calibration is declared once, and nothing after it is adjustable; it is emphatically not a fit of $\Omega_{\rm DM}$ — $\Omega_{\rm DM}$ is the input on the right-hand side. The honest statement is that the *number* is a subtraction identity and the *mechanism that saturates it* is a one-calibration closure; the codebook tags $\Omega_{\rm DM}$ as D^* under the Kaxion-misalignment-completeness conjecture, which says precisely that the dynamical closure is conjectural, not closed.

Objection 51 (O-Cosmos-51) (the matter-sector critic)
"Your framework geometry-selects nearly every number it touches, yet the misalignment angle θ_i is read off the observed abundance as a generic $\mathcal{O}(1)$ number. Why is this one an output? And since $1/(2\pi)$ is exactly the kind of value your ledger likes, is it excluded or merely unmentioned?"

Response. It is an output by construction, not by exception. The Kaxion sector presents two unknowns — the decay constant f_K and the angle θ_i — and one equation; the geometry fixes the ladder f_K sits on but not, at present, its absolute scale. Arisaka (2026, C1-Kaxion) closes the pair by convention, stated once: the scale is calibrated so that a generic angle meets the abundance at the benchmark rung, and the angle is read back, $\theta_i \simeq 1.1$ rad. Naming θ_i a geometric selection as well would be claiming a determination the framework does not have.
On $1/(2\pi)$: it is excluded, not merely unmentioned. Nothing in the OGN-5 ledger selects it for this slot, and it is quantitatively wrong — $1/(2\pi) = 0.159$ rad lies a factor of about seven below the 1.18 ± 0.04 rad that the closure (774) returns against the observed abundance. A geometry-selected θ_i at that value would require the scale to be recalibrated by a factor 50 in the abundance, which is to say it is not what the closure returns.

Objection 52 (O-Cosmos-52) (the matter-sector critic)
"You advertise the Kaxion mass $m_K \approx 8.1$ μeV as a prediction. But that is just the standard QCD-axion mass image of f_K: once you fix the decay constant, $m_K f_K = \sqrt{\chi_{\rm QCD}}$ hands you the mass for free. It is not an independent second prediction; it is the same

number wearing a different hat."

Response. Conceded, and stated the same way in Arisaka (2026, C1-Kaxion). The relation $m_K f_K = \sqrt{\chi_{\rm QCD}}$ is the standard strong-CP axion relation, with $\chi_{\rm QCD}$ fixed by Standard-Model lattice input; once f_K is set, $m_K \approx 8.1\ \mu$eV (a 1.95 GHz tone) is its DFSZ image, not a logically separate output. The genuinely geometric content is not f_K's absolute value — that is calibrated — but the ladder f_K sits on and the count: the frequency *window* 1.59–2.13 GHz of seven rungs on the DFSZ coupling line, and one line standing alone across it, a second in-window line refuting the count and naming a registered mode of the parents. Those, not the mass value, are what carry the falsifiable content; m_K is a consequence of f_K and the measured susceptibility, and the paper should be (and is) read that way.

Objection 53 (O-Cosmos-53) (the matter-sector critic)
"Your baryon asymmetry $\eta_B \simeq 4.89 \times 10^{-10}$ rests on $\epsilon_{\rm geo}^{\rm eff} = e^{-K/2}$, and you admit on the same page that the exponent $K/2$ is a postulate, not derived. Then you present a central value within 20% of the observed η_B and call the chain controlled. That is numerology dressed as a derivation."

Response. Conceded. The section "Baryogenesis via the inherited neutrino and CP ledger" says in plain text that the exponent $K/2$ is a postulate of the present compression — geometrically suggestive as $\sqrt{e^{-K}}$ but not derived — and that it is part of what the Resonant-leptogenesis-normalization conjecture (Conjecture VI) must fix. The washout efficiency and transport normalization were uncomputed when this objection was written; they are since measured by the time-resolved integration of §I.6 ($E = 10.6$), and the compressed chain's $\mathcal{O}(1)$–$\mathcal{O}(10)$ coefficient is now a number, a factor of about ten, stated one-sidedly. For exactly this reason η_B is demoted to an *order-of-magnitude* claim and displayed in the parameter-audit Dev plot with a factor-$\sim$3 theory band, and the text explicitly warns that closeness of central values at this tier must not be read as a precision success. The -20% figure is reported as a central value inside a wide honest band, not as a precision match. It is tagged D* under a named conjecture, not advertised as closed.

Objection 54 (O-Cosmos-54) (the matter-sector critic)
"Your lithium-7 fix now sets the $A = 7$ survival fraction to the closed-packet $1/\pi$ — the *same* factor you use for $\Omega_m = 1/\pi$. So you are plugging two different holes, the cosmic inventory and the lithium deficit, with one unproven number. That is not a derivation of either; it is a single conjecture (Conjecture VIII) doing double duty. And you still have not exhibited the microscopic mechanism that makes the port act on mass-7 at all."

Response. A fair sharpening, and the honest status is exactly as you state it. The $A = 7$ port strength is *not* a tuned coefficient: the surviving mass-7 fraction is the closed-packet $1/\pi$, the same factor that fixes $\Omega_m = 1/\pi$, so the predicted primordial abundance is $(^7{\rm Li/H}) = {\rm SBBN}/\pi \approx 1.6 \times 10^{-10}$ with *no* free knob (the section "The $A = 7$ port and the $1/\pi$ prediction" and the leptogenesis appendix carry the derivation). That removes the "one-parameter cheat" charge — but it does not remove the dependence on Conjecture VIII, and you are right to press

there. Because the port reuses the *same* closed-packet $1/\pi$, the lithium fraction inherits exactly the status of the inventory: Derived-conditional (D*) under the one closed-packet-integration conjecture, not an independent success. One unproven number doing double duty is a more economical claim than two free knobs, but it is no less conditional, and the codebook tags it that way. What remains genuinely open — the open item PS-Cosmos-15 — is not a coefficient to fit but three structural obligations: *why* the port is selective in mass number (acting on $A = 7$ and not $A = 2, 3, 4$), the full coupled BBN-network closure confirming the other abundances do not move, and compatibility with detailed balance. The lithium-7 problem is therefore reported as a parameter-free but conditional prediction, falsifiable by any firm primordial ^{7}Li away from SBBN/π — not as solved.

Objection 55 (O-Cosmos-55) (the matter-sector critic)
"Worse: a selective sink on a single mass number is exactly the kind of thing that violates detailed balance and contradicts well-measured nuclear cross-sections. You cannot just remove ^{7}Be/^{7}Li by fiat without the inverse reactions repopulating it. Your port is unphysical."

Response. The honest answer is that this is exactly the question the BBN-network implementation (PS-Cosmos-15, open) must settle, and the present paper flags it as open rather than claiming to have passed it. The port is proposed as a *selective open-system sink* in the GDOS/GFF ledger — a seam-level loss term acting on the late, tiny $A = 7$ inventory, with surviving fraction fixed by the closed-packet $1/\pi$ rather than by a free coefficient — not as a modification of a measured two-body cross-section, precisely so that it does not touch the robust deuterium and helium yields fixed earlier in the network. Whether such an open-system sink respects detailed balance, or whether the inverse channels repopulate the $A = 7$ abundance and neutralize it, is not something the effective seam-level description can answer by itself; it requires the full network. That check is named as a deferred obligation, and the lithium row carries the Derived-conditional (D*) tag for this reason. We are not asserting it is consistent; we are flagging it as the load-bearing thing the BBN-network implementation (PS-Cosmos-15) must prove — an open item of this paper, carried by no companion.

Objection 56 (O-Cosmos-56) (the matter-sector critic)
"You have two routes to the baryon density and they disagree. The topological route gives $\Omega_b = 1/(2\pi^2) = 0.0507$; the baryogenesis route gives $\Omega_b \approx 0.0393$ at $h = 0.674$, a -22% gap — and your $\Omega_b h^2$ sits $\sim$2.8% high anyway. A framework that derives the same observable two ways and gets two answers has not derived it."

Response. Conceded, and the two-route tension is flagged openly in the section "A consistency chain from η_B to Ω_b." The topological route gives $\Omega_b = 1/(2\pi^2) = 0.05066$; the η_B route gives $\Omega_b h^2 \simeq 0.0179$, i.e. $\Omega_b \approx 0.0393$ at $h = 0.674$, a -22% discrepancy (-21% at the carried $h = 0.6676$; the ladder's 73.5 is not a second rate the framework carries, and converting at it would only widen the gap, to -35%), which sits inside the transport chain's factor-$\sim$3 band. The two are not, however, roads of equal standing, and the response should say so rather than leave the reader to infer it: one is the paper's prediction of the number and the other is a

mechanism checked against it. Converted through the same bridge at the structural rate of Route I, $h = 0.6676$ (zero sky input), the ledger road gives $\eta_{10} = 6.18$ against an observed 6.12, an error of +1.1%; the transport mechanism gives 4.89, which is 79% of that target. (At the density road's $h = 0.666$ the ledger road gives 6.15, +0.57% — but that h takes the measured ω_b and ω_m as inputs and is therefore not independent of the datum, so it is not the number to lead with.) The ledger road is the paper's prediction; the mechanism's shortfall against it — 21% in the closed form, a factor of about ten in the time-resolved integration — is what the finite-temperature hypothesis has to close. The paper states in plain text that the two routes do not yet agree and records "two-route-consistency" as a named open problem alongside the closed-packet-integration conjecture. Ω_b itself is carried as D* under Conjecture VIII, and its +2.76% deviation is the largest precision-register entry in the parameter-audit Dev plot (the η_B entry sits farther out, carried on its order-of-magnitude band), marked as a candidate next-order correction (full entropy history, sphaleron throughput, reheating profile). You are right that two routes disagreeing is not a derivation; the paper's claim is the weaker, honest one — a consistency demonstration with the residual offset named, not erased.

Objection 57 (O-Cosmos-57) (the matter-sector critic)
"Strip away the language and the whole matter sector is "Partially resolved modulo three named conjectures." That phrase means the matter sector is not derived. You are claiming credit for a derivation you have explicitly told me you have not completed."

Response. Agreed on the words, and that is the intended reading, not a slip. The matter sector is a *consistency demonstration up to named conjectures* — the Kaxion-misalignment-completeness, Resonant-leptogenesis-normalization, and closed-packet-integration conjectures — not a closed derivation, and the codebook tags say exactly this: $\Omega_{\rm DM}$, η_B, and Ω_b are all D* (Derived-conditional), with lithium-7 D* under Conjecture X, its network closure an open item (PS-Cosmos-15). The paper's repeated claim is the asymmetric, honest one: a small set of *precise* conjectures, once proved, would close the matter ledger, and no recognized problem is left without a named route forward. That is a different and weaker claim than "derived," and the provenance triage is built specifically so a reader can tell which is which. The credit claimed is for a constrained, falsifiable structure with its open conjectures labeled — not for a finished proof.

Objection 58 (O-Cosmos-58) (the matter-sector critic)
"Two more loose threads. Your resonant TeV leptogenesis produces η_B from a CP-odd invariant you assert is "controlled by geometric phase" — but you never compute the phase, you just assert it lands right. And a misalignment dark matter that solves strong CP should source CMB isocurvature; have you actually checked you are not already excluded?"

Response. Two separate points. On leptogenesis: conceded that the CP-odd invariant is asserted, not computed. The section "Baryogenesis via the inherited neutrino and CP ledger" takes the standard resonant formula and argues that GG-6 *narrows* both the CP-odd phase and the Breit–Wigner enhancement into the inherited geometry, but the actual phase value, washout,

and transport collapse into the same $e^{-K/2}$ throughput — postulated, and since checked by the time-resolved integration of Appendix I.6 — flagged in Objection 53; the neutrino-and-CP ledger it leans on is the Majorana/CP block derived in Arisaka (2026, A9-SM), and the resonant phase is its named, unclosed obligation. On isocurvature: the question must be asked at the right epoch, and an earlier version of this response asked it at the wrong one, quoting a tiny $\beta_{\rm iso}$ from the lock epoch's low Hubble rate. The lock crossing is prompt and imprints nothing; the budget is decided by the primary inflation, whose $H_1 \simeq 1.5 \times 10^{13}$ GeV would randomize a light angle outright. The paper's answer is therefore conditional: under the inflation-era pinning conjecture (C-Kaxion-d, the section "Kaxion dark matter") the angle carries no fluctuation and $\beta_{\rm iso}$ vanishes identically; without it the angle is randomized *during* inflation — there is no symmetry restoration that could randomize it afterward — which decorrelates it in a quarter of an e-fold across a sky spanning seven and deposits an order-unity cold-dark-matter isocurvature, ten orders of magnitude in power above the *Planck* bound; at the bulk scale the angle is quiet but its pre-inflationary fluctuation still exceeds the bound by eight orders, and no decay constant below 5×10^{17} GeV evades it; the conjecture is two-valued, and a detection kills it outright. So the isocurvature worry is answered by a named conjecture with named killers rather than by a computation — and the conjecture is necessary, not optional: if it fails, this paper's compact-angle route to the dark-matter share fails with it; the leptogenesis phase is genuinely open, and is tagged as such.

O6 The gravitation, black-hole, and primordial-gravitational-wave critic

Objection 59 (O-Cosmos-59) (the gravitation critic)
"Your Part 'The six-dimensional bulk' opens from a boundary Einstein equation as though four-dimensional general relativity were a free gift. But you never derive it. You assert it and then dress it up by inserting the FRW metric. A constitutional theory is supposed to earn its gravitational law, not borrow it."

Response. The charge is half right and the half it gets right is one we state ourselves. GG-A10-Cosmos does not derive boundary general relativity; it inherits it. The text of 'The six-dimensional bulk' is explicit that the paper "begins from the boundary Einstein equation rather than postulating it," and the word "rather" is doing real work only because the derivation lives upstream. That derivation is in the companion Arisaka (2026, A5-GDOS), Part IV, theorems T8.1 through T8.4: the boundary Einstein–Hilbert action as the metric zero-mode export of the bulk (T8.1–T8.2), the Newton constant fixed by the scale-overlap integral $M_{\rm Pl}^2 = M_6^4\, 2\pi R_\theta \int e^{-2A} {\rm d}r$ rather than left free, and stress-energy conservation $\nabla_\mu T^{\mu\nu} = 0$ as the boundary diffeomorphism Ward identity (T8.3). So the right reading is a division of labor, not a gap: A5 proves the law, A10 applies it to the homogeneous-isotropic background. We would only be guilty as charged if we claimed to re-prove GR here; we do not, and the codebook traces the inheritance explicitly.

Objection 60 (O-Cosmos-60) (the gravitation critic)
"Even granting A5, the projection you actually use reproduces only the FRW background. You get the two Friedmann equations and the continuity equation by hand-feeding a perfect fluid. That is not the full Einstein equations, not the Bianchi identities, not local conservation off the symmetric background. You are advertising 'boundary GR' while delivering one ODE in the scale factor."

Response. Conceded as a statement of scope, and it is the correct scope. GG-A10-Cosmos operates at the FRW background level because that is all cosmology at large scales requires: homogeneity and isotropy force the maximally symmetric spatial slice, the 00 and trace components of the boundary Einstein equation give the two Friedmann relations, and covariant conservation gives the continuity equation. The full local field equations, the Bianchi identities, and the seam-level conservation law are the remit of Arisaka (2026, A5-GDOS), where conservation is a Ward identity (T8.3) holding off any symmetry restriction. A10 should not be read as re-deriving the full local dynamics; it reduces an already-derived local law to the symmetric sector and then does its real work on the source ledger. If a reader wants the local Einstein equations or the Bianchi structure, the citation is A5, not this paper.

Objection 61 (O-Cosmos-61) (the gravitation critic)
"Your 'Page curve as a GSSB saddle switch' in the black-hole Part is theater. You rename the island formula as a 'GFF-minimizing saddle,' call the transition a 'GSSB event,' and declare victory. There is no replica computation, no island extremization done in your framework, and no quantitative Page time that you compute rather than borrow. Relabeling is not deriving."

Response. Conceded, and it is conceded corpus-wide, not just here. The GSSB language in 'The Page curve as a GSSB saddle switch' names which physical thing is changing at the turnover — the operationally dominant saddle of the open-system entropy ledger — but it does not replace the replica/island computation that establishes the generalized-entropy rule in the first place. We say as much: our entropy functional "is mathematically the same object" as the island generalized entropy, and the island program tells us *how* to compute the saddle while GDOS tells us *why* a branch-minimized functional had to be the answer. The honest quantitative status is the same as in the companion Arisaka (2026, C3-BH), which computes the $A/4$ coefficient and thermality (T-BH-1) but treats the island/replica Page-curve as conceptually closed and quantitatively open. Our own Page-time estimate $M_{\mathrm{Page}} = M_0/\sqrt{2}$ is a structural consequence of the half-entropy condition, and we flag in the text that "what matters physically is not the exact numerical prefactor." The slogan names the saddle switch; it does not perform the replica calculation, and we do not pretend it does.

Objection 62 (O-Cosmos-62) (the gravitation critic)
"The whole black-hole Part leans on the Bekenstein–Hawking thermodynamic fact, yet you defer the actual $A/4$ entropy coefficient to C3. So A10 proves nothing about the very fact it depends on. You are decorating an imported result."

Response. Correct on the facts, and again this is a deliberate division of labor rather than a

hole. GG-A10-Cosmos uses the area law as an established thermodynamic input and builds the open-system reading of information return on top of it; it does not claim to prove $A/4$. The derivation is in Arisaka (2026, C3-BH), Theorem T-BH-1 (Derived-conditional on three named premises), where the horizon entanglement entropy is the induced $A/4G_N$ and the conical-replica method fixes the coefficient $1/4 = 4\pi/16\pi$ regulator-independently for the generic throat. That is exactly the right place for it: C3 is the entropy paper, A10 is the cosmology paper. The objection would bite only if A10 claimed the area law as its own result. It does not; it cites it.

Objection 63 (O-Cosmos-63) (the gravitation critic)
"Your CMB section and your inflation Part both now quote $r = 0.0038$ as if it were obviously right, but the whole number rests on one e-fold count, $N^* = 52.6$, which rests on one reheating calculation. One number out of place and the whole edifice falls."

Response. Yes — and that strength is the point, not an embarrassment. The chain is short and every link is stated: the GDOS-forced plateau class gives $r = 12/N^{*2}$ at leading order, 0.0038 exact; the Planck-strength width gives $T_{\rm rh}$ and hence $N^* = 52.6$; the prompt lock crossing guarantees no dilution. Move any link and the observables move together — a diluted window drags n_s below 0.955 as it drags r up, so the two measurements cross-check each other and the framework cannot retreat to a corner of parameter space it does not have. A theory that can be killed by a single clean measurement is doing exactly what a falsifiable theory should. The fragility is the falsifiability, stated honestly rather than hidden behind a soft bound.

Objection 64 (O-Cosmos-64) (the gravitation critic)
"Your primordial-black-hole discussion is a menu of possibilities with no commitment. You list the channels and never predict a mass function. That is phenomenology theater: unfalsifiable because it predicts nothing specific."

Response. Conceded, and stated as such in the text. Primordial black holes in GG-A10-Cosmos are a feature-dependent supplementary channel, not a committed prediction. The reason is structural: the inherited primordial spectrum is red and smooth, so the collapse fraction is exponentially negligible ($\beta \sim e^{-10^3}$ on the minimal branch), and no committed mass function follows. We do not advertise a PBH mass spectrum, and the Part 'Black holes and primordial relics' frames PBHs precisely as a contingent channel that would only switch on if a non-minimal small-scale feature were inserted — which the minimal ledger does not supply. The honest disposition is: minimal branch predicts essentially no PBHs; any non-minimal branch is a separate, clearly-flagged extension. We claim the former, not a mass function.

Objection 65 (O-Cosmos-65) (the gravitation critic)
"You claim to sidestep the singularity theorems by starting the universe at the finite GUT edge of the throat. But Hawking–Penrose forbid exactly that kind of geodesically complete past under standard energy conditions. You have not satisfied or circumvented their hypotheses; you have simply redefined the initial-value problem so the theorems

do not apply. That is changing the subject."

Response. You have characterized our move accurately, and we say the same thing in 'The beginning of the universe at E_{GUT} rather than E_{Pl}': the paper is "not claiming to remove singularity theorems by algebraic trick" but is "changing the admissible initial-value problem." The lawful beginning sits at the finite boundary surface $\Sigma_{\mathrm{GUT}}^{\mathrm{Bdy}}$ at $\mu = E_{\mathrm{GUT}}$, with finite H_* and finite curvature invariants, replacing the formal $a \to 0$ Planck-singularity datum. Whether this genuinely circumvents the energy-condition hypotheses of the singularity theorems, rather than merely declining to extrapolate past the admissible boundary surface, is a foundational question we flag rather than claim solved. The defensible claim is narrow: backward extrapolation below E_{GUT} is not an admissible boundary operation in the GG-6 projection, so the four-dimensional Planck singularity is not the operative initial datum of this effective description. We do not claim to have proven that no completion violates an energy condition; that remains open.

Objection 66 (O-Cosmos-66) (the gravitation critic)
"Your black-hole entropy uses an induced Newton constant. Is that the same G_N A5 derives from the scale overlap, or a different one you quietly assumed to be equal? If they are merely assumed equal, the $A/4$ coefficient inherits an unproven identity."

Response. They coincide by construction, and the identity is now explicit rather than tacit. Both Arisaka (2026, A5-GDOS) (through the T8 chain) and Arisaka (2026, C3-BH) read $1/G_N$ off the same canonical A2 scale profile (Arisaka, 2026, A2-GG6), so the induced Newton constant entering horizon entropy is the scale-derived one, not an independent parameter glued on to match. C3 carries the cross-reference explicitly: the horizon entanglement entropy *is* the induced $A/4G_N$ with that same G_N. And the coefficient itself is the more robust part of the claim — the conical replica fixes $1/4 = 4\pi/16\pi$ from the heat-kernel ratio, which depends on the ratio of overlap integrals, not on the absolute normalization of G_N. So even were the absolute normalization shifted, the $A/4$ coefficient would not move. The identity you worried might be assumed is the one both papers derive from the single shared scale profile.

Objection 67 (O-Cosmos-67) (the gravitation critic)
"Your 'resolution' of the black-hole information paradox is that the four-dimensional observer holds only a CPTP partial-trace image of a closed unitary bulk, so pure-to-mixed evolution is 'unsurprising.' That proves far too little. Any open system has a CPTP image; you have explained why the observer sees thermality, but you have not shown information actually returns, you have not reconstructed the interior, and you have not ruled out genuine information loss. You have renamed the paradox 'openness' and walked away."

Response. A fair and important limit on what the CPTP framing buys. The partial-trace observation does exactly two things, and we should not be read as claiming more. It dissolves the *ontological* version of the paradox: once the exterior observer correctly holds only $\rho_\partial = \mathrm{Tr}\,\Pi_t[\rho_6]$, pure-to-mixed behavior in that reduced state is the generic signature of openness, not a contradiction with global unitarity — because unitarity lives in the closed 6D bulk, not in the boundary image. That is a genuine but modest result: it says the paradox was posed at the

wrong level, and it does not by itself return information, reconstruct the interior, or exclude loss. Information *return* is a separate claim, carried not by the CPTP map but by the saddle-switch structure of 'The Page curve as a GSSB saddle switch', and that structure inherits the same quantitatively-open status as the island/replica program (Objection 39). So the honest stack is: openness explains thermality and removes the false contradiction (proven, modest); the saddle switch describes how the dominant entropy branch turns over (structural, quantitatively open); full interior reconstruction and a rigorous exclusion of information loss are *not* claimed here. The objection is right that openness alone proves little; we agree, and we do not rest the information-return claim on it.

O7 The quantum-field-theory and string-theory critic

Objection 68 (O-Cosmos-68) (the formal-theory critic)
"You call the cosmological-boundary stratum a holographic projection of a 'bulk' and read the radial coordinate as an RG flow, but A10 never writes down a bulk action or a holographic-renormalization scheme. Without a defined action, boundary counterterms, and a renormalized on-shell value, 'warped bulk = holographic RG flow' is a slogan, not a calculation."

Response. The slogan charge is half right, and the honest split matters. A10 itself does *not* carry a worked holographic renormalization; in the Inheritance Part and throughout, the warped-bulk/holographic-RG reading is imported by citation, with the explicit dictionary footnote stating that the radial coordinate $r = \ln(E_{\mathrm{GUT}}/\mu)$ is the renormalization-group axis rather than a Randall–Sundrum spatial dimension. That is a deliberate inheritance, not an original construction here. What blunts the "pure metaphor" version of the charge is that the companion Arisaka (2026, A2-GG6) does supply an explicit classical six-dimensional action on $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$ with the canonical warped metric $ds_6^2 = e^{-2r}\eta_{\mu\nu}dx^\mu dx^\nu + dr^2 + d\theta^2$, a worked Kaluza–Klein reduction with overlap integrals, and a BRST sector satisfying a homogeneous Slavnov–Taylor identity (its Theorem 7). So there is a bulk action and a quantized gauge sector behind the words. What is genuinely still open — and fair to flag — is the last mile: a holographic renormalization with an exhibited counterterm set and a renormalized boundary functional, mapping the bulk on-shell action to the cosmological-boundary observables, is dictionary-by-citation rather than computed. A10 should be read as inheriting a defined bulk and conceding the renormalization scheme as not-yet-worked, not as asserting a holographic duality it has proved.

Objection 69 (O-Cosmos-69) (the formal-theory critic)
"$\Omega_m = 1/\pi$ cannot be a topological invariant. Topological invariants of a four-manifold — Pontryagin numbers, Euler characteristics, instanton numbers — are integers. $1/\pi$ is irrational. Calling it a 'topological invariant' is a category error."

Response. The terminology is conceded, and the codebook already carries the concession. The objection is correct that a characteristic number is integer-valued, and that $1/\pi$ cannot be

one. What the framework actually claims, once the overreaching word is stripped, is narrower: $\Omega_m = 1/\pi$ is a *residual Pontryagin normalization* — the integer matter-cell weight $1/15$ from the OGN-5 ledger, divided by the π-carrying normalization factor that the Chern–Pontryagin integral of $\mathrm{tr}\, R^2$ contributes ($\int \mathrm{tr}\, R^2 = 8\pi^2 p_1$). It is the integer that is the invariant; the π enters through the normalization of the de Rham representative, not as a characteristic number. Accordingly the codebook classifies all three π-locked fractions to Derived-conditional (D*) under Conjecture VIII, states $\Omega_m = 1/\pi$ as a *conjectured* identity rather than a proved invariant, and the codebook carries the honest tag. The headline theorem (T-Cosmos-17) is titled "$\Omega_m = 1/\pi$ as a topological invariant," and that title overreaches relative to the body; the Proof-status note and the codebook are the governing text, and they do not call it an invariant. The critic is right that the word was wrong; the substance the word was reaching for is the residual-normalization reading, which is what is now carried.

Objection 70 (O-Cosmos-70) (the formal-theory critic)
"Theorem T-Cosmos-17 does not close. Step 4 multiplies and divides OGN-5 integers until $1/15$ cancels and $1/\pi$ falls out, but the cancellation needs a normalization you never exhibit. As printed it is algebra arranged to land on a predetermined answer. A theorem whose key step is circular is not a theorem."

Response. The objection is aimed at a version of the honesty note this paper no longer carries, and the current answer is better than a concession. Step 4's cancellation does need a normalization, and §74.3 now says precisely what that normalization is and what became of it: lemma VIII-b′ is proved outright, because it compares *counts* and a ratio of counts is dimensionless at every cutoff; lemma VIII-b″ compares a length to a volume, Eq. (421), and therefore needs a physical unit; and the unit bridge terminates on the dilution no-go, which *proves* that VIII-b″ cannot close as a statement about a constant. What it closes into is the occupancy identity $u = \pi\Omega_m = 1$, true at exactly one epoch. So the theorem is not algebra arranged to land on a predetermined answer: the coupled π-form is forced by the upstream fiber assignment, and no measured density enters the selection at any step. What is not closed is narrower and better posed — that the epoch at which the running inventory reads the form is the present one — and Ω_m, Ω_Λ, Ω_b are D* for that reason and no other. Two further points of honesty are worth recording. First, a false pointer — which sent the reader to "the companion derivation of GG-A4-Euclid" for the integration — has been removed; GG-A4-Euclid supplies the OGN-5 integers, but contains no closed-packet integration, and that stale pointer was deleted. Second, the theorem is stated as a conditional — *if* the cost-minimum integration exists with the exhibited normalization, *then* $\Omega_m = 1/\pi$ — so it is honest about being a structural ansatz with an open antecedent rather than a closed deduction. The critic's diagnosis and the paper's self-assessment coincide.

Objection 71 (O-Cosmos-71) (the formal-theory critic)
"The 'closed-packet four-form integration' is undefined. You never specify the 4-manifold $\mathcal{M}_4$, never name a characteristic class, never define the measure or the integration domain. The operation appears in A10 and nowhere else in the cohort. An undefined

operation cannot ground a derivation."

Response. Conceded, and this is precisely the content of the open Conjecture VIII. The closed-packet integration is *named* — the integral of the BI-6 four-form X_4 over the cosmological-boundary slice $\mathcal{M}_4$, decomposed into OGN-5 cells — but it is not yet *defined* with the rigor a critic is entitled to demand: $\mathcal{M}_4$ is characterized as the four-dimensional boundary slice from the OGGEP projection rather than exhibited as a specific manifold with a specified topology, and the cell-weight assignment is asserted from the OGN-5 ledger rather than constructed from a defined characteristic class. The critic is right that this is why the operation does not yet do rigorous work. It is also exactly why the three π-locked fractions are D^* and not D: the codebook does not pretend the integration is defined. The honest statement is that A10 introduces a named conjectural operation, states what it would have to evaluate to, and defers the definition — the 4-manifold, the class, the measure — to future work. The critic should read the closed-packet integration as the open research target it is labeled to be, not as machinery the paper claims to possess.

Objection 72 (O-Cosmos-72) (the formal-theory critic)
"The OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ is posited. You write it down and use it everywhere, but a tuple of integers asserted by hand is not a derivation, and the cosmic numbers inherit that arbitrariness."

Response. On this point the inheritance is genuinely sound, so the objection misfires — though it is fair to be precise about what is proved and what is not. The OGN-5 ledger is not posited in A10; it is *proved* to theorem status in the cohort. Arisaka (2026, A2-GG6) derives it (its Theorem 6) as a five-step Diophantine cascade closing through the dual-lock identity $K = LJ = M^2 + N^2 + L^2$, and Arisaka (2026, A4-Euclid) promotes the OGN-5 integer cascade to theorem status as part of the zero-free-parameters discipline. A10 inherits these integers; it does not invent them. What A10 itself does *not* close is the next step — the closed-packet integration that would turn the integer ledger $(1, 3, 5; 7, 35)$ into the irrational fraction $1/\pi$. So the honest division of labor is: the integers are theorems upstream; the map from integers to cosmic fractions is the open Conjecture VIII downstream. The critic is right to demand that the ledger be derived rather than asserted, and it is; the residual openness is in the integration, not in the ledger.

Objection 73 (O-Cosmos-73) (the formal-theory critic)
"The π is convention-dependent. Powers of π float around with your choice of volume normalizations — $\mathrm{Vol}(S^3) = 2\pi^2$, factors of 2π in measures, $8\pi^2$ in instanton normalizations. And your 'uniqueness against competing rational functions of π' rules out $1/(2\pi)$ and $1/\pi^2$ by appeal to the data. That smuggles empirical preference into what you call a topological proof."

Response. Half of this is fair and half is aimed at a superseded reading, and the two halves should be separated. *On uniqueness*, the critic is quoting a concession this paper has withdrawn, because the competitors are now excluded *geometrically*, with no appeal to magnitude: $1/(3\pi)$ would require a holonomy volume $6\pi^2$ that no manifold in the construction carries; $1/\pi^2$ would

require an S^5 that the four-dimensional boundary slice does not have; and $1/(2\pi)$ is not a competitor at all — it is the *baryonic* projection, one unit of $U(1)_B$ charge per revolution of the fiber, which is a different quantity from the matter-to-total ratio and appears in the inventory in its own right. None of those three exclusions mentions a measured density. *On the convention-dependence:* the framework's defense is that the closed-packet integration is supposed to fix the normalization *once*, so that the relative powers of π across the three fractions (π^{-1} for matter, $1-\pi^{-1}$ for vacuum, $(2\pi^2)^{-1}$ for baryons) are forced as a coupled set rather than chosen per-entry. That is a real structural constraint — a post-hoc fit tunes one normalization, not a coupled triple. The residual conditionality is therefore not the normalization and not the uniqueness: it is the identification of the crossing epoch with the present one, and that is what the codebook's D* records (§74.3).

Objection 74 (O-Cosmos-74) (the formal-theory critic)
"$M_{\rm lock} \approx 3.85$ TeV is reverse-engineered. You take $\frac{1}{2}M_{\rm Pl}e^{-35}$, land near a few TeV, and call it a compactification scale. That is dimensional analysis run backwards from a desired answer, not a derived KK threshold."

Response. The phenomenology of $M_{\rm lock}$ is not A10's remit, and the reverse-engineering charge is best answered where the work is done — the companion Arisaka (2026, P1-KK). A10 inherits $M_{\rm lock} \approx 3.85$ TeV only as the top of the radial chronology, the inherited bulk modulus scale of the descent's second saddle — the prompt lock crossing where the dark sector is born, not the inflationary plateau, which sits one e-fold below the GUT edge (Part IV); it does not itself derive the lock spectroscopy. Arisaka (2026, P1-KK) carries the structure: $M_{\rm lock} = \frac{1}{2}M_{\rm Pl}e^{-35}$ is fixed by the OGN-5 integer $K = 35$ (a theorem-level integer, not a scanned parameter), the factor of $\frac{1}{2}$ is the scale half-wave of the radial standing mode rather than a fudge, and the same $M_{\rm lock}$ over-constrains the electroweak vacuum value $v = M_{\rm lock}(2/5)^3 \approx 246$ GeV — an over-determination that a reverse-engineered single number could not produce. That over-constraint is the structural answer to "dimensional analysis run backwards": one cannot reverse-engineer a single scale to hit two independent targets (~ 3.85 TeV at the LHC *and* 246 GeV at electroweak breaking) unless the relation is forced. The honest qualification is that the integer $K = 35$ is the input, and whether $K = 35$ itself is forced rather than chosen is the OGN-5 ledger question answered upstream in Arisaka (2026, A2-GG6); the falsifiability is sharp, since Arisaka (2026, P1-KK) stakes the implementation on a clustered spin-2 signal within 3.85 ± 0.10 TeV.

Objection 75 (O-Cosmos-75) (the formal-theory critic)
"Anomaly freedom and the four-form coefficient vector $(1, -1, -2/5, -2/5)$ are imported, not checked. You quote the modified Bianchi identity and the $-2/5$ Abelian fractions as if they were established, but A10 performs no anomaly computation. How do I know the bulk on which your whole cosmology rests is even consistent?"

Response. Here the inheritance is genuinely solid, and A10 correctly imports rather than re-derives. The BI-6 coefficient vector $(1, -1, -2/5, -2/5)$ and its anomaly freedom are *derived* in Arisaka (2026, A3-BI6), not asserted in A10. The $-2/5$ Abelian fractions follow from an

explicit trace computation — the Georgi–Glashow embedding gives the trace ratio that pins both Abelian coefficients symmetrically to $-2/5$ via the trace-ratio lock $|a/c| = |a/d| = 5/2$ (equivalently $(L - N)/L = 2/5$ from the OGN-5 integers), so no coefficient is free to adjust. Arisaka (2026, A3-BI6) proves both the local anomaly-cancellation theorem and the global anomaly-freedom theorem, with the single-tensor Green–Schwarz factorization $I_8 = \frac{1}{2}X_4 \wedge X_4$ supplying the same X_4 four-form that A10 then integrates. A10's role is to inherit this consistent, anomaly-free bulk as the parent law of the cosmological boundary, and to use the BI-6 four-form construction as the object whose boundary integral is conjectured to yield $1/\pi$. The critic's demand — show me the bulk is consistent — is answered, but in Arisaka (2026, A3-BI6); A10 is right not to repeat the anomaly computation, and right to cite it as the upstream theorem rather than claim it as its own.

Objection 76 (O-Cosmos-76) (the formal-theory critic)
"Naturalness is never addressed. A TeV-scale lock and a meV-scale vacuum energy both sit far below the cutoff, and you say nothing about whether they are radiatively stable. Without a stability analysis your 'derived' scales are exposed to the same fine-tuning problems as everything else — worse, because you claim they are fixed."

Response. Conceded, and flagged open. A radiative-stability analysis of M_{lock} and of the meV-scale vacuum residual is not in the present scope of A10, and the paper does not claim one. This is a real gap, and an honest one to state plainly: showing that $M_{\text{lock}} \approx 3.85$ TeV is protected against large radiative corrections, and that the $\Omega_\Lambda = 1 - 1/\pi$ vacuum fraction is radiatively stable rather than merely posited at the right value, would require a renormalization-group and effective-potential treatment that the framework has not carried out. Two mitigating remarks, offered without overclaiming. First, the framework's stance is that the scales are fixed by integer ledger data ($K = 35$ for the lock, the closed-packet topology for the vacuum fraction) rather than by a balance of continuous parameters, so the framework's *hope* is that there is no continuous dial to detune — but a hope is not a stability proof, and the radiative corrections to a fixed structural value still need to be shown to be controlled. Second, the absence of this analysis is recorded as an open item rather than papered over. The critic is right that a derivation claiming fixed scales owes a stability argument, and the honest status is that this argument is owed and not yet paid.

O8 The methodology, honesty, and numerology critic

Objection 77 (O-Cosmos-77) (the methodology critic)
"The whole program is numerology. You pick pretty numbers — $1/\pi$, $1/(2\pi^2)$, e^7 — and then build topological language around them to rationalize the choice."

Response. The discipline against numerology here is structural, not rhetorical. First, every cosmological observable is triaged in the Cosmological codebook by a published provenance code — Derived (D), Derived-conditional (D*), or Asserted (A) — so that a reader can see at a glance which numbers are claimed as outputs and which are still owed; twenty-one entries, one D,

nineteen D*, one A, with no entry quietly upgraded. Second, the conditional entries are not vague promises but nineteen closures resting on ten *named, falsifiable* conjectures, the dark-matter row on an eleventh besides, the companion's inflation-era pinning (closed-packet-integration, vacuum-residual-exactness, Kaxion-misalignment-completeness, the stratum assignment of the vacuum residue, and the rest), each of which can be disproved by a counterexample. Third, the Post-hoc-fit defense section spells out why one fixed topology forces three coupled fractions at once: a curve-fit explains one number, but the same closed-packet integration that yields $\Omega_m = 1/\pi$ yields $\Omega_\Lambda = 1-1/\pi$ and $\Omega_b = 1/(2\pi^2)$ with no further free step, and the $2\pi^2 = \mathrm{Vol}(S^3)$ factor in the baryon fraction is the wrong combination of powers to be a coincidence. Fourth, and decisively, the same six-dimensional geometry fixes both the particle inventory of (Arisaka, 2026, A9-SM) and the cosmological inventory of this paper with zero continuous dials: a structurally locked derivation cannot be retrofitted, so a single wrong number kills it. That is the opposite of an adjustable fit. The honest concession is only that several of the ten conjectures remain unproved — which is why those rows read D*, not D.

Objection 78 (O-Cosmos-78) (the methodology critic)
"Every empirical comparison in the paper is central-value-to-central-value. There is no joint likelihood, no covariance, no real statistical confrontation with the data."

Response. Conceded in part, and stated openly in the text. The codebook claims are at the level of central values and summary statistics: each entry is compared to a published central value with its quoted error, and the parameter-audit Dev plot displays fractional deviations with empirical and, where applicable, theory error bars, the σ_{th} column. The Boltzmann computation has been run and the likelihood has not: the zero-parameter TT spectrum from the derived parameters gives $\chi^2 = 187$ over the 83 binned Planck points with diagonal errors (Section 40), a demonstration rather than a fit. A global CMB-plus-BAO-plus-distance-ladder fit is named explicitly as required future work and is the dedicated remit of the companion (Arisaka, 2026, C2-Hubble). Until that pipeline exists, the right reading of, say, the 0.27σ agreement on Ω_{DM} or the 2.3σ tension on Ω_b is as a per-parameter audit against marginal posteriors, not as a likelihood ratio for the model as a whole. We do not claim the latter and have flagged the gap rather than papered over it.

Objection 79 (O-Cosmos-79) (the methodology critic)
"'Derived (conditional)' is just 'not derived' with better marketing."

Response. The D* code has a precise and auditable meaning: the entry follows from an *explicitly named, falsifiable conjecture* (two for the dark-matter row), and each conjecture is printed — not hidden in prose, but listed by name in the codebook row and in the conjecture ledger. The reader can see exactly what is owed: $w = -1$ is D* on the vacuum-residual-exactness conjecture; Ω_{DM} is D* on Kaxion-misalignment-completeness together with the companion's pinning conjecture C-Kaxion-d; and so on. This is strictly more honest than an unlabeled assertion, because it converts "we believe this is derivable" into "this is derivable *if and only if* the following stated proposition holds, which you may now attack." A genuinely empty D* would be one where the conjecture is unfalsifiable or unstated; ours are neither. The label is a

debt instrument with the amount written on it, not a euphemism for an unpaid bill of unknown size.

Objection 80 (O-Cosmos-80) (the methodology critic)
"You insist no number was tuned after the fact, yet the Dev plot is full of entries that sit exactly at zero deviation. That is tuning by construction."

Response. The Dev plot is built precisely to make that distinction visible rather than to obscure it. Filled markers are independent predictions — numbers computed without reference to the measured value (the atomic-chain $z_{\rm dec}$, $w = -1$, the Boltzmann ℓ_D, the π-locked fractions, η_B) — and these carry genuine, nonzero deviations: Ω_b at $+2.76\%$, $z_{\rm dec}$ at $+0.15\%$ (the switch e^7, a distinct topological marker, annotated 0.62% above the measured value), w near $+2.7\%$, the Boltzmann ℓ_D at $+0.4\%$. The one open marker is the entry anchored by construction: the amplitude A_s, which pending Lemma V′ echoes the empirical input and *necessarily* sits at zero deviation, and the caption says so in as many words. (Earlier versions anchored the CMB Hubble reading, the early S_8, the tilt and $N_{\rm eff}$ as well; all four are now computed and scored, the tilt at -0.4σ against *Planck* and 3.2σ below the DESI-combined value.) Far from hiding the trivial zeros, the plot color-codes and shape-codes them so the reader never mistakes an anchor for a hit. The honesty of the figure is exactly that it refuses to count the open markers as successes.

Objection 81 (O-Cosmos-81) (the methodology critic)
"The celebrated '0.6%' success of $z_* = e^7$ uses a generous denominator. Measured against the actual Planck error, that offset is a large tension, and you quietly chose the kinder yardstick."

Response. The premise is retired: the paper no longer scores e^7 against the recombination redshift under *any* denominator, generous or harsh. The two-quantity resolution (§42.1) separates what that comparison conflated. $z_{\rm sw} = e^7$ is the topological corridor switch — an integer-ledger marker, not a prediction of the measured decoupling redshift — so neither the visibility width nor the Planck error is its yardstick; reaching for the kinder of the two would indeed be comparator-shopping, which is exactly why the response to Objection 23 (the "27σ" objection) declines that move. The quantity that *is* scored against Planck is the framework's own atomic-chain decoupling prediction, $z_{\rm dec}^{\rm GG} = 1091.5$ (raw, at the framework's own h), which lands on $z_{\rm dec} = 1089.92 \pm 0.25$ at 0.15% with zero fitted parameters — inside the 0.18% modeling systematic of the recombination calculation itself, which is that comparison's honest yardstick. The switch–decoupling offset is a computed number, $\delta = 0.0038$ (Eq. (262)), not an error bar to be gamed. The both-framings presentation your objection quotes survives only in the thermal appendix, explicitly marked there as the pedagogical bridge from the superseded one-quantity reading to the present form. What remains of the objection is a comparison the paper no longer makes.

Objection 82 (O-Cosmos-82) (the methodology critic)
"The eight 'falsification tests' are theater. Several lie beyond any foreseeable instrument,

and the strong modules are propped up next to weak ones to borrow their credibility."

Response. The asymmetry between strong and weak modules is real and is stated, not smoothed over. The strongest single kill-switch is now *two-sided* and within reach, which strengthens the objection's own premise against it: the derived plateau fixes $r = 0.0038$ exact ($12/N^{*2} = 0.0043$ at leading order), a factor nine below the current bound and several times above the Simons-Observatory/LiteBIRD floor, so a measurement inconsistent with it in *either* direction — a detection well above, or an exclusion well below — refutes the derived launch. This replaces the retired one-sided $r \ll 10^{-3}$ reading, which was indeed beyond reach; the new target is decisive this decade. The near-term, genuinely exposable tests are different in character: a high-precision independent recombination-redshift determination, the BAO dilation, predicted flat at the background r_d, and the GW-siren Hubble determination, predicted to converge on the background at $H_0^{\rm GW} \approx 67$. These confront the derivation on the next-decade horizon: a siren tracking the local ladder (≈ 73) rather than the background would refute the π-inventory value. We do not claim every module is equally weak-point-free; the baryogenesis entry is order-of-magnitude only (a factor-$\sim$ 3 theory band), and the local-ladder excess above the background is not derived at all — it is read whole as a calibration systematic the framework predicts (§42.6), and the ladder-free probes test the prediction. Listing eight tests of uneven decisiveness, with the weak ones labeled weak, is the opposite of theater: it is an inventory of where the program is most and least exposed.

O9 The cohort-coherence and philosophy-of-cosmology critic

Objection 83 (O-Cosmos-83) (the cohort-coherence critic)
"The cohort cross-cites in circles. A10 leans on C2 for the Hubble resolution, and C2 leans straight back on A10. The agreement is manufactured, not independent."

Response. The loop is real and we acknowledge it rather than disguise it; the defense is a clean division of labor, stated in the Cross-paper coherence map. A10 owns the *structural claim*: the named theorem that one cosmological-boundary Friedmann background, with its switch event at recombination, is read through two dictionaries and by every ladder-free probe, and that the ladder's higher value is a systematic. The companion (Arisaka, 2026, C2-Hubble) owns the *dedicated phenomenology and the eventual joint fit*: the ladder-channel falsifiers and the global likelihood that would turn the reading into a quantitative resolution. The two papers say the same thing: one derived rate, 66.8, and the ladder's excess read whole as a calibration systematic, priced in the companion against the published corrections and found to exceed even their fully aligned sum. The mutual citation is therefore a pointer to who carries which half of the argument, not a manufactured confirmation: neither paper claims the tension is *resolved*, and both flag the joint fit as still required.

Objection 84 (O-Cosmos-84) (the cohort-coherence critic)
"A10 keeps deferring the hard parts to companion papers — the BBN lithium network, the Kaxion mass cap — and at least one of those companions, the nucleosynthesis

paper, does not even exist yet."

Response. Conceded, and labeled as such throughout. The lithium-7 network closure is explicitly future work: the BBN audit proposes a selective $A = 7$ open-system port and lists the quantitative network calculation as an open item (PS-Cosmos-15) — no companion in the cohort carries it, and none is cited as though one did. The Kaxion mass cap is the cleaner case: the compact-angle-mass-uniqueness conjecture was an open D* in this paper, and the closure has *already* been executed in the existing companion (Arisaka, 2026, C1-Kaxion), calibrating the one scale the conjecture leaves open, declaring the calibration, and delivering the 8.1 μeV line on a ladder of seven rungs; the remainder — the orientation convention behind the calibration and a defining computation for the ladder — is named there, not hidden. So the deferrals are of two kinds, both honestly marked: those discharged by a companion that exists, and those carried on this paper's own open list (PS-Cosmos-15), owed to no companion. What the paper does not do is count an open item's result as a present accomplishment.

Objection 85 (O-Cosmos-85) (the philosophy-of-cosmology critic)
"The 'universal-cosmic-identity, layer-by-layer from electrons to ATP to cells to cosmos' framing is not physics. It is a grand narrative that weakens, rather than strengthens, the formal case."

Response. Conceded as a matter of status, and that is exactly why it is fenced. The universal-cosmic-identity framing is an interpretive synthesis, presented as a preview and explicitly reserved as downstream of the calculations — it lives in future work and in the capstone (Arisaka, 2026, A11-Origin), not in the codebook or the headline theorems. The cosmological results of this paper — the π-locked inventory, $z_* = e^7$, the corridor, the Kaxion sector — stand or fall entirely on their own derivations and falsification tests, with no load borne by the layer-by-layer narrative. We keep the synthesis because we think it points somewhere, but we draw the line so a reader who rejects it loses nothing of the physics. The narrative is offered as interpretation labeled as interpretation, which is the only honest way to offer it.

Objection 86 (O-Cosmos-86) (the philosophy-of-cosmology critic)
"Your contrast with the inflationary multiverse and anthropic selection is just rhetoric — a way of scoring points against a rival picture rather than a substantive argument."

Response. The contrast is methodological, and it cashes out in falsifiability rather than in tone. The substantive claim is structural: a zero-dial geometry, fixed once with no continuous knobs and no ensemble to choose from, is *obliged* to produce exactly one definite value per observable, and is therefore exposed to refutation by any single wrong number — it derives $\Omega_m = 1/\pi$ rather than measuring it. An anthropic ensemble runs the other way: it lets the constants vary across an unobservable population and selects the subset compatible with observers, which accommodates a wide range of values and so predicts none of them sharply. That is not a rhetorical preference; it is the difference between a framework that one wrong number can kill and a framework constructed so that almost no measurement can. ΛCDM, for its part, neither derives nor selects — it measures its six parameters. The three-way contrast is about closure mode and falsification profile, which are the properties a critic should care about.

Objection 87 (O-Cosmos-87) (the philosophy-of-cosmology critic)
"Leaving 'the final metaphysical step to the reader' is evasive. If the framework implies a necessity, say so; if it does not, stop gesturing at one."

Response. The restraint is deliberate, and it is the honest position, not an evasive one. The paper draws the derived/interpreted line explicitly and repeatedly: the codebook marks what is an output (D), what is conditional (D*), and what is asserted (A); the synthesis sections mark themselves as interpretation. Where the framework genuinely implies a value, it asserts it and exposes it to falsification — it does not hedge $\Omega_m = 1/\pi$ or $z_* = e^7$. What it declines to assert is *metaphysical necessity* — the claim that the universe *could not have been otherwise* — because that is a step the calculations do not establish, and asserting it anyway would be the actual overreach. Declining to claim what one cannot prove, while clearly marking the boundary, is intellectual restraint. The alternative the objection seems to want — a confident metaphysical pronouncement — is precisely what an honest paper should refuse.

Objection 88 (O-Cosmos-88) (the cohort-coherence critic)
"The whole inheritance chain rests on companions that are themselves preprints no one outside the program has reviewed. A theorem chain is only as sound as its weakest link, and every link is provisional."

Response. True, and it is the correct way to read the cohort: the chain inherited from (Arisaka, 2026, A1-Constitution), (Arisaka, 2026, A4-Euclid), (Arisaka, 2026, A8-QM), (Arisaka, 2026, A9-SM), and the capstone (Arisaka, 2026, A11-Origin) is a chain of preprints, and the present paper's conditional results are conditional on them in the strict logical sense. We make no claim that the upstream results carry an outside reviewer's authority they do not yet have. Two things make this honest rather than fatal. First, the dependence is explicit: each cosmological entry carries its upstream citation chain (e.g. the OGN-5 ledger to (Arisaka, 2026, A4-Euclid), the baryonic protection to (Arisaka, 2026, A9-SM)), so a reader who doubts a link can trace exactly which downstream results fall with it. Second, the conditionality is structural, not hidden: this is what the D* code and the named conjectures encode — "true if the stated upstream proposition holds." A theorem chain of preprints that labels every dependency and never asserts more than its links support is doing the most a preprint cohort can honestly do. The remedy is the ordinary one — the links get tested, or they fail and the dependent rows fail with them, exactly as the provenance triage already advertises.

Objection 89 (O-Cosmos-89) (the philosophy-of-cosmology critic)
"Your 'second horizon problem' is not new. Zeldovich and Kibble posed the random-vacuum-choice problem in the 1970s; Linde listed the uniqueness of the universe among inflation's targets and even asked why all properties of the universe are uniform; and axion cosmology has always known the misalignment angle is uniform across our patch because inflation stretched one patch. You have renamed the domain problem."

Response. The pedigree is real, and the paper cites it in exactly those terms (§72.2) — Zeldovich–Kobzarev–Okun and Kibble for the sharpest early fragment, Linde's "uniqueness of the universe" and his "all properties" remark as the nearest miss in ninety years, the axion literature for the

pinning mechanism stated almost verbatim. What is claimed as new is deliberately narrow, and each piece can be checked against the quoted sources. First, the naming and the ranking: the domain problem was posed as a relic nuisance and solved by dilution; nobody ranked the sameness of things as the *harder* half of the 1981 question, parallel to and deeper than temperature. Second, the extension: the 1970s fragment concerns the orientation of one degenerate vacuum inside one fixed theory; the second horizon problem asks about the whole inventory — masses, couplings, abundances, and the angle at once. Third, the temporal answer: the field's answer has always been spatial (we live deep inside one stretched domain); the ordering — inflation preceded the existence of matter, so identity was inherited rather than negotiated — has a universally known premise and, as far as we can find, a nowhere-drawn conclusion. Fourth, the reduction: cosmic identity and vacuum uniqueness are shown to be one puzzle, not two. Renaming the domain problem would be claiming the question; what this paper claims is the shape of the question — and the half of the answer that is temporal rather than spatial. A critic who finds any of the four steps in print is invited to send the reference; the pedigree subsection was written to make that check easy.

Objection 90 (O-Cosmos-90) (the formal-theory critic)
"You keep contrasting your unique vacuum with the string landscape, yet you never say the landscape is refuted. If the second horizon problem is as serious as you claim, why not draw the obvious conclusion and declare the 10^{500}-vacua picture dead?"

Response. Because refutation is the wrong verdict, and the accurate one is stronger. Two objects travel under one name: string theory as a candidate quantum-gravity *formalism* — worldsheet finiteness, the dualities, the AdS/CFT correspondence — and the *landscape* as a vacuum-selection program. Nothing here bears on the formalism, whose Green–Schwarz anomaly mechanism this cohort in fact adapts to six dimensions as BI-6 (GG-A3-BI6); the remarks concern only the second.

On the second: the observed identity of the low-energy constants across causally disconnected regions of our own past light cone is a datum — the second horizon problem (§72). A framework with a *unique* vacuum explains that datum as necessity, the same constants everywhere because there is nowhere else for them to be, and is thereby *exposed*: a single confirmed spatial variation of a dimensionless constant such as α or m_e/m_p would falsify it outright. A framework with of order 10^{500} vacua and no dynamical selection principle (Douglas, 2003) cannot supply that explanation as necessity, because the vacua are vacua of the theory meant to describe *our* universe — so within a single horizon volume the agreement of the constants is a physical question, not a metaphysical one about other worlds. The available answer is that our whole observable patch descends from one Coleman–De Luccia nucleation event (Coleman & De Luccia, 1980) in an eternally inflating background: internally consistent, since tunneling between high-barrier vacua is exponentially suppressed and domain walls inside our horizon are generically negligible, and therefore *not contradicted* by the datum.

The issue is not contradiction but cost. That repair commits the program to eternal inflation, and with it to the measure problem, in which the relative frequencies of observations are undefined without an additional, non-derived regulator (Freivogel, 2011); predictivity for any

further quantity is then not recovered but deferred to that choice. The one standing quantitative success of the anthropic reading, Weinberg's bound on the cosmological constant (Weinberg, 1987), remains a single post-diction. In the language of research programs (Lakatos, 1978) the picture is best described not as refuted but as *non-progressive*: each datum accommodated by adjustment, none forecast. GG-Theory takes the opposite structural wager — one vacuum, pinned upstream by the OGN-5 ledger, obliged to return one value for each constant and holding no freedom to absorb a discrepant measurement. We record the difference as one of closure mode — obligation versus accommodation, one measurement decides versus unfalsifiable by construction — and leave its weight to the reader.

Objection 91 (O-Cosmos-91) (the model-building critic)
"Your leptogenesis assumes exactly what every model-builder assumes: singlet Majorana neutrinos and a type-I seesaw. Dressing the standard choice in geometric language is not a derivation."

Response. It was an assumption until Proposition 36.2 (T-Cosmos-18); it is now fiber arithmetic. A Majorana self-pairing exists iff the fiber weight vanishes — on every charged sector the Majorana bilinear carries weight $2n$ and is forbidden, so lepton number is exact throughout the dial-carrying ledger and the sole $\Delta L = 2$ door is the one field the circle does not move. The LHN portal is likewise the unique fiber-neutral trilinear touching that door, and the heavy scale is not chosen either: the throat integer $K = 35$ pins the triplet to $M_{\rm lock} \simeq 3.85$ TeV with $O(1/K)$ splittings. What remains genuinely open is stated in the module: the detailed values of the splittings δ_i within their geometric window. The singlet's real, dial-free character is developed as entry R10 of the real-bulk series of Arisaka, 2026, A8-QM.

Objection 92 (O-Cosmos-92) (the curved-space critic)
"There is no preferred vacuum in an expanding universe. Your particle inventories — the N_i abundances, the thermal histories, the relic numbers — are convention-dependent quantities dressed as predictions."

Response. The ambiguity is real and has a fixed address (§12.1). On the canonical bulk geometry the constitutional time is globally static and the bulk-level positive-frequency structure is unique; observer-adapted vacua are *dictionary* data — which modes an instrument counts is boundary information, part of the OGGEP package by definition — and a change of dictionary is a Gaussian channel acting on the boundary description, not an ambiguity in the bulk state. Every inventory quoted in this paper is computed in the comoving dictionary, consistently, and the thermal content attributed to horizons is produced by the GDOS partial trace itself, at the standard temperature. What remains conjectural — the factorization for strongly gravitating bulk configurations — is named in the real-bulk treatment of Arisaka, 2026, A8-QM rather than waved through here.

Objection 93 (O-Cosmos-93) (the CP-provenance critic)
"Your baryon asymmetry rests on a CP phase you cannot derive; 'geometric phase' is a euphemism for a fitted parameter."

Response. The CP-odd inputs of this module are inherited, real, lattice-locked angles, derived

upstream and quoted here in completed tense: the quark phase $\delta_{\mathrm{CKM}} = 2\pi/5$ at first order (two pentagonal steps on the single compact circle, with no second flavor torus on which a separate angle could float), and the leptonic set $\delta_{\mathrm{PMNS}} = \pi + 2\arcsin(1/7)$, $\alpha_{31} = \pi$ (Arisaka, 2026, A9-SM). Their rephasing invariance is holonomy invariance: in the real presentation a physical phase is a closed-loop angle of the coupling texture — the flavor analog of an Aharonov–Bohm angle — which no per-family rephasing can remove. At the resonance the Yukawa magnitude cancels and $\epsilon^{\mathrm{res}} \simeq \frac{1}{2}\sin(2\phi_{\mathrm{geo}})$ is controlled by these angles alone. The falsifier is correspondingly sharp: a measured drift of the flavor phases away from the lattice values retires the account — there is no continuous dial here to absorb one.

O10 The curious reader returns: the multiverse and the three horizon problems

The critic sections above interrogated the machinery. The four questions below return to where this Part began — the multiverse — because the reader who has now met the second and third horizon problems (§72, §73) owns a sharper instrument than the one Objections 1 and 2 could assume, and deserves to see what it does and does not cut.

Objection 94 (O-Cosmos-94) (the curious reader, returning)
"You derive three horizon problems — the same temperature, the same electrons, the same choice of matter over antimatter, across regions that never exchanged a signal. If everything is supposed to vary from universe to universe, doesn't that observed sameness falsify the multiverse outright?"

Response. It is tempting to say yes, and the temptation should be resisted, because the truth is sharper. The landscape does not contradict the three samenesses; it accommodates them, and it does so using the same mechanism this paper uses: the choice of vacuum and of CP sign is made once, early, and inflation then stretches that single decided patch into everything anyone can see (§72, §73). The variety the multiverse trades in is stored outside the horizon, exactly where no measurement reaches — so the observed sameness cannot execute the hypothesis, because nothing observable could. What the three problems *do* kill is every mechanism in which the choosing happens late, after the regions have fallen out of contact: such a sky would carry domain boundaries and matter–antimatter annihilation seams, and it carries none (Cohen et al., 1998). That fence forces every framework, string theory included, into the early-choice corner — and surviving there has a price. In that corner the landscape can no longer explain why *this* vacuum, *this* sign, *this* yield; it can only report that each happens somewhere, and hand the remainder to observer selection (Objection 12). The right verdict is therefore not falsification but exile: the sameness of the sky drives the multiverse to a place where it can neither explain nor be tested. This framework's answer lives inside the horizon instead — determinism plus inheritance — where the falsifiers can reach it.

Objection 95 (O-Cosmos-95) (the curious reader, returning)
"If the multiverse cannot be falsified, is it really worse off than your 'uniqueness'? Both sound like bets about things nobody will ever see."

Response. The two bets have opposite exposures, and that is the whole difference. Uniqueness is indeed a conjecture (Objection 2 marks it so), but it is a conjecture that stakes its life on numbers inside our horizon: an extra light family, cosmic shares settling away from the π-locked fractions, a tensor ratio far from 0.0038, a decoupling redshift off the atomic chain — any one of these kills it (Objection 22). Plurality stakes nothing comparable: whatever is measured here is compatible with some address in the ensemble, and the ensemble itself is beyond reach. So one hypothesis about the unseen forbids things in the seen, and the other forbids nothing anywhere. Both are bets about places nobody will visit; only one can lose at home. A theory that cannot lose is not winning — it has merely stopped playing. The asymmetry is not rhetoric; it is the operational content of the word *scientific*, and it is the reason this paper spends its pages staking falsifiers rather than multiplying worlds.

Objection 96 (O-Cosmos-96) (the curious reader, returning)
"But 'every possibility happens somewhere' does cover the facts. Why do you refuse to count that as an explanation?"

Response. Because it has the same logical shape as a claim this framework also declines, from the other side of the aisle: *the world was designed*. Both statements are compatible with anything a telescope could ever find. Neither names a step that could have failed, forbids a measurement, or risks anything on tomorrow's data: whatever is observed, the designer could have designed it, and the ensemble could have contained it. That is precisely why neither functions as an explanation of the particular world in front of us. An explanation, in the sense this paper tries to practice, is a chain with exposed links: name the mechanism, compute the number, state what reading of what instrument breaks the chain. The same razor cuts both comfortable alternatives, the theological and the statistical, and for the same reason. That it cuts symmetrically is worth noticing: this framework is not the design hypothesis wearing equations (Objection 5), and its refusal of the multiverse is not a matter of taste. It declines both because both are, in the exact and technical sense, unlosable — and it accepts the discipline of being losable instead.

Objection 97 (O-Cosmos-97) (the curious reader, returning)
"Is there anywhere the two pictures meet the same measurement and predict differently — an experiment, not a philosophy seminar?"

Response. One, and it is dated. This paper's launch is a derived plateau with tensor-to-scalar ratio $r = 0.0038$ (Part IV), squarely in the range the Simons Observatory and LiteBIRD are built to decide this decade. Inside the string program there is an actively debated internal criterion — the "de Sitter swampland" conjecture — whose strong forms disfavor exactly the flat, positive-energy plateaus that drive inflation of this type. The criterion is contested within that community, so a detection would not refute string theory; conjectures can be dropped. But the situation is asymmetric in an instructive way. A measurement of r near 0.0038 would confirm a two-sided prediction this framework cannot adjust, while landing on the disfavored

side of a criterion the other program proposed for itself. A measurement far from 0.0038, in either direction (beyond the launch's own spread, 0.0037–0.0040), kills this paper's launch outright — and inconveniences the landscape not at all, since there is always another vacuum. One program has placed a bet that a named, funded instrument settles; the other retains the right to reinterpret the outcome. The reader may decide which posture better deserves the word *prediction* — and should note the date: this is not a wager that can be left on the table indefinitely.

Part XVII
References

Cited works, each with a relevance note linking it to the present paper

---

Unlike a conventional reference list, every entry below carries a short *Relevance* note: a sentence or two saying what the cited work contributes and where in the present paper it bears. The aim is to turn the bibliography from a bare list of sources into a navigable companion — a reader can move through the references and see, at a glance, which derivation, table, or section each one underwrites.

The notes are written for two readers at once. For the newcomer, they trace the intellectual lineage of the present derivation: from the ΛCDM concordance model and the cosmological-constant problem, through inflation, primordial nucleosynthesis, and the growth of structure, to the present-day H_0 and S_8 tensions — and then to the twenty-two-paper GG-Theory cohort, in which GG-A10-Cosmos reads cosmology as the boundary projection of the six-dimensional bulk, one story above the Standard Model of GG-A9-SM. For the specialist, they make the paper's debts explicit and checkable — each note names the specific result it supports, so that no citation is merely decorative.

The GG-Theory cohort papers (GG-A1-Constitution–GG-A11-Origin, together with the phenomenology and foundations P-series and the consequence C-series) are cited throughout under the canonical keys and titles fixed by the cross-paper coherence map of the program; GG-A10-Cosmos's place in that cohort — the cosmological-boundary member of the 4D-Physics stratum, the story above the Standard Model — is set out in the cohort overview of this paper. The references are listed alphabetically by first author.

A. Papers by Katsushi Arisaka (the GG-Theory cohort)

K. Arisaka. *The GG-Constitution: Four Axioms and Seven Postulates of the GG-Theory Program, From Which GG-6, BI-6, OGN-5, and GDOS Descend as Theorems*, Arisaka-GG-A1-Constitution. Preprint, available at https://preprints.arisaka-gg.org/a1-constitution/. doi:10.5281/zenodo.22063867.

Relevance. Supplies the constitutional charter — the four axioms, eleven postulates, admissibility rules, and OGGEP semantics — that GG-A10-Cosmos inherits as its first cause. AX-4 MSMF (no hidden continuous moduli) and POS-6 OGN are the upstream clauses that force the zero-free-parameter cosmic inventory; OGGEP is the boundary-projection law that the present paper specializes one story up, at the cosmological boundary, to derive the modified Friedmann dynamics and the π-locked split $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$. Invoked throughout as the legal source of every admissibility argument in the cosmology paper. First cited in §6.

K. Arisaka. *GG-Theory with GG-6: Geometry–Gauge Locking at the 6D Bulk (GG-6) at the Heart of Grand Geometric Theory (GG-Theory)*, Arisaka-GG-A2-GG6.
Preprint, available at https://preprints.arisaka-gg.org/a2-gg6/.
doi:10.5281/zenodo.22064298.
Relevance. Constructs the 6D bulk geometry $\mathcal{M}_4 \times I_r \times S^1_\theta$ on which the entire GG-Theory program runs, together with the BI-6 closure, the OGN-5 integer ledger, and the compact-holonomy ontology. GG-A10-Cosmos descends directly from this stage: the cosmic bulk geometry, the radial chronology axis I_r along which the one background's two dictionaries are read, and the topologically protected S^1_θ winding that hosts the Kaxion compact-angle pseudoscalar all originate here. First cited in §9.1.

K. Arisaka. *Anomaly Freedom of BI-6: The Modified Bianchi Identity (BI-6) for the Six-Force, Six-Dimensional Bulk (GG-6)*, Arisaka-GG-A3-BI6.
Preprint, available at https://preprints.arisaka-gg.org/a3-bi6/.
doi:10.5281/zenodo.22064300.
Relevance. Establishes the complete BI-6 legality theorem — the Green–Schwarz-modified closure identity $dH_3 = X_4$ that fixes the unique anomaly polynomial of the 6D bulk. GG-A10-Cosmos cites GG-A3-BI6 to certify that the cosmological use of projected Ward identities, anomaly inflow, and global consistency is admissible: the boundary cosmology inherits a consistent, anomaly-free bulk, which is what licenses the OGGEP-port treatment of horizons and the closure of the boundary stress ledger. First cited in §9.1.

K. Arisaka. *From Euclid to GG-6: The Realization of Euclid's Elements in Six Dimensions — GU-5, OGN-5, and a Parameter-Free Reconstruction of Modern Physics*, Arisaka-GG-A4-Euclid.
Preprint, available at https://preprints.arisaka-gg.org/a4-euclid/.
doi:10.5281/zenodo.22064302.
Relevance. Fixes the zero-free-parameter Euclidean input ledger — the GU-5 dimensionful constants $(c, G_N, \hbar, k_B, Q_B)$, the structural constants π and e, and the proved OGN-5 integer ledger. The π-locked cosmic inventory of the present paper ($\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$) ultimately rests on this discipline: π enters the inventory as a topological input certified here, not as a fitted number, so GG-A10-Cosmos introduces no free continuous parameter beyond the cohort ledger. First cited in §9.1.

K. Arisaka. *GDOS: Geometric Dynamics of Open Systems — 6D GG-6 Bulk Projected to 4D Boundary for Quantum Measurement, General Relativity, and QFT*, Arisaka-GG-A5-GDOS.
Preprint, available at https://preprints.arisaka-gg.org/a5-gdos/.
doi:10.5281/zenodo.22064304.
Relevance. Provides the exact boundary-reduction engine of the cohort: the GDOS construction of the canonical boundary algebra on $\partial\mathcal{M}$ and the open-system dynamics that follow from tracing the bulk down to its boundary shadow. GG-A10-Cosmos inherits its 4D observables, its cosmic state-manifold evolution, and its open-system arrow of time from GDOS; the modified Friedmann equations of the present paper are the cosmological-stratum specialization of the GDOS boundary dynamics under OGGEP. First cited in §6.

K. Arisaka. *The GRIP on GDOS — GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind*, Arisaka-GG-A6-GRIP.
Preprint, available at https://preprints.arisaka-gg.org/a6-grip/.
doi:10.5281/zenodo.22064306.
Relevance. The immediate predecessor for the runtime constitution used in the cosmos paper. Supplies the Gi-4 four-pillar backbone (TGG, CCI, GFF, GRIP), the free-energy split $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$, Wilson-loop GSSB, and SIP; it also defines the two runtime regimes GRIP-In/GRIP-Out adopted in the cosmic GRIP grammar of the present paper. GG-A10-Cosmos runs cosmic history on this engine: contraction toward attractors under GRIP-In and necessary escape under GRIP-Out organize the radial chronology and the relic/vacuum endpoints. First cited in §6.

K. Arisaka. *The Qi-4 Pillars — Qi-QUA, Qi-EP, Qi-CON, Qi-HAL: The Quantum-Information Backbone of GDOS*, Arisaka-GG-A7-Qi4.
Preprint, available at https://preprints.arisaka-gg.org/a7-qi4/.
doi:10.5281/zenodo.22064308.
Relevance. Supplies the Qi-4 quantum-information completion of the GDOS runtime — the Qi-QUA, Qi-EP, Qi-CON, Qi-HAL pillars, the SeOS/QCI language, the information balance IH + OH = EIC, and the five-story pagoda. GG-A10-Cosmos uses this apparatus in its later interpretive Parts: the cosmological boundary is read as an information channel, and the Qi-QUA ledger anchors the quantitative bookkeeping that the cosmic inventory shares with the rest of the cohort. First cited in §6.

K. Arisaka. *Quantum Mechanics Derived from GG-Theory as Projection from 6D Bulk to 4D Boundary*, Arisaka-GG-A8-QM.
Preprint, available at https://preprints.arisaka-gg.org/a8-qm/.
doi:10.5281/zenodo.22064310.
Relevance. The quantum-mechanics member of the 4D-Physics Trio: reads textbook 4D quantum mechanics as the boundary projection of the 6D bulk, deriving the Born rule, Bell/CHSH violation, and decoherence as theorems of the non-invertible 6D-to-4D OGGEP projection. The present paper applies the same OGGEP boundary projection one story up, at the cosmological boundary, so GG-A8-QM is the structural sibling that fixes the projection grammar GG-A10-Cosmos reuses — bulk to boundary on the QM stratum there, bulk to boundary on the cosmological stratum here. First cited in §6.

K. Arisaka. *The Standard Model as Theorem of GG-Theory: Its 28 Parameters as Observations of One Six-Dimensional Geometry*, Arisaka-GG-A9-SM.
Preprint, available at https://preprints.arisaka-gg.org/a9-sm/.
doi:10.5281/zenodo.22064312.
Relevance. The immediate particle-to-cosmos bridge: fixes the Standard-Model ledger, the compact-angle ontology, the neutrino block, and the parameter constraints as theorems of GG-6, BI-6, GDOS, GRIP, and GFF. GG-A10-Cosmos inherits this ledger as the matter content of the cosmological inventory — the baryonic and neutrino budgets that populate $\Omega_m = 1/\pi$ — and

reuses GG-A9-SM's compact-angle reading to identify the Kaxion as a non-spatial compact-angle degree of freedom of the same bulk. First cited in §9.1.

K. Arisaka. *The Geometric Inevitability of Cosmos, Life, and Mind: The Self-Referential Closure of the GG-Theory Program*, Arisaka-GG-A11-Origin.
Preprint, available at https://preprints.arisaka-gg.org/a11-origin/.
doi:10.5281/zenodo.22064319. No result of the present paper depends on it.
Relevance. Defines the Ground Geometric Attractor (GGA) and proves T-GGA-as-Ground-State, the ground-state refinement of the GRIP terminus that the present paper uses to distinguish transient cosmic histories (GA's) from persistent relics and final states (GGA's) — specifically the Kaxion dark matter and the de Sitter vacuum. GG-A11-Origin also proves the regime theorems T-GRIP-In-Lyapunov-Descent and T-GRIP-Out-Necessity that underlie the GRIP-In/GRIP-Out reading of cosmic history adopted here, and supplies the four-move operational anatomy (DGI → GFF Morse flow → GSSB → GA) that the cosmic GRIP grammar specializes. First cited in §9.1.

K. Arisaka. *The GG-6 Lock Cluster: A Non-Spatial Kaluza–Klein Resonance near 3.85 TeV as the Standing LHC Falsifier of GG-Theory*, Arisaka-GG-P1-KK preprint.
Preprint, available at https://preprints.arisaka-gg.org/p1-kk/.
doi:10.5281/zenodo.22064321.
Relevance. The first paper of the Predictions Branch: locates the Kaluza-Klein lock cluster at $M_{\text{lock}} = (1/2)\, M_{\text{Pl}}\, \mathrm{e}^{-35} \approx 3.85$ TeV, with explicit collider falsifiers. GG-A10-Cosmos inherits this scale as the upper end of the inherited scale chain and the high-energy edge of the radial cosmic chronology — the earliest, deepest-bulk rung of the radial timeline I_r along which the boundary cosmology unfolds. First cited in §9.1.

K. Arisaka. *Why Three Generations: A Three-Leg Derivation of $N_g = 3$ from BI-6, MSMF, and GDOS Morse Stability*, Arisaka-GG-P2-N3 preprint.
Preprint, available at https://preprints.arisaka-gg.org/p2-n3/.
doi:10.5281/zenodo.22064323.
Relevance. The dedicated family-number companion: executes the three-leg derivation of the generation count $N_g = 3$ that the present paper reserves — the MSMF lower bound from CP-violation combinatorics, the GDOS capacity upper bound on $\mathbb{C}P^{N_g-1}$ (the capacity theorem Cap $= \pi N_g$ with exact saturation, backed by the integrality theorem L-INT and by two obstruction theorems; conditional on one declared topological residue, R-INT), and the codebook consistency check — together with the GRIP-In reachability reading: a fourth family's ground state would be dynamically unreachable within the stratum's instanton budget. That dynamical content is the cosmological three-generation count GG-A10-Cosmos uses when it populates the neutrino and lepton budgets of the cosmic inventory. First cited in §9.1.

K. Arisaka. *The Electron as the Ground State of Nature: The Six Oldest Questions about the Electron, Closed and Fenced within GG-Theory*, Arisaka-GG-P3-Electron preprint.
Preprint, available at https://preprints.arisaka-gg.org/p3-electron/.
doi:10.5281/zenodo.22064325.

Relevance. The third paper of the Predictions Branch: reads the electron as the ground-state Geometric Attractor of nature — identity, absolute stability, the two excited copies, the positron as conjugate readout, pointlikeness, and field response derived as six faces of ground-state-ness, importing GG-P2-N3's $N = 3$ verbatim. GG-A10-Cosmos cites it as the particle-stratum instance of the universal-identity principle that the present paper extends to the cosmological inventory: the same attractor logic that makes every electron identical makes the cosmic relics and the vacuum the unique GGA endpoints of their strata. First cited in §5.

K. Arisaka. *The Real Geometry Beneath the Complex Numbers of Physics: Positive Frequency is Holomorphy — from Cardano's Cubics to Twistor Space, via the Sixth Dimension of GG-Theory*, Arisaka-GG-P4-Real preprint.
Preprint, available at https://preprints.arisaka-gg.org/p4-real/.
doi:10.5281/zenodo.22064327.
Relevance. The origin-of-i member paper of the Foundations Branch: the standalone, full-length development of the real-bulk series this paper leans on at the observer-indexed vacuum and at the dial-free door (T-Cosmos-18) — the two real sources of the derived complex structure, the holomorphy bridge at its Hardy, tube, and twistor levels, and the open items stated as open. First cited in §9.1.

K. Arisaka. *The Mirrors of Matter: Why Every Particle Has a Twin, Why Nature's Mirrors Break, and Why the Universe Is Made of Matter*, Arisaka-GG-P5-CPT preprint.
Preprint, available at https://preprints.arisaka-gg.org/p5-cpt/.
doi:10.5281/zenodo.22064329.
Relevance. The CPT member paper of the Foundations Branch and the GG-A9-SM–GG-A10-Cosmos interface of the cohort map: C, P, and T as orientation reversals of the three real bulk factors, the chiral-pentagon CP phases, and the matter–antimatter asymmetry through the weight-zero door. The ledger phases this paper's baryon-asymmetry chain consumes — the PMNS Dirac phase and the Majorana parity $\alpha_{31} = \pi$, derived in GG-A9-SM — reach Part V and Appendix I through it, at the grade recorded there. First cited in §9.1.

K. Arisaka. *The Proton as the Ground Geometric Attractor of Matter: The Eleven Questions of the Proton, Closed and Fenced within GG-Theory*, Arisaka-GG-P6-Proton preprint.
Preprint, available at https://preprints.arisaka-gg.org/p6-proton/.
doi:10.5281/zenodo.22064331.
Relevance. The Predictions Branch member that runs the ground-attractor method of GG-P3-Electron on the baryonic stratum. Its bearing on GG-A10-Cosmos is a contrast rather than an inheritance, and the contrast is worth stating because this paper's matter sector depends on the distinction. The baryon asymmetry of Part V is a *number* that survived, computed from a chain with a named input and a postulated exponent; the proton's stability in GG-P6-Proton is a *winding class* that no local process can change. Protection by an amount and protection by a topological class are different mechanisms on one geometry, and only the second is exact. Cited at roster level; nothing in GG-A10-Cosmos depends on its results. First cited in §9.1.

K. Arisaka. *The Kaxion: Cold Dark Matter as the Non-Spatial Compact-Angle Pseudoscalar of GG-6*, Arisaka-GG-C1-Kaxion preprint.
Preprint, available at https://preprints.arisaka-gg.org/c1-kaxion/.
doi:10.5281/zenodo.22064333.
Relevance. The dedicated dark-matter companion, executing the closure that the present paper's Conjecture III would make derived: the π-locked abundance closure (bulk normalization $f_{\rm bulk} = 3.04\times10^{13}$ GeV, calibrated and declared; single physical decay constant $f_K = f_{\rm bulk}/43 \simeq 7.1\times10^{11}$ GeV), the integrated misalignment normalization adopted in Eq. (155), the laboratory window $\nu_K \in [1.59, 2.13]$ GHz of seven rungs ($m_K = 8.1\,\mu$eV, $\pm1.3\%$ at fixed f_K), the exact anomaly ratio $E/N = 8/3$ placing the Kaxion on the DFSZ line, the single-patch history carried as Conjecture C-Kaxion-d, and the count — one line, alone — as the identification test that replaced the withdrawn twin-peak and 5:2 readouts. GG-A10-Cosmos treats it as the phenomenology that converts its $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ dark-matter share into a falsifiable haloscope target. First cited in §4.1.

K. Arisaka. *The Hubble Tension in the Light of GG-Theory: The Hubble Constant as One Reading of a Derived Expansion History $H(t)$*, Arisaka-GG-C2-Hubble preprint.
Preprint, available at https://preprints.arisaka-gg.org/c2-hubble/.
doi:10.5281/zenodo.22064336.
Relevance. The dedicated Hubble companion: a program-level reading of the H0DN Local Distance Network consensus $H_0 = 73.50 \pm 0.81$ km/s/Mpc (Casertano et al., 2026) through the derived background rate of the present paper — the two-route geometric determination of H_0, the void channel priced and not carried, and the whole excess as the named calibration systematic, priced against the published corrections. It carries the falsifier registry F-Hubble-1 to F-Hubble-11 (including the pre-registered $\zeta_\Lambda = \pi/175$ target and the standard-siren kill test), with the falsifiable exposure to future releases of the distance network itself. GG-A10-Cosmos supplies the one background and its theorem, of which GG-C2-Hubble is the dedicated observational study. First cited in §9.1.

K. Arisaka. *Black-Hole Entropy under GG-Theory: The GDOS Reading of the Horizon as an OGGEP Port — with the Exact Bekenstein–Hawking $A/4G_N$ Derived from Induced Gravity*, Arisaka-GG-C3-BH preprint.
Preprint, available at https://preprints.arisaka-gg.org/c3-bh/.
doi:10.5281/zenodo.22064338.
Relevance. The cohort's strong-gravity anchor: derives the Bekenstein–Hawking entropy $A/4$ with the horizon treated as an OGGEP port, the coefficient following from induced gravity up to the species closure it registers as F-BH-9. GG-A10-Cosmos cites it in its black-hole-information and primordial-relic Part: the same OGGEP-port treatment that yields $A/4$ governs how horizons exchange information with the cosmological boundary, tying the strong-gravity ledger into the boundary cosmology of the present paper. First cited in §9.1.

K. Arisaka. *The Holonomic Origin of Leptogenesis: CP Violation from the Chiral Pentagon to the Baryon-to-Photon Ratio*, Arisaka-GG-C4-Leptogenesis preprint.

Preprint, available at https://preprints.arisaka-gg.org/c4-leptogenesis/.
doi:10.5281/zenodo.22190046.
Relevance. The cohort's dedicated leptogenesis paper: the flavored Boltzmann integration of the resonant chain on the inherited phases, and the resolution of the CP phase $\alpha_{31} = \pi$ as conserving in the flavor sum and violating in the flavor differences. GG-A10-Cosmos's Part V and Appendix I compute the same chain; the state count and the resulting η_B are carried here on the two-state count of GG-A9-SM, which the companion's deposited version predates.

K. Arisaka. *The Acoustic Peaks as a Theorem: The Microwave Background from a Six-Dimensional Geometry with No Adjustable Parameter — What Is Derived, What Is Inherited, and What Would End It*, Arisaka-GG-C5-CMB preprint (2026), cited without a concept DOI.
Relevance. The dedicated microwave-background companion, written from this paper's Part VI and the Hubble companion against a source ledger. It carries this paper's CMB numbers at this paper's grades and is the one companion that has confronted the 2025 tilt measurements; §28.10 follows its comparison.

K. Arisaka. *The Galaxies That Came Too Early: Cosmic Dawn Through JWST and the Rare-Halo Tail in the Light of GG-Theory*, Arisaka-GG-C6-JWST preprint (2026), cited without a concept DOI.
Relevance. The dedicated early-galaxy companion. It executed this Part's own kill criterion F-JWST-2 on the early-residue boost, which is why §44.3 withdraws that derivation; this paper leans on that execution and cites it here for it. Its dark-matter and fifth-force material predates the boundary-condition termination of $U(1)_B$ and the pinning conjecture carried here, and does not yet reflect them.

K. Arisaka. *Symmetry, Conservation, and Geometry: Seven Steps from Noether, Bianchi, to Hopkins–Singer, and the Ledger the Open-System Turn Pays Out*, Arisaka-GG-W1-Symmetry preprint.
Preprint, available at https://preprints.arisaka-gg.org/w1-symmetry/.
doi:10.5281/zenodo.22064340.
Relevance. The White-Paper Branch survey of how conservation moved from symmetry to cohomology. It supplies the context for a habit this paper relies on without arguing for it: the cosmic inventory is treated as a set of classes forced by the geometry rather than as a set of quantities conserved by a symmetry, and the difference is why an entry in the codebook can be Derived without any symmetry being named beside it. The survey is conceptual and carries no computation GG-A10-Cosmos uses. Cited at roster level. First cited in §9.1.

K. Arisaka. *The Circle That Remembers, the Interval That Forgets: GG-Theory and the Sixteen Problems of Particle Physics and Cosmology*, Arisaka-GG-W2-Holonomy preprint.
Preprint, available at https://preprints.arisaka-gg.org/w2-holonomy/.
doi:10.5281/zenodo.22064342.
Relevance. The White-Paper Branch treatment of protection by topology, and the closest of the three to GG-A10-Cosmos's own argument. The compact phase circle whose holonomy this

paper reads as the seed of the Kaxion, and whose winding carries the baryon number the matter Part accounts for, is one rung of the ladder of protection mechanisms classified there. A reader who wants to know why that rung and not another should read the classification rather than this paper's use of it. Its sixteen open problems overlap the named conjectures of the codebook. Cited at roster level. First cited in §9.1.

B. Other references

L. F. Abbott and P. Sikivie. A cosmological bound on the invisible axion. *Phys. Lett. B* **120** (1983) 133–136.

Relevance. The second of the three 1983 misalignment papers, deriving the cosmological bound relating the axion mass and decay constant to its relic abundance. GG-A10-Cosmos's Kaxion appendix uses the same misalignment relic-abundance scaling (set by θ_i, m, and the oscillation-onset epoch) to connect the Kaxion parameters to its dark-matter share, and cites this paper for the abundance–mass relation. First cited in §M.

M. Abdul Karim et al. (DESI Collaboration). DESI DR2 results II: measurements of baryon acoustic oscillations and cosmological constraints. *Phys. Rev. D* **112** (2025) 083515; arXiv:2503.14738.

Relevance. The BBN-calibrated baryon-acoustic determination $H_0 = 68.51 \pm 0.58\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$, the sharpest ladder-free value in hand, quoted in §42.4 because it is 2.8σ above this paper's derived background and a prediction about the ladder-free probes must be stated against the ladder-free probes that have reported. Its inference takes r_d from BBN-calibrated ω_b and is therefore not free of the acoustic chain.

E. Abdalla *et al.* Cosmology intertwined: a review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies. *JHEAp* **34** (2022) 49–211.

Relevance. The 2022 "Cosmology Intertwined" review is the broad community survey linking the H_0 and S_8 tensions to each other and to a range of cosmological anomalies. GG-A10-Cosmos cites it in the S_8 section as the reference establishing that the clustering tension is a genuine, multi-probe discrepancy, which the framework's computed $S_8 = 0.825$ shares rather than resolves. First cited in §4.

Ade, P., et al. (The Simons Observatory Collaboration). (2019). The Simons Observatory: science goals and forecasts. *Journal of Cosmology and Astroparticle Physics*, *2019*(02), 056; arXiv:1808.07445.

Relevance. The Simons Observatory forecast, $\sigma(r) \simeq 0.003$ for the baseline survey and 0.002 for the goal configuration, is the near-term design floor against which the derived $r = 0.0038$ is placed in Parts IV and VII: a signal of that size sits at one to two standard deviations of this instrument and is decided only by the next one. First cited in §28.

Aharonov, Y., & Bohm, D. (1959). Significance of electromagnetic potentials in the quantum theory. *Physical Review, 115*(3), 485–491.

Relevance. The 1959 Aharonov–Bohm paper, the classical demonstration that a gauge-invariant phase accumulated around a loop is physically real even where the local field vanishes. GG-A10-Cosmos cites this as historical context for the Wilson-loop and holonomy reading of cosmic branch selection: the cosmic-grip grammar treats the compact-angle holonomy of the inherited GG-6-6 bulk, not a local order parameter, as the genuine gauge-invariant coordinate that labels which cosmological branch is realized. First cited in §14.1.

A. Almheiri, N. Engelhardt, D. Marolf, and H. Maxfield. The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole. *J. High Energy Phys.* **12** (2019) 063.

Relevance. With Penington (2020), the 2019 derivation of the Page curve from the generalized entropy of an evaporating black hole, introducing the island contribution. Part VII's island formula is this result; the companion Arisaka (2026, C3-BH) carries the full attribution.

K. Akita and M. Yamaguchi. A precision calculation of relic neutrino decoupling. *JCAP* **08** (2020) 012; arXiv:2005.07047.

Relevance. The Standard-Model value $N_{\rm eff} = 3.044$ from which the codebook's derived-null row D12 starts, with non-instantaneous decoupling and finite-temperature QED corrections included; Froustey, Pitrou and Volpe (*JCAP* **12** (2020) 015) and Bennett et al. (*JCAP* **04** (2021) 073) reach the same value independently. First cited in §68.1.

Albrecht, A., & Steinhardt, P. J. (1982). Cosmology for grand unified theories with radiatively induced symmetry breaking. *Physical Review Letters, 48*(17), 1220–1223.

Relevance. Together with Linde (1982), the 1982 "new inflation" proposal that replaced the false-vacuum picture of Guth (1981) by a slow-rolling scalar field and thereby repaired its graceful-exit problem. GG-A10-Cosmos cites this as historical background for the inflation section, which retains the slow-roll structure of new inflation but constrains the inflationary modulus to a specific direction in the inherited bulk-modulus space and locates the inflationary scale at TeV rather than GUT energies. First cited in §2.

Alpher, R. A., Bethe, H., & Gamow, G. (1948). The origin of chemical elements. *Physical Review, 73*(7), 803–804.

Relevance. The 1948 founding paper of hot-Big-Bang nucleosynthesis, proposing that the light elements are built up by neutron-proton interactions in a cooling, expanding plasma. GG-A10-Cosmos cites this as the historical opening of the BBN section, which retains the standard light-element calculation (D, ^{3}He, ^{4}He, ^{7}Li) and asks only whether the inherited ledger modifies any BBN input in a way that addresses the residual lithium-7 anomaly without disturbing the other three abundances. First cited in §2.

Amari, S.-I., & Nagaoka, H. (2000). *Methods of Information Geometry.* Providence, RI: American Mathematical Society.

Relevance. The standard monograph on information geometry, establishing the Fisher-Rao metric and dual-connection structure on statistical manifolds. GG-A10-Cosmos cites this for orientation in the later Qi-runtime and EIC parts, where the state-manifold and selector language of the boundary runtime is described; it is a comparison vocabulary, not a derivational input to the cosmology. First cited in §106.

Asgari, M., Lin, C.-A., Joachimi, B., et al. (2021). KiDS-1000 cosmology: Cosmic shear constraints and comparison between two point statistics. *Astronomy & Astrophysics*, *645*, A104.
Relevance. The benchmark low-redshift cosmic-shear determination of the clustering parameter, $S_8 = 0.759^{+0.024}_{-0.021}$. Cited as the cosmic-shear comparator for S_8 in the parameter-audit Dev plot and data table. The framework's computed $S_8 = 0.825 \pm 0.010$ sits $+2.5\sigma$ above this measurement — the S_8 tension, which the framework shares with ΛCDM rather than resolving; an earlier version quoted a corridor-response prediction 0.78 ± 0.01 that agreed with this measurement at the 0.9σ level, the same two-branch structure proposed as a candidate resolution of the Hubble tension. First cited in §4.1.

J. M. Bardeen. Gauge-invariant cosmological perturbations. *Phys. Rev. D* **22** (1980) 1882–1905.
Relevance. Bardeen's 1980 paper introduced the gauge-invariant formulation of cosmological perturbation theory, resolving the gauge ambiguities that plagued the Lifshitz approach. GG-A10-Cosmos's growth appendix works in the conformal-Newtonian gauge with the gauge-invariant potentials Φ, Ψ that Bardeen's construction justifies, ensuring that the growth equation and the resulting (σ_8,S_8) normalization are physical rather than gauge artifacts. First cited in §3.

J. M. Bardeen, P. J. Steinhardt, and M. S. Turner. Spontaneous creation of almost scale-free density perturbations in an inflationary universe. *Phys. Rev. D* **28** (1983) 679–693.
Relevance. The 1983 Bardeen–Steinhardt–Turner paper provided the gauge-invariant, conserved-curvature-perturbation treatment that put the inflationary spectrum calculation on rigorous footing. GG-A10-Cosmos's inflation appendix relies on the conservation of the curvature perturbation on super-horizon scales it established to connect the spectrum at horizon exit to its re-entry in the growth appendix. First cited in §3.

J. M. Bardeen, J. R. Bond, N. Kaiser, and A. S. Szalay. The statistics of peaks of Gaussian random fields. *Astrophys. J.* **304** (1986) 15–61.
Relevance. The 1986 "BBKS" paper provides the standard fitting form for the cold-dark-matter transfer function and the peak statistics of Gaussian random fields. GG-A10-Cosmos's growth appendix uses the BBKS transfer function $T(k)$ to pass from the primordial spectrum to the processed matter power spectrum $P(k)$ whose $8\,h^{-1}$Mpc variance defines σ_8. First cited in §3.

D. Baumann. *Cosmology.* Cambridge University Press, Cambridge, (2022).
Relevance. The modern graduate cosmology text: background, thermal history, inflation, perturbations, and the CMB, in the notation closest to the present course. Recommended as a primary reference for the constructive appendices (E–K) of the graduate course in Part XV. First cited in §A.

Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D,* 7(8), 2333–2346.
Relevance. The 1973 paper that assigned an entropy proportional to horizon area to a black hole and so opened black-hole thermodynamics. GG-A10-Cosmos cites this to frame the horizon-ledger problem in the Part V audit, where the area-entropy bookkeeping is reinterpreted through the GFF free-energy functional and the GDOS open-system dynamics rather than taken as a primitive. First cited in §4.1.

G. Bertone and D. Hooper. History of dark matter. *Rev. Mod. Phys.* **90** (2018) 045002.
Relevance. The 2018 Bertone–Hooper review surveys the observational and theoretical history of dark matter and the landscape of candidate models. GG-A10-Cosmos cites it in the Kaxion section as the contextual reference situating the Kaxion misalignment candidate within the broader dark-matter program and its abundance constraints. First cited in §3.

BICEP/Keck Collaboration (Ade, P. A. R., et al.). (2021). Improved constraints on primordial gravitational waves using *Planck*, WMAP, and BICEP/Keck observations through the 2018 observing season. *Physical Review Letters*, *127*(15), 151301.
Relevance. The BK18 bound on the tensor-to-scalar ratio, $r_{0.05} < 0.036$ at 95% confidence, is the present bound against which the derived $r = 0.0038$ is scored — a factor of nine below it — in the parameter-audit table and in Part IV; the 2025 joint analysis (arXiv:2512.10613, 2025) tightens it to 0.034. First cited in Table 2.

S. Birrer *et al.* (TDCOSMO Collaboration) (2025). TDCOSMO XVII. New time-delay cosmography results using lens model constraints from JWST and ALMA. *Astronomy & Astrophysics* **704**, A63; arXiv:2506.03023.
Relevance. The current strong-lensing time-delay determination of H_0, $71.6^{+3.9}_{-3.3}$, superseding H0LiCOW 2020: ladder-free and sound-horizon-free, it reads the background, and at 1.5σ from the derived rate it decides nothing yet; a converged value at 73 would refute the one-background reading. First cited in §42.

Bianchi, L. (1902). Sui simboli a quattro indici e sulla curvatura di Riemann. *Rendiconti della R. Accademia dei Lincei, 11*, 3–7.
Relevance. Founding paper of the Bianchi identity $dR = 0$, the geometric ancestor of the modified Bianchi identity BI-6 of (Arisaka, 2026, A3-BI6); cited in the framing remark and in the universal-cosmic-identity section as historical lineage for the constitutional descent $dF = 0 \to dR = 0 \to dX_4 = 0$. First cited in §E.

G. Bressi, G. Carugno, F. Della Valle, G. Galeazzi, G. Ruoso, and G. Sartori. Testing the neutrality of matter by acoustic means in a spherical resonator. *Phys. Rev. A* **83** (2011) 052101.
Relevance. The sharpest laboratory audit of exact neutrality: the charge per hydrogen molecule bounded at the $10^{-21}\,e$ level. In GG-A10-Cosmos this is the precision anchor of the neutrality theorem (§36.13) — the exactness that the trace identity of the unified embedding holds, not protects.

E. Camphuis, W. Quan, L. Balkenhol *et al.* (2025). Joint analysis of SPT-3G, *Planck*, and ACT CMB temperature and polarization data. arXiv:2506.20707.

Relevance. The 2025 three-instrument microwave-background combination, $H_0 = 67.24 \pm 0.35$, the sharpest early-universe reading and the comparator the Hubble companion scores against; 1.1σ above the derived 66.8. First cited in §42.

S. M. Carroll. The cosmological constant. *Living Rev. Relativity* **4** (2001) 1.
Relevance. Carroll's 2001 Living Review is a standard pedagogical reference on the cosmological constant and the vacuum-energy problem. GG-A10-Cosmos's dark-energy section cites it for the statement of the fine-tuning and coincidence problems that the framework's late vacuum term — a descendant of the same anomaly-locked geometric stage — is meant to address rather than to insert by hand. First cited in §3.

Casertano, S., Anand, G. S., Anderson, R. I., et al. (H0DN Collaboration) (2026). The Local Distance Network: A community consensus report on the measurement of the Hubble constant at ∼1% precision. *Astronomy & Astrophysics, 708*, A166.
Relevance. The covariance-weighted Local Distance Network consensus $H_0 = 73.50 \pm 0.81$ km/s/Mpc — the codebook's ladder comparator, promoted from the single-team Riess et al. (2022) value; the companion GG-C2-Hubble is the program-level study of this measurement. First cited in §4.1.

A. G. Cohen, A. De Rújula, and S. L. Glashow. A matter–antimatter universe? *Astrophys. J.* **495** (1998) 539.
Relevance. The observational closure of the domain alternative: contact zones between matter and antimatter regions would shine in annihilation gamma rays, and the diffuse backgrounds exclude such zones out to scales comparable to the visible universe. GG-A10-Cosmos cites it in the third-horizon Section (§73) as the empirical fact that every disconnected region made the *same* choice — the observation that the determinism of the cascade must, and does, explain.

S. Cole *et al.* (2dFGRS Team). The 2dF Galaxy Redshift Survey: power-spectrum analysis of the final data set and cosmological implications. *Mon. Not. R. Astron. Soc.* **362** (2005) 505–534.
Relevance. The 2005 2dFGRS power-spectrum analysis is one of the two first detections of the baryon acoustic feature in the galaxy distribution. GG-A10-Cosmos cites it in the expansion section as observational confirmation of the standard ruler in which the one-background reading predicts *no* residue at any redshift; a percent-level residue counts against it. First cited in §3.

E. J. Copeland, M. Sami, and S. Tsujikawa. Dynamics of dark energy. *Int. J. Mod. Phys. D* **15** (2006) 1753–1936.
Relevance. The 2006 Copeland–Sami–Tsujikawa review is the comprehensive reference for the dynamics of dark-energy models, including scalar-field quintessence and the equation-of-state phenomenology. GG-A10-Cosmos's dark-energy section cites it for the parametrization of the dark-energy equation of state against which the framework's residual-vacuum branch is compared. First cited in §3.

R. H. Cyburt, B. D. Fields, K. A. Olive, and T.-H. Yeh. Big bang nucleosynthesis: present status. *Rev. Mod. Phys.* **88** (2016) 015004.

Relevance. The 2016 Cyburt *et al.* review is the authoritative modern status report on BBN, with up-to-date reaction rates, abundance predictions, and observational comparisons. GG-A10-Cosmos's BBN appendix cites it for the current best-fit light-element yields and uncertainties against which the worked Y_p computation is checked. First cited in §3.

P. S. Bhupal Dev, P. Millington, A. Pilaftsis, and D. Teresi. Flavour covariant transport equations: an application to resonant leptogenesis. *Nucl. Phys. B* **886** (2014) 569–664 [arXiv:1404.1003].

Relevance. The flavor-covariant density-matrix transport formalism for resonant leptogenesis, in which the flavor coherences of the heavy and charged-lepton sectors are dynamical variables. It is the natural framework in which condition C4 of Appendix I must be discharged; in the regimes it treats, the effective regulator is of the heavy-sector class, which is why C4 is carried open rather than assumed. First cited in Appendix I.

E. Di Valentino, O. Mena, S. Pan, L. Visinelli, W. Yang, A. Melchiorri, D. F. Mota, A. G. Riess, and J. Silk. In the realm of the Hubble tension — a review of solutions. *Class. Quant. Grav.* **38** (2021) 153001.

Relevance. The 2021 Di Valentino *et al.* review is the comprehensive catalog of proposed solutions to the Hubble tension, classified into early-time, late-time, and systematics-based approaches. GG-A10-Cosmos's Hubble-tension section cites it to position the one-background reading within this taxonomy and to motivate the BI-6 refusal to scan over arbitrary new fluids. First cited in §4.

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. (1965). Cosmic black-body radiation. *Astrophysical Journal, 142*, 414–419.

Relevance. The 1965 companion paper to Penzias and Wilson (1965), interpreting the 4080 Mc/s excess as the cooled black-body afterglow of a hot early universe. GG-A10-Cosmos cites this as historical background for the CMB section, which retains the standard acoustic, damping, and projection structure of the microwave background and asks only where the inherited ledger reshapes the source assumptions behind it. First cited in §2.

Coleman, S. & De Luccia, F. (1980). Gravitational effects on and of vacuum decay. *Phys. Rev. D* **21**, 3305–3315.

Relevance. The bubble-nucleation event by which a metastable vacuum decays in a gravitational background. GG-A10-Cosmos cites it in the Common Objections (Objection O-Cosmos-90) as the mechanism by which a landscape framework must have our entire observable patch descend from a single nucleation — the internally consistent repair whose cost the second-horizon discussion prices. First cited in §O9.

M. Dine and W. Fischler. The not-so-harmless axion. *Phys. Lett. B* **120** (1983) 137–141.

Relevance. The third of the three founding 1983 misalignment papers, emphasizing the cosmological constraints on the axion decay constant from the relic-abundance overclosure bound. GG-A10-Cosmos's Kaxion appendix cites it alongside the other two as the joint origin of the misalignment dark-matter calculation it carries out. First cited in §M.

M. Dine, W. Fischler, and M. Srednicki. A simple solution to the strong CP problem with a harmless axion. *Phys. Lett. B* **104** (1981) 199–202.
Relevance. The 1981 "DFSZ" paper is one of the two canonical invisible-axion model constructions, embedding the Peccei–Quinn symmetry with a large decay constant. GG-A10-Cosmos's Kaxion appendix cites it as a reference model for how a phenomenologically viable axion sector is realized, against which the Kaxion's geometric (Wilson-line) origin is contrasted. First cited in §M.

S. Dodelson. *Modern Cosmology.* Academic Press, San Diego, (2003).
Relevance. Dodelson's 2003 textbook is the standard pedagogical source for the Boltzmann-hierarchy treatment of perturbations, the growth of structure, and the CMB acoustic peaks. GG-A10-Cosmos follows its line-by-line style in the perturbation/growth appendix (the master equation $\ddot{\delta} + 2H\dot{\delta} - 4\pi G\rho_m\delta = 0$ and the transfer function) and in the CMB acoustic-peak appendix; the framework leaves this standard linear theory unmodified and acts only on the late-time growth history. First cited in §3.

R. Durrer. *The Cosmic Microwave Background*, 2nd ed. Cambridge University Press, Cambridge, (2020).
Relevance. The dedicated CMB monograph: full Boltzmann hierarchy, polarization, and lensing. Recommended companion for the acoustic-peak appendix (K), whose derived-parameter confrontation (Part VI) hands its inputs to exactly the machinery this book develops. First cited in §A.

Douglas, M. R. (2003). The statistics of string/M theory vacua. *J. High Energy Phys.* **05**, 046. [arXiv:hep-th/0303194].
Relevance. The canonical statistical statement of the string vacuum landscape — of order 10^{500} consistent flux vacua with no dynamical selection principle. GG-A10-Cosmos cites it in Objection O-Cosmos-90 as the precise form of the vacuum-selection freedom the unique-vacuum wager is set against. First cited in §O9.

Einstein, A. (1917). Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, 142–152.
Relevance. The 1917 paper that introduced the cosmological field equations and the constant Λ in order to permit a static universe. GG-A10-Cosmos cites this as historical background both for the beginning-of-universe section, which retains the field-equation lineage from Einstein through Friedmann and Lemaître, and for the dark-energy section, which retains the empirical Λ but reinterprets its scale through the residual GFF ledger of the inherited bulk. First cited in §1.

D. J. Eisenstein *et al.* (SDSS Collaboration). Detection of the baryon acoustic peak in the large-scale correlation function of SDSS luminous red galaxies. *Astrophys. J.* **633** (2005) 560–574.
Relevance. The 2005 SDSS detection of the BAO peak in the LRG correlation function is the landmark measurement of the low-redshift standard ruler. GG-A10-Cosmos cites it in the expansion section as the empirical anchor for the BAO distance scale; the paper's falsifiable prediction of *no*

BAO residue at any redshift on the one background is stated directly against this measurement class. First cited in §3.

G. F. R. Ellis. Issues in the philosophy of cosmology. In *Philosophy of Physics (Handbook of the Philosophy of Science)*, Elsevier, Amsterdam (2006) 1183–1285.
Relevance. Ellis's survey is the standard statement of the epistemic foundations of cosmology, including the observation that "cosmology proceeds by assuming the laws of physics are the same everywhere" — an assumption only partially testable from inside one horizon. GG-A10-Cosmos cites it in the second-horizon-problem pedigree as the clearest published acknowledgment that the uniformity of the laws is an assumption of the standard framework: noticed by philosophers of cosmology, never mechanized. First cited in §72.2.

R. P. Feynman. The development of the space-time view of quantum electrodynamics. Nobel Lecture, December 11 (1965).
Relevance. The lecture preserves Wheeler's 1940 telephone call proposing the one-electron universe — all electrons identical because they are one electron threading back and forth through time. GG-A10-Cosmos cites it in the second-horizon-problem pedigree as the famous anticipation that proves the rule: the identity of the electrons was once felt to demand a mechanism, and the episode survives as an anecdote, not as a named problem. First cited in §72.2.

B. D. Fields. The primordial lithium problem. *Ann. Rev. Nucl. Part. Sci.* **61** (2011) 47–68.
Relevance. Fields' 2011 review is the canonical statement of the primordial lithium problem — the persistent factor-of-three discrepancy between the BBN-predicted and observed ^{7}Li abundance. GG-A10-Cosmos's BBN appendix flags the lithium problem honestly as an unresolved standard-model tension that the framework does not claim to solve, distinguishing it from the well-fit helium and deuterium yields. First cited in §3.

W. L. Freedman. Measurements of the Hubble constant: tensions in perspective. *Astrophys. J.* **919** (2021) 16.
Relevance. Freedman's 2021 paper presents the tip-of-the-red-giant-branch (TRGB) calibration of H_0, which yields a value intermediate between the Cepheid and CMB determinations and bears directly on whether the tension is physical or systematic. GG-A10-Cosmos's Hubble-tension section cites it as the key alternative local calibration against which the predicted ladder systematic can be tested. First cited in §42.

Freivogel, B. (2011). Making predictions in the multiverse. *Class. Quantum Grav.* **28**, 204007. [arXiv:1105.0244].
Relevance. A review of the measure problem of eternal inflation: the relative frequencies of observations are undefined without an additional, non-derived regulator. GG-A10-Cosmos cites it in Objection O-Cosmos-90 as the reason the landscape's one-bubble repair defers rather than restores predictivity. First cited in §O9.

Friedmann, A. (1922). Über die Krümmung des Raumes. *Zeitschrift für Physik, 10*(1), 377–386.
Relevance. The 1922 paper that discovered the dynamical, time-dependent homogeneous-isotropic solutions of Einstein's equations and established what became the FRW family — the kinematic

backbone of every cosmology paper since. GG-A10-Cosmos cites this as historical background for the beginning and expansion-and-cooling sections, both of which retain the Friedmann dynamics intact, and it underwrites the explicit FRW-derivation appendix §E, which reconstructs the metric, Christoffel symbols, and Friedmann equation that feed the inherited source ledger. First cited in §1.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127–138.
Relevance. The 2010 statement of the free-energy principle, casting perception and action as variational minimization of a surprise functional. GG-A10-Cosmos cites this as an external comparison point for the inference and internal-reconstruction language imported from GG-A7-Qi4 (the Qi-4 pillars) into the later EIC discussion; it is an orienting analog, not a derivational input. First cited in §106.

M. Fukugita and T. Yanagida. Baryogenesis without grand unification. *Phys. Lett. B* **174** (1986) 45–47.
Relevance. The 1986 Fukugita–Yanagida paper introduced leptogenesis: a lepton asymmetry from heavy right-handed-neutrino decays, reprocessed into a baryon asymmetry by sphalerons. GG-A10-Cosmos's baryogenesis section cites it as the mechanism that ties the baryon asymmetry to the inherited neutrino and CP structure that BI-6 requires, anchoring the $\eta_B \leftrightarrow \Omega_b$ chain. First cited in §4.1.

M. Garny, A. Kartavtsev, and A. Hohenegger. Leptogenesis from first principles in the resonant regime. *Annals Phys.* **328** (2013) 26–63 [arXiv:1112.6428].
Relevance. First-principles closed-time-path (Kadanoff–Baym) treatment of the resonant regime, deriving the in-medium effective regulator of the self-energy asymmetry — the class of computation the C4 work order of Appendix I names as decisive. First cited in Appendix I.

G. W. Gibbons and S. W. Hawking. Cosmological event horizons, thermodynamics, and particle creation. *Phys. Rev. D* **15** (1977) 2738–2751.
Relevance. The 1977 Gibbons–Hawking paper established that a de Sitter cosmological event horizon carries a temperature and an area entropy, extending black-hole thermodynamics to the cosmological horizon. GG-A10-Cosmos cites it in the dark-energy section as the basis for treating the late vacuum-dominated, de Sitter-like phase thermodynamically, and it underlies the horizon-entropy reasoning of the black-hole entropy appendix. First cited in §5.

A. H. Guth and S.-Y. Pi. Fluctuations in the new inflationary universe. *Phys. Rev. Lett.* **49** (1982) 1110–1113.
Relevance. The 1982 Guth–Pi paper showed how quantum fluctuations during inflation are stretched and frozen into classical density perturbations with a calculable spectrum. GG-A10-Cosmos's inflation appendix cites it within the 1982 cluster that fixed the standard inflationary prediction of a nearly scale-invariant, Gaussian, adiabatic spectrum. First cited in §G.

Guth, A. H. (1981). Inflationary universe: A possible solution to the horizon and flatness problems. *Physical Review D, 23*(2), 347–356.

Relevance. The 1981 paper that introduced inflation, showing that an early epoch of accelerated expansion driven by a metastable false vacuum solves the horizon, flatness, and unwanted-relic problems of the hot Big Bang. GG-A10-Cosmos cites this to set the standard problem-setting that the present paper reinterprets geometrically rather than discarding, retaining the slow-roll content of its "new inflation" descendants while deriving the launch geometrically, at the GUT edge of the inherited bulk. First cited in §2.

Hawking, S. W. (1971). Gravitationally collapsed objects of very low mass. *Monthly Notices of the Royal Astronomical Society, 152*(1), 75–78.
Relevance. The 1971 founding paper of primordial-black-hole physics, proposing that large density perturbations in the very early universe could collapse directly into black holes without passing through stellar evolution. GG-A10-Cosmos cites this as historical background for the PBH section, which asks which slice of the resulting mass spectrum, if any, is consistent with the inherited ledger and could contribute to dark-matter, gravitational-wave, or seed-black-hole physics without conflicting with the rest of the cosmic audit. First cited in §49.1.

Hawking, S. W. (1975). Particle creation by black holes. *Communications in Mathematical Physics, 43*(3), 199–220.
Relevance. The 1975 paper deriving thermal Hawking radiation and establishing the temperature–horizon relation that completed black-hole thermodynamics. GG-A10-Cosmos cites this as the origin of the thermal-horizon problem used as the black-hole audit in Part V, where the evaporation ledger is reread through GFF and GDOS rather than taken as an isolated semiclassical result. First cited in §4.1.

S. W. Hawking. The development of irregularities in a single bubble inflationary universe. *Phys. Lett. B* **115** (1982) 295–297.
Relevance. Hawking's 1982 paper is one of the founding 1982 calculations relating the inflationary spectrum of density perturbations to the slow-roll dynamics of the inflaton. GG-A10-Cosmos's inflation appendix cites it among the converging derivations that fixed the amplitude and near-scale-invariance of the primordial spectrum used downstream. First cited in §G.

Hill, J. C., McDonough, E., Toomey, M. W., & Alexander, S. (2020). Early dark energy does not restore cosmological concordance. *Physical Review D*, *102*(4), 043507.
Relevance. Demonstrates that full-likelihood early-dark-energy fits drive the net S_8 shift positive; cited as the mandatory sign caveat on the retired corridor's S_8 suppression branch (record, Appendix K.10). First cited in §43.

S. Hollands and R. M. Wald. An alternative to inflation. *Gen. Rel. Grav.* **34** (2002) 2043–2055.
Relevance. Hollands and Wald argue that inflation's explanations of homogeneity, isotropy, and flatness are not fully satisfactory, and that a proper account requires a deeper theory of the initial state. GG-A10-Cosmos cites this in the second-horizon-problem Section to locate the present claim *outside* that critical tradition: there the target is a deeper half of the *first* (conditions) problem, in the initial state and its entropy, whereas here inflation's 1981 success is granted in full and a

second, differently shaped problem — about things, not conditions — is posed beside it. First cited in §72.2.

Hosotani, Y. (1983). Dynamical mass generation by compact extra dimensions. *Physics Letters B, 126*(4), 309–313.
Relevance. The 1983 paper showing that a Wilson line wrapping a compact extra dimension can dynamically break gauge symmetry and generate mass — the Hosotani mechanism. GG-A10-Cosmos cites this as historical context for the holonomy-first reading of cosmic branch selection, in which the compact-angle Wilson line of the inherited GG-6-6 bulk, rather than a local potential, fixes which branch the cosmos occupies. First cited in §14.1.

Hosotani, Y. (2005). Dynamical gauge symmetry breaking as the Hosotani mechanism. *International Journal of Modern Physics A, 20*(22), 5068–5104.
Relevance. The 2005 review consolidating gauge-Higgs unification and the dynamics of Wilson-line symmetry breaking. GG-A10-Cosmos cites this as historical context for the Wilson-loop GSSB and holonomy variables of the cosmic-grip grammar, supplying the established field-theory background against which the compact-angle branch coordinates of the present paper are introduced. First cited in §14.1.

Inflation at the end of 2025: constraints on r and n_s using the latest CMB and BAO data, arXiv:2512.10613 (2025).
Relevance. The current joint inflationary constraint from SPT-3G, *Planck*, ACT and BICEP/Keck, with and without DESI baryon-acoustic oscillations: $r < 0.034$ at 95%, which leaves the derived 0.0038 (0.0043 at leading order) untouched, and $n_s = 0.9682 \pm 0.0032$ (CMB alone) or 0.9728 ± 0.0029 (with DESI), against which the derived tilt sits 1.5σ and 3.2σ low (§28.10). Cited because the comparison must be made against the data that exist, not only against *Planck* 2018.

W. Hu and N. Sugiyama. Anisotropies in the cosmic microwave background: an analytic approach. *Astrophys. J.* **444** (1995) 489–506.
Relevance. The 1995 Hu–Sugiyama paper provided the analytic acoustic-oscillation theory — sound horizon, baryon loading, and peak locations — that underlies modern CMB phenomenology. GG-A10-Cosmos's CMB appendix follows their analytic treatment closely: the sound speed c_s, the sound horizon r_s, the baryon-loading parameter R, and the peak-location formula $\ell_n \approx n\pi D_A/r_s$ are all developed in their framework. First cited in §3.

W. Hu and N. Sugiyama. Small-scale cosmological perturbations: an analytic approach. *The Astrophysical Journal* **471** (1996) 542–570.
Relevance. The standard analytic fitting form for the decoupling redshift, accurate to better than a percent against full Boltzmann codes. Used in §42.1 as the cross-check of the framework's atomic-chain computation of $z_{\rm dec}$ — the quantity whose 0.1% zero-parameter agreement with the measured value is that section's result. First cited in §42.1.

Hubble, E. (1929). A relation between distance and radial velocity among extra-galactic nebulae. *Proceedings of the National Academy of Sciences, 15*(3), 168–173.

Relevance. The 1929 paper that gave the first observational evidence for cosmic expansion through a roughly linear distance-redshift relation for extra-galactic nebulae. GG-A10-Cosmos cites this as historical background for the expansion-and-cooling section, which retains the Friedmann equations and the Hubble law as the leading-order operational law and places all new content in the inherited source ledger feeding the right-hand side. First cited in §1.

W. Huggins and W. A. Miller. On the spectra of some of the fixed stars. *Phil. Trans. R. Soc.* **154** (1864) 413–435.

Relevance. The founding demonstration that the stars shine with the same spectral lines as terrestrial elements — the same atomic physics across the visible universe. GG-A10-Cosmos cites it as the opening datum of the second horizon problem: the evidence for the sameness of *things* predates the evidence for the sameness of *conditions* (the 1965 uniform temperature) by a full century, yet the temperature's paradox was named within a few years and the identity's never was. First cited in §72.1.

T. Jacobson. Thermodynamics of spacetime: the Einstein equation of state. *Phys. Rev. Lett.* **75** (1995) 1260–1263.

Relevance. Jacobson's 1995 paper derived the Einstein field equations as an equation of state from the proportionality of horizon entropy to area plus the Clausius relation, recasting gravity as thermodynamics. GG-A10-Cosmos cites it in the beginning/dark-energy sections as the conceptual precedent for reading the Einstein/Friedmann equations as a coarse-grained seam-level balance rather than a fundamental postulate, resonant with the CCI interpretation of FRW continuity. First cited in §5.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review, 106*(4), 620–630.

Relevance. The 1957 paper recasting equilibrium statistical mechanics as maximum-entropy inference under constraints, bridging entropy, information, and constrained maximization. GG-A10-Cosmos cites this for orientation in the discussion of GFF, the Qi energy-principle, and inner reconstruction; the maximum-entropy reading is a conceptual analog for the free-energy bookkeeping, not a derivational premise of the cosmology. First cited in §106.

H. Jukkala, K. Kainulainen, and P. M. Rahkila. Flavour mixing transport theory and resonant leptogenesis. *JHEP* **09** (2021) 119 [arXiv:2104.03998].

Relevance. A local quantum transport theory for the mixing heavy sector, resolving the coherence information that the flavor-measurement hypothesis of Appendix I turns on; among the sharpest existing statements of how the heavy-pair coherence is damped in-medium. First cited in Appendix I.

S. Kanemura and S.-P. Li. Resonant forbidden CP asymmetry from soft leptons. *Phys. Rev. Lett.* **134** (2025) 211802.

Relevance. Opens the 2025–26 thermal-resonant line: thermal resummation of *Standard-Model soft leptons* produces a resonantly enhanced CP asymmetry stable under finite-temperature perturbation theory — a flavor-odd source clocked by lepton-sector rates. Cited in Appendix I as

published support for the structure of conditions C1–C2 (and as the opening of the competitor-mechanism question the C4 work order registers). First cited in Appendix I.

T. W. B. Kibble. Topology of cosmic domains and strings. *J. Phys. A* **9** (1976) 1387–1398.
Relevance. Kibble's 1976 paper made the causal argument general and quantitative: at any cosmological symmetry-breaking transition the vacuum choice is made independently in regions with no prior causal contact, the correlation length is bounded by the horizon, and defects form where the choices disagree. GG-A10-Cosmos cites it in the second-horizon-problem pedigree as the sharpest early posing of the identity question — for the *orientation* of a degenerate vacuum, framed as a relic-abundance problem — fifty years before the question is named and answered in its own right in §72. First cited in §72.2.

J. Klarić, M. Shaposhnikov, and I. Timiryasov. Uniting low-scale leptogenesis mechanisms. *Phys. Rev. Lett.* **127** (2021) 111802 [arXiv:2008.13771].
Relevance. Demonstrates that resonant leptogenesis and leptogenesis via oscillations are limits of one kinetic framework spanning the GeV–TeV heavy-mass range that includes the lock scale; the unified equations are a benchmark any C4-discharging computation must reduce to. First cited in Appendix I.

H. Kodama and M. Sasaki. Cosmological perturbation theory. *Prog. Theor. Phys. Suppl.* **78** (1984) 1–166.
Relevance. Kodama and Sasaki's 1984 review is the comprehensive systematic reference for the gauge-invariant multi-fluid theory of cosmological perturbations. GG-A10-Cosmos's growth appendix follows its organization of the continuity, Euler, and Poisson equations into the second-order master growth equation; the framework adopts this standard apparatus unmodified and acts only at the late-time normalization. First cited in §3.

E. W. Kolb and M. S. Turner. *The Early Universe.* Addison-Wesley, Redwood City, (1990).
Relevance. Kolb and Turner's 1990 monograph is the canonical reference for early-universe thermodynamics: the relativistic degree-of-freedom counting $g_*(T)$, entropy conservation $g_{*s}a^3T^3 =$ const, freeze-out, and the relic-abundance machinery. GG-A10-Cosmos uses it throughout the thermal-history appendix and in the Kaxion and BBN appendices (freeze-out temperatures, the neutron-to-proton ratio, and the misalignment relic estimate). First cited in §3.

Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *Annals of Mathematical Statistics, 22*(1), 79–86.
Relevance. The 1951 paper introducing the relative-entropy (Kullback–Leibler) divergence as a measure of statistical distinguishability. GG-A10-Cosmos cites this for the entropy and inference language surrounding the CEP sector of the free-energy functional F_{GFF}, where divergence-style bookkeeping frames the cost of distinguishing boundary states; it is orienting vocabulary, not a derivational input. First cited in §106.

V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov. On anomalous electroweak baryon-number non-conservation in the early universe. *Phys. Lett. B* **155** (1985) 36–42.

Relevance. The 1985 Kuzmin–Rubakov–Shaposhnikov paper established that electroweak sphaleron processes violate baryon-plus-lepton number and are unsuppressed at high temperature, a cornerstone of modern baryogenesis. GG-A10-Cosmos's baryogenesis section reads sphaleron conversion as one of the seam-level source terms of the open projected ledger and cites this paper for the sphaleron transport physics. First cited in §4.1.

Lakatos, I. (1978). *The Methodology of Scientific Research Programmes: Philosophical Papers, Volume 1* (eds. J. Worrall & G. Currie). Cambridge University Press.
Relevance. The distinction between *progressive* and *degenerating* research programs — forecasting novel facts versus accommodating known ones by adjustment. GG-A10-Cosmos cites it in Objection O-Cosmos-90 to state the verdict on the vacuum-selection program precisely: non-progressive rather than refuted. First cited in §O9.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development, 5*(3), 183–191.
Relevance. The 1961 paper establishing the thermodynamic cost of logically irreversible operations — the Landauer bound linking erasure to dissipation. GG-A10-Cosmos cites this as external context for why information-bearing subsystems in the boundary runtime still obey energetic and entropic bookkeeping, reinforcing the GFF ledger framing without entering the cosmological derivation as a premise. First cited in §106.

Lemaître, G. (1927). Un univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Société Scientifique de Bruxelles, A47*, 49–59.
Relevance. The 1927 paper that independently derived the expansion law from observational considerations and supplied the operational language of a beginning later sharpened into the primeval atom. GG-A10-Cosmos cites this as historical background for the beginning-of-universe and expansion-and-cooling sections, both of which retain the FRW dynamics of Friedmann and Lemaître intact and ask only the narrower question of at what scale the first physically meaningful boundary surface is initialized. First cited in §1.

S.-P. Li and A. Pilaftsis. Dominant thermal resonant mechanism for low-scale leptogenesis. arXiv:2601.15921 [hep-ph] (2026).
Relevance. The thermal-resonant mechanism in full: flavor-odd CP sources clocked by charged-lepton Yukawa rates, with gauge damping dropping out of the CP-odd channel — the strongest published support for the lepton-sector clock of Appendix I, while leaving the cross-sector matching of condition C4 open. Registered there as a potential competitor to the ΔM-filtered source at $T \sim M_N$. First cited in Appendix I.

S.-P. Li and A. Pilaftsis. Low-scale leptogenesis from resonant thermal lepton flavour coherences. arXiv:2604.06493 [hep-ph] (2026).
Relevance. Companion of the preceding entry: the resonant enhancement carried by *thermal lepton flavor coherences* rather than heavy-sector quasi-degeneracy. Cited with it in the C4 work order of Appendix I. First cited in Appendix I.

A. R. Liddle. *An Introduction to Modern Cosmology*, 3rd ed. Wiley, Chichester, (2015).
Relevance. A gentle, physically transparent first course on the expanding universe and the inventory; recommended as the first-pass companion to Appendices E and F for students meeting cosmology for the first time. First cited in §A.

LiteBIRD Collaboration (Allys, E., et al.). (2023). Probing cosmic inflation with the *LiteBIRD* cosmic microwave background polarization survey. *Progress of Theoretical and Experimental Physics*, *2023*(4), 042F01; arXiv:2202.02773.
Relevance. The LiteBIRD forecast, a total uncertainty $\delta r < 0.001$ for $r = 0$, is the satellite design floor named beside the Simons Observatory's in Parts IV and VII; at $\sigma(r) \simeq 10^{-3}$ the derived $r = 0.0038$ is a several-standard-deviation detection or a refutation, which is the two-sided statement the tensor test makes. First cited in §28.

A. R. Liddle and D. H. Lyth. *Cosmological Inflation and Large-Scale Structure*. Cambridge University Press, Cambridge, (2000).
Relevance. The standard in-depth treatment of the inflationary sector and its perturbation output; recommended companion to Appendix G, whose Step-1 slow-roll apparatus follows its conventions before the geometric Step 3 replaces the potential's freedom with the derived plateau. First cited in §A.

E. M. Lifshitz. On the gravitational stability of the expanding universe. *J. Phys. (USSR)* **10** (1946) 116.
Relevance. Lifshitz's 1946 paper is the foundational relativistic treatment of the gravitational instability of an expanding universe, establishing that density perturbations grow as a power law rather than exponentially in a matter-dominated FRW background. GG-A10-Cosmos's growth appendix recovers exactly this result — the growing-mode $D \propto a$ in the Einstein–de Sitter regime — as the leading line of the linear-growth derivation. First cited in §3.

Linde, A. D. (1982). A new inflationary universe scenario. *Physics Letters B, 108*(6), 389–393.
Relevance. The 1982 "new inflation" paper that replaced the false-vacuum picture of Guth (1981) with a slow-rolling scalar field, repairing the graceful-exit problem. GG-A10-Cosmos cites this to frame the standard scalar-potential paradigm that the present paper retains in slow-roll form but replaces at the level of mechanism by a geometric launch-and-exit picture, fixing the inflationary modulus within the inherited bulk-modulus space. First cited in §2.

A. D. Linde. Chaotic inflation. *Phys. Lett. B* **129** (1983) 177–181.
Relevance. Linde's 1983 chaotic-inflation paper showed that successful inflation does not require a fine-tuned false vacuum but can proceed from generic large-field initial conditions. GG-A10-Cosmos's inflation appendix cites it as the alternative large-field paradigm against which the framework's inherited plateau-type stage and its predicted small r are contrasted. First cited in §2.

A. D. Linde. *Particle Physics and Inflationary Cosmology*. Harwood Academic, Chur (1990); arXiv:hep-th/0503203.

Relevance. Linde's monograph lists "the problem of the uniqueness of the universe" — why the entire universe sits in the same broken-symmetry minimum — among the problems inflation addresses, and contains the nearest published approach to the second horizon problem: the remark that no anthropic selection explains why all properties of the universe are approximately uniform over its entire observable part. GG-A10-Cosmos cites it in the pedigree of §72 as the closest near miss: the sentence was written, and not followed — the uniformity of *things* never joined the canonical problem list, and the answer on offer remained spatial (one stretched domain), not temporal (the ordering). First cited in §72.2.

D. H. Lyth and A. Riotto. Particle physics models of inflation and the cosmological density perturbation. *Phys. Rept.* **314** (1999) 1–146.
Relevance. The 1999 Lyth–Riotto review is the standard reference connecting particle-physics model building to the observable inflationary perturbation spectrum, including the slow-roll consistency relations $n_s = 1 - 6\epsilon + 2\eta$ and $r = 16\epsilon$. GG-A10-Cosmos's inflation appendix uses exactly these relations and cites this review for the systematic mapping from a potential $V(\phi)$ to (n_s, r). First cited in §3.

R. A. Lyttleton and H. Bondi. Physical consequences of a general excess of charge. *Proc. Roy. Soc. Lond. A* **252** (1959) 313.
Relevance. The classic charged cosmology: a proton–electron mismatch of two parts in 10^{18} would drive the expansion electrostatically. The proposal marks the first serious recognition that a gauged excess cannot hide, and the neutrality theorem of §36.13 states why the excess is exactly zero rather than merely small.

Maldacena, J. (1998). The large-N limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics, 2*(2), 231–252.
Relevance. The 1998 paper conjecturing the AdS/CFT correspondence, the canonical statement of holographic gauge-gravity duality. GG-A10-Cosmos cites this as context for the black-hole and hologram discussion that situates the Part V horizon audit and the island comparison; it frames the holographic literature without being used as a premise of the present derivation. First cited in §4.1.

Maldacena, J. (1999). Eternal black holes in anti-de Sitter. *Journal of High Energy Physics, 2003*(04), 021.
Relevance. The paper on eternal AdS black holes and the thermofield-double construction relating two boundary copies through a shared horizon. GG-A10-Cosmos cites this to situate the horizon discussion of Part V within the broader holographic literature on black-hole structure; like Maldacena (1998) it is comparative context, not a derivational input. First cited in §48.11.

D. J. E. Marsh. Axion cosmology. *Phys. Rept.* **643** (2016) 1–79.
Relevance. Marsh's 2016 review is the comprehensive modern reference for axion cosmology, covering the misalignment mechanism, structure formation with axion dark matter, and observational constraints. GG-A10-Cosmos's Kaxion appendix cites it for the up-to-date relic-abundance

formulae and the cosmological constraints that any Kaxion dark-matter candidate must satisfy. First cited in §3.

C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. W. H. Freeman, San Francisco, (1973).

Relevance. The 1973 "MTW" text is the standard reference for the tensor-calculus machinery (metric, connection, curvature, Bianchi identities) on which the FRW background appendix and the black-hole entropy appendix rest. GG-A10-Cosmos follows its conventions in carrying out the explicit Christoffel-symbol and Einstein-tensor computation without skipping index work; the contracted Bianchi identity it supplies underwrites the cohort reading of FRW continuity as the isotropic reduction of CCI. First cited in §3.

V. F. Mukhanov and G. V. Chibisov. Quantum fluctuations and a nonsingular universe. *JETP Lett.* **33** (1981) 532–535.

Relevance. The 1981 Mukhanov–Chibisov paper is the first calculation of the quantum-generated, nearly scale-invariant spectrum of cosmological perturbations from inflation. GG-A10-Cosmos's inflation appendix recovers the same curvature-perturbation spectrum and uses it to seed the growth appendix; the framework adopts the quantum origin of the spectrum and reads it as a boundary projection. First cited in §2.

V. Mukhanov. *Physical Foundations of Cosmology*. Cambridge University Press, Cambridge, (2005).

Relevance. Mukhanov's 2005 textbook is a primary reference for the gauge-invariant theory of cosmological perturbations and the quantum origin of inflationary fluctuations. GG-A10-Cosmos draws on it in the inflation appendix (slow-roll parameters, the curvature-perturbation spectrum, n_s and r) and the growth appendix; the framework adopts the textbook generation of the primordial spectrum and reads it as the boundary projection of the inherited bulk inflationary stage. First cited in §3.

V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger. Theory of cosmological perturbations. *Phys. Rept.* **215** (1992) 203–333.

Relevance. This 1992 Physics Reports review unified the gauge-invariant perturbation formalism with the quantum theory of inflationary fluctuations and the late-time growth of structure. GG-A10-Cosmos cites it as the connective reference linking the inflation appendix (origin of the primordial spectrum) to the growth appendix (its subsequent evolution into the matter power spectrum and σ_8). First cited in §2.

D. N. Page. Information in black hole radiation. *Phys. Rev. Lett.* **71** (1993) 3743–3746.

Relevance. Page's 1993 paper introduced the "Page curve" for the entanglement entropy of Hawking radiation, the benchmark statement of unitary black-hole evaporation and information recovery. GG-A10-Cosmos's black-hole information discussion and the entropy appendix cite it for the Page-curve criterion that any information-preserving evaporation history — including the cohort's bulk-boundary reading — must reproduce. First cited in §4.1.

J. A. Peacock. *Cosmological Physics*. Cambridge University Press, Cambridge, (1999).
Relevance. The encyclopaedic one-volume treatment of relativistic cosmology and structure formation; recommended for in-depth study alongside Appendices E, F, and J. First cited in §A.

Peccei, R. D., & Quinn, H. R. (1977). CP conservation in the presence of pseudoparticles. *Physical Review Letters, 38*(25), 1440–1443.
Relevance. The 1977 paper that solved the strong-CP problem by introducing a new global U(1) symmetry whose spontaneous breaking dynamically relaxes the QCD vacuum angle to zero. GG-A10-Cosmos cites this as historical background for the Kaxion dark-matter section, whose compact-angle pseudo-scalar inherits the Peccei–Quinn template through the GG-6-6 compact holonomy rather than through a postulated low-energy U(1). First cited in §4.1.

P. J. E. Peebles and B. Ratra. The cosmological constant and dark energy. *Rev. Mod. Phys.* **75** (2003) 559–606.
Relevance. The 2003 Peebles–Ratra review is the authoritative survey of the cosmological constant and dynamical dark-energy (quintessence) models. GG-A10-Cosmos's dark-energy section cites it to delimit the menu of standard explanations — Λ, quintessence, modified gravity — against which the cohort's one-ledger reading of the late vacuum residue is set. First cited in §3.

P. J. E. Peebles and J. T. Yu. Primeval adiabatic perturbation in an expanding universe. *Astrophys. J.* **162** (1970) 815–836.
Relevance. Peebles and Yu's 1970 paper predicted the baryon acoustic oscillation imprint in both the matter and the photon power spectra, establishing the sound horizon as a standard ruler. GG-A10-Cosmos cites it in the expansion section as the theoretical origin of the BAO scale that the one background must respect, and it underlies the CMB acoustic-peak appendix. First cited in §3.

P. J. E. Peebles. *Principles of Physical Cosmology*. Princeton University Press, Princeton, (1993).
Relevance. Peebles' 1993 treatise is a foundational reference spanning the FRW background, thermal history, BBN, and the statistics of large-scale structure. GG-A10-Cosmos cites it as a pedagogical and historical anchor for the FRW and thermal-history appendices and for the definition of the matter power spectrum and its normalization (σ_8) in the growth appendix. First cited in §3.

G. Penington. Entanglement wedge reconstruction and the information paradox. *J. High Energy Phys.* **09** (2020) 002.
Relevance. The quantum-extremal-surface derivation of the Page curve from semiclassical gravity, one of the two 2019 island papers. Cited at Part VII's island formula, which earlier versions of this paper sourced to Maldacena (1998, 1999); those are the holographic background, not the result.

Penrose, R. (1994). *Shadows of the Mind*. Oxford, UK: Oxford University Press.
Relevance. Penrose's 1994 book arguing that gravitation and a non-computable element are essential to consciousness, with objective state reduction set near a gravitational threshold. GG-A10-Cosmos cites this for the black-hole and consciousness context surrounding the later

pagoda and EIC discussion; it is an orienting reference for the boundary-runtime framing, not a derivational input. First cited in §51.

R. Penrose. *The Road to Reality.* Jonathan Cape, London (2004).

Relevance. Penrose's critique of inflationary explanation — that appealing to a thermalization process cannot explain the specialness of the initial state, whose entropy must have been extraordinarily low — is the best-known statement of the critical tradition that seeks a deeper half of the *first* horizon problem. GG-A10-Cosmos cites it in the second-horizon-problem Section to draw the boundary: the present claim grants inflation's 1981 success and poses a second, differently shaped problem about identity, rather than reopening the conditions problem Penrose targets. First cited in §72.2.

Penzias, A. A., & Wilson, R. W. (1965). A measurement of excess antenna temperature at 4080 Mc/s. *Astrophysical Journal, 142*, 419–421.

Relevance. The 1965 serendipitous discovery of the cosmic microwave background as an isotropic "excess antenna temperature at 4080 Mc/s". GG-A10-Cosmos cites this, together with the companion interpretation of Dicke et al. (1965), as the historical opening of the CMB section, which retains the standard acoustic, damping, and projection structure of the microwave background and asks only where the inherited ledger reshapes its source assumptions. First cited in §1.

Perlmutter, S., et al. (1999). Measurements of Ω and Λ from 42 high-redshift supernovae. *The Astrophysical Journal, 517*(2), 565–586.

Relevance. The 1999 Supernova Cosmology Project measurement that, with Riess et al. (1998), established the accelerating late-time expansion and the observational reintroduction of Λ. GG-A10-Cosmos cites this as the historical evidence for cosmic acceleration behind the dark-energy benchmark inherited by the paper, whose section retains the empirical Λ but reinterprets its scale through the residual GFF ledger of the inherited bulk. First cited in §1.

C. Pitrou, A. Coc, J.-P. Uzan, and E. Vangioni. Precision big bang nucleosynthesis with improved Helium-4 predictions. *Phys. Rept.* **754** (2018) 1–66.

Relevance. The 2018 Pitrou *et al.* review provides the most precise modern BBN predictions, in particular a high-accuracy ^{4}He computation including subtle weak-rate and radiative corrections. GG-A10-Cosmos's BBN appendix cites it as the precision benchmark for Y_p and for the neutron lifetime and freeze-out treatment underlying the helium yield. First cited in §3.

Planck Collaboration. (2018). *Planck 2018 results. VI. Cosmological parameters.* ESA/Planck legacy release.

Relevance. The Planck 2018 cosmological-parameter release, the canonical high-precision ΛCDM determination of H_0, Ω_m, Ω_Λ, S_8, and the CMB acoustic structure. GG-A10-Cosmos uses this as the principal empirical baseline that the present paper seeks to inherit and reinterpret: it supplies the comparators in the codebook and the parameter-audit Dev plot against which the one background and the computed clustering amplitude are scored. First cited in §29.

Poulin, V., Smith, T. L., Karwal, T., & Kamionkowski, M. (2019). Early dark energy can resolve the Hubble tension. *Physical Review Letters*, *122*(22), 221301.

Relevance. Source of the early-dark-energy response slopes (~ 0.4 per unit $f_{\rm EDE}$) adopted as representative coefficients in the retired corridor's linear-response form; the retired corridor module was a structured re-parametrization of this class of models, and the paper now carries no early residue ($f_{\rm EDE} = 0$). First cited in §42.

J. Preskill, M. B. Wise, and F. Wilczek. Cosmology of the invisible axion. *Phys. Lett. B* **120** (1983) 127–132.
Relevance. One of the three founding 1983 papers establishing that a light axion-like field acquires a cosmological cold-dark-matter abundance through the vacuum-misalignment mechanism. GG-A10-Cosmos's Kaxion appendix builds exactly this mechanism — the field frozen by Hubble friction at angle θ_i until $H \sim m$, then oscillating with $\rho \propto a^{-3}$ — to estimate the Kaxion relic density. First cited in §M.

W. H. Press and P. Schechter. Formation of galaxies and clusters of galaxies by self-similar gravitational condensation. *Astrophys. J.* **187** (1974) 425–438.
Relevance. The 1974 Press–Schechter formalism is the foundational analytic prescription for the abundance of collapsed halos as a function of the linear density-field variance. GG-A10-Cosmos's growth appendix invokes it to motivate why the smoothed amplitude σ_8 (and hence S_8) is the cosmologically decisive number: halo counts depend exponentially on it, which is exactly why the clustering tension is sharp. First cited in §3.

Randall, L., & Sundrum, R. (1999). A large mass hierarchy from a small extra dimension. *Physical Review Letters, 83*(17), 3370–3373.
Relevance. The 1999 paper that generated a large mass hierarchy from a extra dimension, with physical scales redshifted exponentially along the throat. GG-A10-Cosmos cites this for intuition about throat geometry and gravitational redshift, while the actual scale-geometry derivation used in the architecture section remains internal to the inherited GG-6-6 locking and is not imported from the Randall–Sundrum construction. First cited in §8.

Riess, A. G., et al. (1998). Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal, 116*(3), 1009–1038.
Relevance. The 1998 High-z Supernova Search measurement that, with Perlmutter et al. (1999), first established the accelerating expansion and reintroduced Λ observationally. GG-A10-Cosmos cites this as the standard evidence for cosmic acceleration and the late-time vacuum-inventory problem addressed in Part II, where the residual GFF ledger is asked to account for the inherited Λ scale. First cited in §1.

Riess, A. G., Yuan, W., Macri, L. M., et al. (2022). A comprehensive measurement of the local value of the Hubble constant with 1 km s^{-1} Mpc^{-1} uncertainty from the Hubble Space Telescope and the SH0ES team. *The Astrophysical Journal Letters*, *934*(1), L7.
Relevance. The SH0ES distance-ladder determination $H_0 = 73.04 \pm 1.04$ km/s/Mpc, quoted in this paper (rounded to 73.0 ± 1.0) as the leading single-team ladder determination; the codebook and the parameter-audit Dev plot score the one background against the H0DN consensus (Casertano et al., 2026), the ladder the predicted outlier. Distinct from the 1998 acceleration-discovery paper

by the same first author, which is cited separately for the historical reintroduction of Λ. First cited in §4.

B. Ryden. *Introduction to Cosmology*, 2nd ed. Cambridge University Press, Cambridge, (2017).
Relevance. The friendliest rigorous introduction to FRW kinematics, the inventory, and the thermal history; recommended as the gentlest entry point to the graduate course of Part XV (Appendices E and F). First cited in §A.

S. Ryu and T. Takayanagi. Holographic derivation of entanglement entropy from AdS/CFT. *Phys. Rev. Lett.* **96** (2006) 181602.
Relevance. The holographic formula equating entanglement entropy with the area of a minimal bulk surface, of which the quantum extremal surface is the quantum generalization. Cited at Part VII's island formula.

R. K. Sachs and A. M. Wolfe. Perturbations of a cosmological model and angular variations of the microwave background. *Astrophys. J.* **147** (1967) 73–90.
Relevance. The 1967 Sachs–Wolfe paper established how gravitational-potential perturbations at last scattering imprint temperature anisotropies on the CMB, fixing the large-angle (low-ℓ) plateau. GG-A10-Cosmos's CMB appendix uses the Sachs–Wolfe relation as the boundary condition that sets the zero point of the acoustic-peak structure below the first peak. First cited in §2.

Sakharov, A. D. (1967). Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe. *JETP Letters, 5*, 24–27.
Relevance. The 1967 paper stating the three necessary conditions — baryon-number violation, C and CP violation, and departure from thermal equilibrium — that any dynamical baryogenesis must satisfy. GG-A10-Cosmos cites this as historical background for the baryogenesis section, which asks not which mechanism is freely chosen but which Sakharov-compatible realization is admissible once the neutrino sector and the CP ledger have been fixed by the upstream particle paper. First cited in §4.1.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*(3), 379–423.
Relevance. The 1948 founding paper of information theory, defining entropy as the measure of a source's information content and the channel-capacity limit on reliable transmission. GG-A10-Cosmos cites this for the OH/IH/EIC language and the measurement/reconstruction interpretation of the boundary runtime in the later parts; it supplies the information-theoretic vocabulary, not a derivational premise of the cosmology. First cited in §106.

P. Sikivie. Experimental tests of the invisible axion. *Phys. Rev. Lett.* **51** (1983) 1415–1417.
Relevance. Sikivie's 1983 paper proposed the haloscope and helioscope detection strategies that define the experimental axion program. GG-A10-Cosmos's Kaxion appendix cites it for the experimental detectability of a misalignment dark-matter axion, relevant to whether the Kaxion is observationally accessible. First cited in §M.

J. Silk. Cosmic black-body radiation and galaxy formation. *Astrophys. J.* **151** (1968) 459–471.

Relevance. Silk's 1968 paper identified the photon-diffusion ("Silk") damping that exponentially suppresses CMB anisotropies below the diffusion scale, producing the high-ℓ damping tail. GG-A10-Cosmos's CMB appendix derives the Silk damping scale and the damping-tail multipole ℓ_D directly, citing this paper as the physical origin. First cited in §3.

G. F. Smoot *et al.* (COBE Collaboration). Structure in the COBE differential microwave radiometer first-year maps. *Astrophys. J. Lett.* **396** (1992) L1–L5.
Relevance. The 1992 COBE/DMR detection of CMB temperature anisotropies at the 10^{-5} level is the landmark confirmation that primordial perturbations of the predicted amplitude exist. GG-A10-Cosmos's CMB appendix cites it as the empirical anchor for the large-angle Sachs–Wolfe plateau and for the overall normalization of the primordial spectrum. First cited in §2.

Starobinsky, A. A. (1979). Spectrum of relict gravitational radiation and the early state of the universe. *JETP Letters, 30*, 682–685.
Relevance. The 1979 paper that first computed the spectrum of relict primordial gravitational waves from a de Sitter epoch and argued that they survive as a present-day stochastic background detectable through CMB B-modes or direct interferometry. GG-A10-Cosmos cites this as historical background for the primordial-gravitational-waves section, which uses the Starobinsky-channel amplitude-versus-scale argument that retired the once-assumed TeV-scale reading in favor of the derived GUT-adjacent plateau, turning the tensor sector into a sharp falsification target. First cited in §2.

A. A. Starobinsky. A new type of isotropic cosmological models without singularity. *Phys. Lett. B* **91** (1980) 99–102.
Relevance. Starobinsky's 1980 R^2 model is the first concrete inflationary scenario and remains among the best fits to the measured scalar spectral index and tensor-to-scalar ratio. GG-A10-Cosmos's inflation appendix references it as the prototype plateau potential and as the benchmark (n_s, r) target against which the framework's inherited inflationary stage is compared. First cited in §2.

G. Steigman. Primordial nucleosynthesis in the precision cosmology era. *Ann. Rev. Nucl. Part. Sci.* **57** (2007) 463–491.
Relevance. Steigman's 2007 review is the standard precision-era reference for the relation between the baryon-to-photon ratio η, the baryon density $\Omega_b h^2$, and the primordial ^{4}He and deuterium abundances. GG-A10-Cosmos's BBN appendix uses its tight $Y_p(\eta)$ and D/H(η) relations as the link in the $\eta_B \leftrightarrow \Omega_b$ consistency chain that the baryogenesis section must satisfy. First cited in §3.

R. A. Sunyaev and Ya. B. Zeldovich. Small-scale fluctuations of relic radiation. *Astrophys. Space Sci.* **7** (1970) 3–19.
Relevance. The 1970 Sunyaev–Zeldovich paper predicted the acoustic oscillations in the photon-baryon fluid that appear as the harmonic peak structure of the CMB power spectrum. GG-A10-Cosmos's CMB appendix builds the driven-oscillator description of the photon-baryon fluid and the resulting peak series ($\ell_1 \approx 220$) that this paper first anticipated. First cited in §3.

J.-P. Uzan. Varying constants, gravitation and cosmology. *Living Rev. Rel.* **14** (2011) 2.
Relevance. Uzan's Living Review is the canonical survey of observational bounds on time and space variation of the fundamental constants, including the quasar-absorption spectroscopy that constrains any spatial variation of the fine-structure constant across the observable universe at the 10^{-5}–10^{-6} level. GG-A10-Cosmos cites it as the measurement arm of the second horizon problem: a precision certification of identity uniformity, framed throughout that literature as an equivalence-principle test — uniformity as null hypothesis, never as a datum requiring a mechanism. First cited in §72.1.

L. Verde, T. Treu, and A. G. Riess. Tensions between the early and late Universe. *Nature Astron.* **3** (2019) 891–895.
Relevance. The 2019 Verde–Treu–Riess perspective is the concise canonical statement that early-universe (CMB-inferred) and late-universe (distance-ladder) determinations of H_0 disagree at high significance. GG-A10-Cosmos's Hubble-tension section cites it to frame the tension as a clash between two inference pipelines — precisely the "two dictionaries reading one background" distinction the present paper draws. First cited in §4.1.

R. V. Wagoner, W. A. Fowler, and F. Hoyle. On the synthesis of elements at very high temperatures. *Astrophys. J.* **148** (1967) 3–49.
Relevance. The 1967 Wagoner–Fowler–Hoyle paper is the foundational systematic computation of primordial nucleosynthesis, establishing the reaction network and the dependence of the light-element yields on the baryon-to-photon ratio. GG-A10-Cosmos's BBN appendix follows its network (the $n + p \to d + \gamma$ bottleneck through to ^{4}He) and its η-dependence to derive the helium mass fraction Y_p. First cited in §3.

Wald, R. M. (1984). *General Relativity.* Chicago, IL: University of Chicago Press.
Relevance. The standard graduate monograph on general relativity, fixing the modern conventions for curvature, horizons, and the Einstein equations. GG-A10-Cosmos cites this for the classical boundary baseline — the FRW and horizon machinery — that the present paper inherits rather than replaces, the geometric scaffold on which the architecture section builds the inherited GG-6-6 locking. First cited in §3.

Weinberg, S. (1972). *Gravitation and Cosmology.* New York, NY: Wiley.
Relevance. Weinberg's 1972 textbook, a canonical reference for the FRW/Friedmann dynamics and the observational grammar of cosmology. GG-A10-Cosmos cites this for the FRW/Friedmann baseline and the observational conventions the present paper retains, supplying the standard background against which the architecture section introduces the inherited bulk source ledger. First cited in §3.

Weinberg, S. (1978). A new light boson? *Physical Review Letters, 40*(4), 223–226.
Relevance. With Wilczek (1978), the 1978 identification of the axion as the light pseudo-Nambu-Goldstone boson necessarily produced by the spontaneous breaking of the Peccei–Quinn symmetry. GG-A10-Cosmos cites this as historical background for the Kaxion section, whose

compact-angle pseudo-scalar plays the cosmological dark-matter role of the axion while inheriting its Peccei–Quinn template through the GG-6-6 compact holonomy. First cited in §35.1.

Weinberg, S. (1987). Anthropic bound on the cosmological constant. *Phys. Rev. Lett.* **59**, 2607–2610.
Relevance. The one quantitative anthropic prediction — an upper bound on Λ from the requirement that gravitationally bound structures can form. GG-A10-Cosmos cites it in Objection O-Cosmos-90 as the single standing post-diction of the anthropic reading, set against a framework that derives the inventory. First cited in §O9.

Weinberg, S. (1989). The cosmological constant problem. *Reviews of Modern Physics, 61*(1), 1–23.
Relevance. The 1989 review surveying the dramatic mismatch between the observed vacuum energy and naturalness estimates from quantum field theory — the cosmological-constant problem. GG-A10-Cosmos cites this to frame the late-time cosmic ledger that the present paper attempts to reinterpret through the residual free-energy functional F_{GFF}, relocating the smallness of Λ from a tuning problem to a bulk-ledger residual. First cited in §3.

S. Weinberg. *Cosmology*. Oxford University Press, Oxford, (2008).
Relevance. Weinberg's 2008 graduate textbook is the modern reference derivation of the FRW background equations, thermal history, perturbation theory, and the CMB transfer function. GG-A10-Cosmos follows its conventions in the FRW background appendix (the Christoffel/Ricci/Einstein-tensor bookkeeping leading to the two Friedmann equations) and in the thermal-history appendix; the cohort reading then reinterprets the unmodified textbook FRW source terms as the isotropic reduction of the seam conservation law CCI. First cited in §2.

Wilczek, F. (1978). Problem of strong P and T invariance in the presence of instantons. *Physical Review Letters, 40*(5), 279–282.
Relevance. The 1978 companion to Weinberg (1978), independently identifying the axion as the pseudo-Nambu-Goldstone boson of the Peccei–Quinn mechanism. GG-A10-Cosmos cites this as historical background for the Kaxion section, whose compact-angle dark-matter candidate inherits the axion template through the GG-6-6 compact holonomy rather than a postulated low-energy U(1). First cited in §35.1.

Wilson, K. G. (1974). Confinement of quarks. *Physical Review D, 10*(8), 2445–2459.
Relevance. The 1974 paper that introduced the Wilson loop as the gauge-invariant order parameter for confinement in lattice gauge theory. GG-A10-Cosmos cites this to motivate Wilson-line and holonomy variables as genuine gauge-invariant branch coordinates in the cosmic-grip grammar, where the compact-angle holonomy of the inherited GG-6-6 bulk labels the realized cosmological branch. First cited in §14.1.

K. C. Wong *et al.* (H0LiCOW Collaboration). H0LiCOW XIII. A 2.4% measurement of H_0 from lensed quasars. *Mon. Not. R. Astron. Soc.* **498** (2020) 1420–1439.

Relevance. The 2020 H0LiCOW time-delay-cosmography measurement is an independent, distance-ladder-free determination of H_0 from strongly lensed quasars that favors the high (late-universe) value. GG-A10-Cosmos's Hubble-tension section cites it because lensing time-delay determinations are named as one of the ladder-free late probes predicted to converge on the background. First cited in §42.

Ya. B. Zeldovich, I. Yu. Kobzarev, and L. B. Okun. Cosmological consequences of a spontaneous breakdown of a discrete symmetry. *Zh. Eksp. Teor. Fiz.* **67** (1974) 3–11 [*Sov. Phys. JETP* **40** (1975) 1–5].

Relevance. The first explicit posing of vacuum identity as a causal problem: during the cooling of "causally disjointed points" of the universe the sign of the condensate is chosen at random, producing a domain structure of the vacuum — and the authors even ask how early the walls must be removed to make the whole universe a single domain. GG-A10-Cosmos cites this 1974 paper as the opening of the fifty-year pedigree traced in §72: the identity question, posed for one corner of identity, and filed under relic abundances. First cited in §72.2.