

The Hubble Tension in the Light of GG-Theory

The Hubble Constant as One Reading of a Derived Expansion History $H(t)$

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 a century, “how fast is the universe expanding?” was a question answered by measuring. By 2026 the measuring had split: read from the afterglow of the Big Bang, the rate comes out near 67 (in astronomers’ units); read from the local ladder of stellar distances, 73.50 ± 0.81 — a seven-sigma gap that no single method explains.

This paper reads the disagreement through GG-Theory, in which the expansion rate was never a number to be measured but a function of time to be derived: a six-dimensional geometry fixes the cosmic recipe as exact expressions in π , Einstein’s equations turn that recipe into the entire history $H(t)$, and today’s value is one snapshot of the curve — expressible in a single analytic line in the three laboratory constants $\{c, G_N, \hbar\}$ and two integers of the theory’s ledger, giving 66.8 ± 0.2 — on the afterglow side of the split, 1.1σ from the CMB combination. More than ten of the standard model’s fitted numbers re-emerge, unfitted, at the percent level. The early measurement needs no new physics.

The local 73.5 is then the puzzle — and not, this paper’s audit finds, a settled consensus: the 2026 landscape’s forty-five determinations span 67–73, the seven sigma certifying one superb network’s internal agreement. The paper reads the local value as the framework predicts it: one background rate, and a ten-percent calibration systematic in the ladder’s rungs that the framework predicts and does not derive — confronted with the field’s own published corrections, which do not yet reach it.

The referee is scheduled: gravitational-wave *standard sirens* read the rate with no ladder and no afterglow, and GG-Theory predicts they converge near 67 — about twenty-two events from now. Near 73, this paper is wrong, and says so in advance. Either way the century-old “constant” has finished its work: what remains to be tested is not a number but a law.

Executive Summary

The Hubble tension is usually told as a duel between two measurements, and the reader is asked to pick a side. This report tells it the other way around: in the GG-Theory series the expansion rate is not measured and then explained — it is *derived*, and once it is derived the tension becomes a diagnosis: one measurement is reading the true rate, and the other carries a bias whose size, sign, and possible origins can be written down and priced. This summary walks that argument in plain language; every step is worked in full inside.

Three results organize this report. *First, the Hubble parameter is derived directly, as a function of time.* The geometry fixes the universe’s contents in closed form, Einstein’s equations turn the contents into the whole history $H(t)$, and the famous “constant” is demoted to a snapshot — the value the falling curve happens to take today. That snapshot collapses to one analytic line,

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

three laboratory constants, two integers of the theory’s ledger ($L = 5$, $K = 35$), and the two transcendentals π and e — checkable on a calculator in under a minute. And the same ledger keeps going: more than ten of the numbers Λ CDM must fit — the matter, vacuum, and baryon fractions, the recombination epoch, the age, the damping scale, the growth amplitude among them — re-emerge, unfitted, at the percent level, the growth amplitude re-emerging with the same S_8 tension the standard model has. *Second, the seven sigma is not the whole field.* The celebrated 7σ belongs to one network’s internal consistency; this paper’s audit of all forty-five determinations published by 2026 finds them spanning 67–73, organized into three bands, with a middle band near 70 steadily filling as the field’s own corrections land. *Third, the century-old concept has finished its work.* If the rate is a derived function, the constant is obsolete the way measured molecular weights became obsolete after the periodic law — and cosmology can move on to a new, falsifiable precision program beyond Λ CDM, with a dozen dated tests in place of one contested decimal.

The Hubble “constant” is not a constant. The expansion rate is a *function of time*: enormous in the hot early universe, 66.8 today, settling forever onto a floor of 55.2. What astronomers call H_0 is simply where the falling curve crosses the present — a shorthand kept from the century when only that one point could be measured. Much of the tension’s mystery dissolves once the question is restated: not “what is the constant?” but “what is the curve, and are we reading one of its points wrongly?”

The early measurement needs no new physics. The derived rate lands 1.1σ from the value inferred from the cosmic microwave background — and the test is sharper than agreement: fed the exact π -locked recipe with no background parameter tuned to the sky, the theory reproduces the afterglow’s acoustic pattern to about a percent — a consistency at the percent level, not a fit; taken as exact the chart is excluded, and Part III prints that exclusion itself, at full strength — and where it misses, it misses *coherently*: seventeen observables, one direction, one small named step. A lucky

fit scatters its errors; a derivation displaced by one named correction lands on one side — which is what the data show. Nothing needs to be added to the young universe.

The local measurement, priced rather than blamed. The distance ladder’s 73.50 ± 0.81 is a magnificent measurement — thirty-eight specialists, every method cross-checked against every other — and this paper does not wave it away. It prices it. The framework derives none of the ten-percent excess and refuses to invent physics for any of it: the linear response of a local under-density is computed, and the depth that would let it carry a share of the excess is not observed, so no lift is carried. What the framework does claim is where the excess lives — on one background every rung-free probe reads the same rate, so the displaced instrument must be the one with rungs — and it names the only places such a bias could live, the shared calibrations no cross-checking can see, and works through them against the ladder’s own error budget: a dust law the infrared will strip away, a statistical prior a simulation can test, a choice of calibrator galaxies a complete JWST sample will settle; each with a published correction, a published rebuttal, and now a dated test.

The whole field, on one page. Rather than argue measurement by measurement, the paper draws them all: every 2026 determination of the expansion rate — forty-five values from nine kinds of measurement — in one chart and one table. The field organizes itself into three bands: a tight early-universe band around 67–68, the distance-ladder band at 73–74, and — new this decade — a middle band near 70 filling steadily with corrected ladders, dust-free supernovae, fast radio bursts, and gravitational-wave values. The derived rate sits at the floor of the first band; the corrected ladders at the floor of the middle one, still two to three units above it. Six families of proposed fixes are read against the same chart, and each is asked the one question the bands make unavoidable: not “67 or 73?” but “*why three bands — and why is the middle one filling?*”

The referee, and the date. It can be settled from outside. A merging pair of neutron stars broadcasts its own absolute distance in gravitational waves — Einstein’s equations themselves certify the luminosity, so there is no ladder, no calibration, no afterglow anywhere in the chain. One such “standard siren” has already been read, straddling both camps with fifteen-percent precision. About twenty-two such events separate the camps at 3σ , about sixty-one at 5σ , and the coming decade’s observing runs are expected to supply their equivalent — bright if the counterparts cooperate, dark and statistical if not. The prediction is printed and cannot be revised: the sirens converge near 67. Near 73, this paper is finished — and a catalog converging *between* the camps would kill both readings at once.

Beyond the constant: the new precision cosmology. Retiring the constant is not a loss; it is a promotion — for the whole field. The same geometry that prints 66.8 prints a falsifiable number nearly everywhere the coming decade will look: the primordial gravitational-wave amplitude ($r = 0.0038$, two-sided); an exact dark-energy equation of state ($w = -1$); a dark-matter particle with a laboratory address ($8.1 \mu\text{eV}$, near 1.95 GHz, one line standing alone); the growth amplitude ($S_8 = 0.825$, sharing the standard model’s tension rather than resolving it); the primordial tilt ($n_s = 0.9634$, today the framework’s one failing test, 3.2σ below the DESI-combined value); the

expansion history at every redshift, down to the centimeters-per-second-per-decade drift the largest telescopes will one day watch live. Twelve dated appointments — from Gaia’s astrometry this December to the siren verdict of the 2030s — replace the dispute over one decimal. This, the paper argues, is where the real excitement begins: precision cosmology whose predictions can *fail*, not precision metrology of a number no theory owned.

What is proved, and what is still assumed. The series keeps an explicit ledger, and this paper inherits its discipline. The π -locked recipe is a theorem *conditional on one named geometric lemma*, still open, and is labeled so everywhere it is used. The closed form’s last coefficient is selected by the theory’s own naming rules, its closure pre-registered as a falsifier. The local excess is not derived at all: the framework predicts it as a calibration systematic in the ladder, computes the one environmental channel it could have credited and finds the required depth unobserved, and hands the verdict to the sirens. Nothing in this summary is graded stronger inside the paper than it is stated here.

Why it matters. The deepest constants of physics have retired before: the speed of light was measured for two centuries until Maxwell derived it; the Rydberg was spectroscopy’s best number until Bohr computed it; the W and Z masses were predictions before they were measurements. Each retirement marked the moment a field stopped doing natural history and started doing law — and each time, the excitement moved from the next decimal to the places the new law could fail. The Hubble constant, on the reading defended here, is next in that lineage: the last measured constant of cosmology, retiring into a derived curve that began with a predicted burst, runs today on a predicted budget, and will end on a predicted calm. And there is a quieter thought beneath the argument, worked out in the closing Part: the same geometry that times the expansion timed the stars that forged our atoms and the chemistry that built our instruments — so a measurement of H_0 was never an outsider’s act performed on the universe from nowhere. We are not the audience of the expansion; we are the part of it that has begun to keep the log. That claim is either true or false, and unusually for claims of its size, the falsifying measurements are named, numbered, scheduled, and priced inside. Either the sirens read 67, and the oldest number in cosmology turns out to have been a theorem all along — or they read 73, and this paper becomes what a wrong theory should be: a clean, public, instructive ruin. Both outcomes are progress. Only one of them is possible.

Keywords: GG-Theory program; Hubble tension; H_0 ; derived Hubble parameter $H(t)$; closed-form H_0 ; H0DN; Local Distance Network; distance ladder; 2026 H_0 landscape; π -locked inventory; sound horizon; standard sirens; local void, priced and not carried; distance-ladder systematics; backreaction bound; S_8 tension; kill criteria; decision calendar; cosmological benchmarks; OGN-5 integer ledger; six-dimensional bulk; GDOS; beyond- Λ CDM precision cosmology.

**The 2026 H_0 Landscape: Every Reported Value in the Audited Reference Base
(GG predictions on top; the CMB combination as guide)**




Figure 1: The 2026 H_0 landscape at a glance: all forty-five audited determinations, the three bands, and the GG-6 predictions on top. Full version, table, and sources: Figure 20 in Part V.

The Hubble Parameter $H(t)$ Is Derived — H_0 Is One Snapshot of a Falling Curve




Figure 2: The rate is a derived function of time: the full history on log–log axes (top) and the late universe in linear time (bottom), falling monotonically to the de Sitter floor $H_\infty = 55.2 \text{ km s}^{-1} \text{Mpc}^{-1}$; the familiar $H_0 = 66.8$ is one reading of this curve, at $t_0 = 13.88$ Gyr. Full version: Figure 13 in Part III.

One Displacement, Seventeen Shadows: the Coherent π -Offset

Zero CMB input: every entry predicted by GG-6/GDOS/GRIP; the sky enters only as the comparison column.




Figure 3: One displacement, seventeen shadows: every entry is a GG-6/GDOS/GRIP prediction — zero CMB input, the sky only the comparison column — and the single $+1\%$ displacement of $1/\pi$ above Planck's fitted Ω_m propagates with the expected sign through all seventeen. A fit scatters its residuals; a derivation lands coherently. Full version: Figure 16 in Part III.

Ten Numbers, Each With Its Grade and Its Falsifier

The paper a reader can carry away. Ten quantities, each with its value, its honesty grade, and the observation that would retire it. Every row carries the grade and the falsifier in the same breath, because that pairing — a claim and the measurement that would kill it, printed together — is the discipline the whole twenty-two-paper series is built on. D* marks a value the ledger derives subject to the measure lemma; the four entries that are not D* say what they are instead. Two of the ten are retired by the same standard-siren result: near 73 this paper is wrong, publicly and quickly, and by its own printed rule. Nothing here is negotiable once its datum arrives, and most of the data arrive within the decade. In full: Table 23, p. 178.

Quantity	Value	Grade	Falsified by
Background rate $H_0^{(0)}$	66.8 ± 0.2	Derived (cond.)	a converged siren at 73
Matter fraction Ω_m	$1/\pi$	D*	Ω_m settling away from $1/\pi$
Dark energy Ω_Λ	$1 - 1/\pi$	D*	a measured $\Omega_\Lambda \neq 1 - 1/\pi$
Baryons Ω_b	$1/2\pi^2$	D*	$\Omega_b/\Omega_m \neq 1/2\pi$
Equation of state w	-1 exactly	Derived-cond.	a converged $w \neq -1$ at 3σ
Decoupling z_{dec}	1091.5	D*	a measured z_{dec} off the atomic chain by $> 1.4\%$
de Sitter floor H_∞	55.2	D*	follows from Ω_Λ
Present age t_0	13.9 Gyr	D*	inconsistent stellar/cosmic ages
Ladder excess ϵ_{syst}	$\approx 10\%$	Named, not derived	a ladder-free probe at 73.5
Siren prediction H_0^{GW}	$\rightarrow 67$	prediction	a converged siren at 73

Overview of Contents

Part I. Introduction	20
Part II. The Hubble Tension	35
Part III. The GG-Theory of the Hubble Parameter	50
Part IV. Reconciling the Hubble Tension with GG-Theory	86
Part V. The 2026 Measurement Landscape and the Rival Explanations	113
Part VI. Kill Switches, the Decision Calendar, and Wider Implications	133
Part VII. Forecasting the Verdict: The Standard-Siren Decade	149
Part VIII. Discussion: Significance, Uniqueness, and Innovation	157
Part IX. Conclusions	177
Part X. Technical Appendices	188
Part XI. Common Objections and Responses	238
Part XII. References	262

List of Figures

1	The 2026 H_0 landscape (showcase; full version in Part V).	5
2	The derived Hubble parameter $H(t)$ (showcase; full version in Part III).	6
3	The coherent π -offset Dev chart (showcase; full version in Part III).	7
4	A century of the Hubble constant: fall, stalemate, convergence, split.	23
5	The twenty-two-paper cohort map.	31
6	Roadmap of the twelve Parts of GG-C2-Hubble.	33
7	Anatomy of the two routes to H_0 .	40
8	The open-system funnel shared across the GG-Theory cohort.	42
9	The Hopf fibration behind Theorem T-Hubble-1.	59
10	The running budget $\Omega_i(t)$, shown twice.	61
11	The derived scale factor $a(t)$: the time dependence beneath $H(t)$.	63
12	The two routes to the present rate: geometry determines, the sky confirms.	71
13	The derived Hubble parameter $H(t)$, of which H_0 is one point.	74
14	Landmarks on the derived $H(t)$: nineteen decades on one curve.	75
15	The zero-parameter temperature spectrum against the Planck data.	82
16	The coherent π -offset: the parameter-audit Dev chart.	85
17	The lift a void would give: three published readings and the implied response.	96
18	The bridge from the derived rate to both measured values, with the residual's candidate anatomy.	105
19	The backreaction / cosmic-variance bound on the local excess.	106
20	The 2026 H_0 landscape: all 45 audited determinations, the three bands, and the GG predictions.	115
21	The decision calendar as a kill-test timeline for the derived Hubble reading.	140
22	The siren forecast: events to verdict, and the simulated catalog.	154
23	How this paper's names nest inside the program.	189

List of Tables

1	The twenty-two-paper GG-Theory cohort.	28
2	The two routes to H_0 , compared.	41
3	What the same five integers have already bought (the ledger’s track record).	52
4	The codebook selection of ζ_Λ , applied <i>before</i> any comparison with data.	66
5	The honesty triage for the derived Hubble rate: every object in the two-route derivation. . . .	73
6	The eight landmarks of the derived expansion history: epoch, the derived value and its origin. .	77
7	The cosmic-microwave-background observables predicted by GG-Theory (grade D* throughout). .	79
8	The scenario ledger for a local-void lift, recorded and not carried.	95
9	Systematic budget of the void channel itself (reproducibility format): each input, its range. .	95
10	Candidate common-mode calibration channels for the named residual: nature, published status. .	102
11	Cross-ledger coherence: the ingredients of the derived rate have sibling predictions in other. .	108
12	Sign structure of the growth reading on one background, as it now stands.	109
13	The 2026 H_0 landscape, tabulated: identifiers, dates, values, and status remarks.	117
14	The siren discriminator, three ways.	128
15	Six classes of Hubble-tension fix versus GG-Theory.	129
16	Status of each class after the 2026 landscape: the one-line verdict the audits.	130
17	Kill criteria for the GG-Theory reading of the H0DN measurement, in four tiers.	138
18	The derived expansion history at the redshifts the curve-reading channels probe.	143
19	GG-Theory-inherited cosmological benchmarks and their observational status as of 2026. . . .	147
20	Events to verdict (Eq. (51); open-source script <code>c2_siren_forecast.py</code>).	152
21	The systematic budget of the referee.	153
22	The open-question program of the derived Hubble reading.	171
23	The paper in ten numbers: value, grade, location, and falsifier.	178
24	The course map of these appendices.	188
25	Every headline number of the derived Hubble reading, its grade, and where it is derived. . . .	190
26	Inputs to the derived background rate, their provenance, and their contribution.	231
27	The open-source scripts released with this paper and what each reproduces.	232
28	Full list of H0DN baseline and variant values from Casertano et al. (2026).	233
29	The internal source ledger: each headline object’s home in the main text, and the appendix. .	233
30	What the derived Hubble reading inherits from the cohort.	234
31	The named theorems and conjectures the derived Hubble reading depends on.	235
32	What would promote each conditional or asserted entry of the derived Hubble reading.	236

Contents

Executive Summary	2
List of Figures	9
List of Tables	9
Preface	19
 Part I. Introduction — The expanding universe, the Hubble parameter $H(t)$, and a geometric answer to a century-old measurement	20
1 The expanding universe: a short history	20
2 The Hubble parameter, the Hubble constant, and the tension	21
3 The GG-Theory approach — a preview	22
4 Three reading paths	24
5 Summary, and the road to Part II	25
6 The cohort inheritance and the roadmap of the twelve Parts	27
6.1 The cohort inheritance	27
6.2 Roadmap of the twelve Parts	32
7 Part I in plain English: the two numbers, and why one of them stopped being a number	33
 Part II. The Hubble Tension — A century of measuring H_0; the dark-energy surprise; the two definitions; the 7σ H0DN consensus; and why the tension is hard to explain	35
8 Defining and measuring H_0 : a century of observation	35
9 The surprise: dark energy and accelerating expansion	36
10 Anatomy of the two routes	38
10.1 The late route: climbing the ladder	38
10.2 The early route: computing the rate from the oldest light	39
10.3 The comparison, and the three possible homes	40
11 Why this measurement matters	41
12 The central claim: the Hubble rate as an open-system observable	41
13 Architecture of the Local Distance Network	42
13.1 Definitions	42
13.2 The baseline pre-registration	43
14 The baseline result and its structural features	43
15 The early-universe network: Planck, ACT, SPT, and DESI	44
16 The statistics of 7.1σ : what the number does and does not say	45
17 What the H0DN paper does <i>not</i> do	46
18 Summary: the target for any theory, and the road to Part III	47
19 Part II in plain English: the problem as any cosmologist would state it, before any geometry is spent	48
 Part III. The GG-Theory of the Hubble Parameter — The Hubble parameter as a GDOS/GRIP observable; the direct derivation of $H(t)$ and H_0; and why it reproduces the CMB value	50
20 The GG-Theory framework in one view	50
20.1 What a theory of geometry may output: the GU-5 dimensional bridge	52
20.2 The recast cosmology GG-A10-Cosmos	53
21 The Hubble parameter as a GDOS/GRIP observable	54
21.1 The Friedmann equations as a theorem	54

22	The π -locked inventory: derived, not assumed	56
23	From the law to the curve: $H(t)$ in six steps	58
24	Geometrical determination of H_0 : geometry first, the sky second	62
24.1	Route I — the rate from the geometry alone: the vacuum-energy seesaw	62
24.2	The selection of ζ_Λ : by the codebook, not by the data	65
24.3	The closed form: the expansion rate from $\{c, G_N, \hbar\}$ and the ledger alone	67
24.4	Route II — the precision confirmation: the microwave sky reads the derived rate off two densities	69
24.5	The concordance, and the canonical rate	70
24.6	Provenance: what is derived, what is conditional, what is input	72
25	Landmarks on the derived $H(t)$ curve	72
26	The microwave background as an over-determined test	77
26.1	The h -free ratio: the baryon-loading test	78
26.2	The core observables, computed and confronted	78
26.3	The full spectrum and the χ^2	78
26.4	The coherent π -offset: one systematic, visible everywhere	80
27	Summary: the derived history, and the road to Part IV	81
28	Part III in plain English: how a handful of integers becomes an expansion history with nothing left to adjust	84

Part IV. Reconciling the Hubble Tension with GG-Theory — The origin of the larger local value; the quantitative estimate; the limit of backreaction as a remedy; and a call to revisit Λ CDM in the light of the derived $H(t)$ **86**

29	The task of this Part, and the strategy	86
30	The Hubble tension as a question of which instruments read the derived rate	87
30.1	The early side needs no new physics	87
30.2	The late side: one predicted systematic	88
30.3	The age of the universe: the time-axis image of the rate	89
30.4	The model-independent referee: standard sirens	89
31	The recombination switch and the two dictionaries	89
32	A worked estimate of the local excess	91
33	The local void, priced and not carried: evidence, linear theory, and budget	92
33.1	The evidence: how deep, how wide, how contested	93
33.2	Linear theory: three steps to the response	93
33.3	What the ladder actually averages, and what the corrections remove	94
33.4	The scenario ledger and the budget	95
34	The systematic against the error budget: where can ten percent hide?	97
34.1	The requirement, stated exactly	97
34.2	What the 1.09% contains, and what it structurally cannot	97
34.3	The two medium channels: dust, and the statistics of the ladder itself	98
34.4	The remaining channels, one by one	100
34.5	The verdict of the arithmetic, in both directions	101
35	The limit of backreaction as a remedy	104
36	The why-now window: the coincidence as a factor-of-a-few timing	106
37	Why the derived reading is not an ad-hoc fit	107
37.1	The strong- CP contrast	108
38	Cross-ledger coherence: why this is the hard test	108
39	The S_8 tension: shared, not resolved	109
39.1	The sign structure, as it now stands	109
40	Revisiting Λ CDM in the light of a derived $H(t)$	110
41	Summary: the reconciliation, and the road to Part V	110

42	Part IV in plain English: why the local ladder reads high, and what the framework does and does not claim about it	112
----	--	-----

Part V. The 2026 Measurement Landscape and the Rival Explanations — Every reported H_0 of the audited reference base, drawn and tabulated; the three bands; the six rival classes read against them; and the one measurement that bisects the field **113**

43	The measurement landscape in 2026	113
43.1	The three bands, and how to read them	113
43.2	The consensus that is not a consensus	119
43.3	The early side is closed	120
43.4	The high geometric counterweight	120
44	Statement of the comparison problem	121
45	Class 1 — Λ CDM plus measurement systematics: a named hypothesis	122
46	Class 2 — phenomenological early dark energy	123
47	Class 3 — modified recombination	124
48	Class 4 — interacting dark sectors and extra radiation	125
49	Class 5 — modified gravity and evolving dark energy	126
50	Class 6 — local voids and inhomogeneity	126
51	The siren discriminator: one measurement, three worlds	127
52	Comparison summary	128
53	Why “uniqueness” means more than “agreement”	130
54	Summary: the landscape, the rivals, and the road to Part VI	131
55	Part V in plain English: all forty-five measurements on one chart, and the six ways people have tried to close the gap	132

Part VI. Kill Switches, the Decision Calendar, and Wider Implications — The eleven named falsifiers with their pre-declared rules; the dated decision calendar 2026–2035 and what the paper commits to at each date; and what the H0DN measurement means for the broader GG-Theory program **133**

56	The falsifier registry: F-Hubble-1 through F-Hubble-11	133
56.1	Program-level kill tests	133
56.2	Parameter-block-level kill tests	134
56.3	Channel-level kill tests: the residual anatomy	136
56.4	The coefficient-level test	137
56.5	Kill-criteria summary	137
57	The decisive observation: how the tension will be settled	137
58	The decision calendar, 2026–2035	139
58.1	The dated entries	139
58.2	What the paper commits to at each date	141
59	Over-determining the curve: chronometers, calibrated rulers, and the redshift drift	142
60	Observational priorities opened by the H0DN result	143
60.1	Priority 1: Joint CMB+BAO+ H_0 analysis within the inherited ledger	144
60.2	Priority 2: The injection–recovery campaign (F-Hubble-9)	144
60.3	Priority 3: The near-infrared ladder and the sibling samples (F-Hubble-8)	144
60.4	Priority 4: Increasing the number of geometric anchors	144
60.5	Priority 5: Megamaser distances beyond NGC 4258	144
60.6	Priority 6: JWST extension of the primary ladders	145
60.7	Priority 7: DESI $w(z)$ constraints	145
60.8	Priority 8: Kaxion haloscope search in the window	145
61	Wider implications: S_8 , early galaxies, and the H0DN program	145

61.1	The S_8 tension	145
61.2	JWST-era early galaxies	146
61.3	The meaning of the agreement for the H0DN program itself	146
62	Cosmological benchmark sheet	146
63	Part VI in plain English: the eleven ways this reading can be killed, and the dates they are scheduled for	148

Part VII. Forecasting the Verdict: The Standard-Siren Decade — The physics of the siren; the referee’s report card 2017–2026; the events-to-verdict forecast on empirical rates; the referee’s own systematic budget; and the three worlds the catalog can choose among **149**

64	The principle: a distance from the waveform alone	149
65	Three pipelines, disjoint by construction	149
66	The referee’s report card, 2017–2026	150
67	The forecast: events to verdict	151
68	The referee’s own budget: can the siren itself be wrong by ten percent?	152
69	The verdict, simulated — and what failure looks like	154
70	The sharpened stake: three outcomes, three worlds	155
71	Reproducibility and the pre-registered rule	156
72	Part VII in plain English: the one instrument that owes nothing to either side, and when it will have spoken	156

Part VIII. Discussion: Significance, Uniqueness, and Innovation — What a derived expansion history means as a scientific claim; the problems it resolves and the moves it registers; the honest boundary of what it computes; and its open questions, predictions, and the road ahead **157**

73	What this Part does	157
74	The essence: one chain, one observable	157
75	Is H0DN a good support for GG-Theory?	158
75.1	What is strongly supported	158
75.2	What is moderately supported	159
75.3	What is not yet tested	159
76	Significance and uniqueness, in ten points	160
77	Comparison with the leading responses	162
78	Historical precedents — including the cautionary ones	163
79	The problems resolved: PS-Hubble-1 through PS-Hubble-7	164
80	The structural innovations: I-Hubble-1 through I-Hubble-10	166
81	The honest boundary: proved, conditional, open	167
82	The tier ladder: what this paper computes, and what it does not	168
83	The cost of being right: zero parameters, maximal exposure	169
84	The cumulative evidence base	170
85	Is the GG-Theory reading necessary?	170
86	Open questions, stated as a program	171
87	Predictions and future directions	173
88	The nobility clause	175
89	The bottom line	175
90	Part VIII in plain English: what the result is worth, what it is not, and where the debts are	176

Part IX. Conclusions — The retirement of a measured constant — one derived

history, a dozen ways the universe can answer, and what it means that we can ask	177
91 The summit thesis	177
92 What this paper has established	177
93 The retirement of a measured constant	178
94 One history, three acts, all predicted	179
95 The dozen ways the universe can answer	180
96 What it would mean if it is right	182
97 What it would mean if it is wrong	182
98 The objective limitations, one last time	183
99 How to check this paper from your desk	183
100 Why there is someone here to time the universe	184
101 A personal note, and a message to my colleagues	184
102 Part IX in plain English: what changes if a measured constant turns out to have been derivable	185
Envoi	186
Acknowledgments	187

Part X. Technical Appendices — A graduate course in the Hubble tension read through GG-Theory — reference dictionaries, a conceptual on-ramp, and from-scratch derivation courses that review each topic at textbook level, name what the standard treatment cannot explain, and derive the geometric resolution, with every step shown **188**

A	Reference dictionaries	189
A.1	The cohort naming convention	189
A.2	Comprehensive summary table: every headline number and where it is earned	189
A.3	Abbreviations and symbol glossary	191
A.4	The honesty grades: a reader’s key	192
B	The conceptual on-ramp	192
B.1	Five primers: the objects of this paper, in plain words	192
B.1.1	The expansion rate $H(t)$	192
B.1.2	What “distance” means in an expanding universe	193
B.1.3	The sound horizon	193
B.1.4	The standard siren	193
B.1.5	The Local Hole	193
B.2	Foundational questions about the Hubble tension	194
B.2.1	What is the Hubble tension, physically?	194
B.2.2	Why is the background rate 66.8 and not 73?	194
B.2.3	If it is one rate, why do early and late disagree?	194
B.2.4	What is ζ_A , and why is the seesaw “soft”?	194
B.2.5	Why does a standard siren settle it when neither ladder nor CMB can?	194
B.2.6	What would kill the derived reading?	195
B.2.7	What exactly is a standard siren, and why is it immune to both routes?	195
B.2.8	What is the Local Hole, and how sure are we it exists?	195
B.2.9	Why can’t a void, or averaging effects, explain the whole tension?	195
B.2.10	The acoustic angle is measured to four digits and the chart misses it by half a percent. How is that a success?	195
B.2.11	If Ω_m is exactly $1/\pi$, why do fits return 0.315?	195
B.2.12	What happens to the universe, on the derived curve?	196
B.2.13	What exactly does “Asserted-external” commit the paper to?	196

	B.2.14	Why print two combination rules — quadrature and stack? . . .	196
	B.2.15	Why is the recombination switch at e^7 , of all numbers?	196
	B.2.16	How can $w = -1$ be <i>exact</i> when every fitted value has error bars? 196	
C		Derivation course: the Friedmann background from the metric up	197
	C.1	The FRW metric and its conventions	197
	C.2	Christoffel symbols: every non-vanishing component	197
	C.3	The Ricci tensor and the Ricci scalar	198
	C.4	The Einstein tensor	199
	C.5	The perfect-fluid source	199
	C.6	The two Friedmann equations, the acceleration equation, and continuity .	199
	C.7	Curvature, the critical density, and the Ω bookkeeping	200
	C.8	Worked numerics: the age, the acceleration epoch, and the $H(z)$ rows . .	201
	C.9	What the standard derivation assumes — and cannot explain	201
	C.10	The geometric resolution: the Friedmann system as a theorem	202
D		Derivation course: distances, the ladder, and the century of H_0	203
	D.1	The five distances of an expanding universe	204
	D.2	The ladder, rung by rung, with its error budget	204
	D.3	Worked example: the common-mode algebra	206
	D.4	What the ladder cannot decide from inside	206
	D.5	The geometric reading of the same budget	206
E		Derivation course: the microwave background as an H_0 instrument	207
	E.1	The tightly coupled plasma and its sound speed	207
	E.2	The sound horizon, evaluated	208
	E.3	The acoustic angle and the peak series	209
	E.4	Why the microwave sky never measures h	209
	E.5	The audit: zero background freedom against the whole spectrum	209
	E.6	The sound-horizon lever, and its cap	210
	E.7	The geometric resolution: the rate from two densities, and the coherent offset	210
	E.7.1	Inputs	210
	E.7.2	The two anchors	210
	E.7.3	Combination and error	210
F		Derivation course: the vacuum energy of the universe and the seesaw	212
	F.1	The vacuum energy of quantum fields, computed honestly	212
	F.2	Seesaw structures: how physics makes small numbers	212
	F.3	What the standard mechanisms cannot do	213
	F.4	The geometric resolution: the seesaw, derived and closed	213
G		Derivation course: the local universe — peculiar velocities, the Hole, and the lift a void would give	214
	G.1	Linear peculiar velocities from the continuity equation	214
	G.2	The master formula: what an under-density does to a local rate	214
	G.3	The measured Hole, and the velocity audit	215
	G.4	What the environment cannot explain	215
	G.5	The geometric resolution: the response law, and the closed exotic channel .	216
	G.5.1	Setting	216
	G.5.2	Effective local Friedmann equation	216
	G.5.3	Magnitude estimate: the first-pass result and the sixty-fold gap .	217
	G.5.4	Why this is not a large local void	217
H		Derivation course: backreaction and averaging	218
	H.1	Buchert averaging, from the definition up	218
	H.2	The size of the terms on the derived inventory	218
	H.3	What averaging cannot buy	219
	H.4	The record: the bound at the exact target	219
I		Derivation course: the growth of structure and S_8	220

I.1	The perturbed universe: three equations	220
I.2	The growth factor, era by era	220
I.3	From modes to observables: σ_8 and S_8	221
I.4	What the standard framework cannot say about its own tension	221
I.5	The geometric reading: one ledger, both amplitudes	221
J	Derivation course: standard sirens from the quadrupole formula up	222
J.1	Gravitational waves in one page	223
J.2	The chirp: one mass from the phasing	223
J.3	The amplitude: an absolute distance	223
J.4	Worked example: GW170817, end to end	224
J.5	What limits the referee	224
J.6	The geometric stake: the forecast algebra, re-derived	224
K	Derivation course: time, age, and the curve-reading channels	225
K.1	Cosmic time, derived in closed form	225
K.2	Chronometers and the drift, from first principles	225
K.3	What a one-number cosmology cannot do	226
K.4	The geometric chronology, as commitments	226
K.5	Annex: the thermal clock — scaling laws, thermal history, and freeze-out	226
L	Derivation course: the statistics of a tension	227
L.1	What a sigma means when systematics exist	228
L.2	Why precision amplifies bias: the decade’s paradox, derived	228
L.3	The bracket theorem: combining channels with unknown correlations	228
L.4	What a covariance-weighted consensus certifies	229
L.5	Reading a landscape: bands, not rows	229
L.6	The geometric posture: statistics with the freedom removed	229
M	Reproducibility and the released pipeline	230
M.1	Reproducibility: inputs, justifications, and the systematic budget	230
M.2	Scripts released with this paper	231
M.3	Reproducing the decisive test: the standard-siren H_0	231
M.4	The H0DN variant ladder: the reproducibility table	232
N	Registries, audits, and roadmaps	232
N.1	Internal source ledger: where each object lives	232
N.2	The falsification catalog, compact	234
N.3	Cross-paper inheritance table	234
N.4	Theorem-and-conjecture deck	235
N.5	Promotion roadmap	235
N.6	Notation	235
	Author’s Note	237

Part XI. Common Objections and Responses — Twenty big questions for the curious reader, plainly answered; then the hardest questions a skeptical cosmologist would ask of the derived-rate reading, answered in the framework’s own terms **238**

O	For the curious reader: the big questions, plainly answered	238
P	On the derived rate	245
Q	On the local ladder and the residual	247
R	On the geometry and the π -fractions	248
S	On the comparison with other solutions	250
T	On falsifiability and scope	250
U	On the statistics of the tension	251
V	On the siren and the forecast	252
W	On the program and its manner	253

X	On the 2026 landscape and the residual anatomy	254
Y	The seven-persona hostile review	256
Z	The second exchange	257
Z.1	The third exchange: the audit of the amendments	259

Part XII. References — The curated bibliography, split into the twenty-two-paper GG-Theory cohort (Part A) and all other works (Part B), each entry carrying a Relevance note and its first in-text citation	262
--	------------

Preface

This Preface is addressed, more than to anyone else, to the working cosmologist who is tired of the Hubble tension — who has watched a decade of proposed resolutions each add a field, a fluid, or a knob, and wondered whether the argument can end except by exhaustion. It is also addressed to the student deciding whether the tension is worth a thesis (it is — though not for the usual reason).

Consider where the field stands. The tension is real: after the H0DN consensus — the H_0 Distance Network, the thirty-eight-author ladder combination — no single method, anchor, or instrument can be blamed, and the formal significance is beyond argument. Yet the response of theory has been, overwhelmingly, to negotiate — to add early dark energy here, extra radiation there, each proposal purchasing its ten percent with a parameter tuned after the fact, each retireable only slowly, by accumulating disfavor. Ten years of this has produced a literature in which everything is possible and nothing is decisive. The missing ingredient was never another mechanism; it was a *stake*: a number some theory needed to be true, printed before the measurement that could kill it.

This paper hands you one. The expansion history is derived — the whole curve, from the first second to the de Sitter floor, with the present value at 66.8 ± 0.2 and not one continuous dial anywhere — and everything that follows from it is priced in the open: the microwave sky it must fit with no background parameter tuned (it does, at the percent level, coherently), the landmarks it must pass from nucleosynthesis to the present age (it does, all within a percent), the local excess it must explain (it explains a third and declines, in print, to invent the rest), and the ladder’s error budget against which the remainder is confronted channel by channel. Where the paper cannot derive, it says so in a graded ledger; where it can be killed, it says how, with numerical rules fixed in advance.

What should you do with it on Monday morning? Three things, in ascending order of effort. Recompute the headline: the background rate is two published densities divided by two exact fractions — a calculator exercise, and either it lands in the anchors’ 66.4–67.0 span, with the geometry’s own 66.8 ± 0.2 inside it, or this paper has a very short future. Run the scripts: ten of them ship with the paper, one per figure and forecast, each asserting its own arithmetic before it draws; the fastest referee report in cosmology is `python c2_siren_forecast.py`. And if you work on the ladder or the sirens: the shell test (H_0 versus $z_{\min}$), the twenty-two-event forecast, and the channel table of Part IV were written to be sharpened by exactly your expertise — adversarially, please.

Two warnings, so the invitation is honest. This is not the community’s consensus framework: it is one deductive structure, published as a twenty-two-paper cohort, and the present paper leans on theorems and two named conjectures proved or stated elsewhere in it — the honesty triage and the objections Part tell you exactly which links carry which grade, and the seventy-nine objections there include every charge we could construct against ourselves, and the twenty big questions a curious reader asks, answered without softening. And the falsifiers cut both ways by design: a converged siren at 73, a firm inventory below $1/\pi$, a confirmed drifting equation of state, or a single proton decay each ends this reading outright, with no fallback prepared. We have written down the dates.

That is the wager. The tension has waited ten years not for another mechanism but for a prediction worth scheduling instruments around; what follows is that prediction — and whichever way the sirens rule, the argument will end the way physics arguments are supposed to end.

Part I

Introduction

The expanding universe, the Hubble parameter $H(t)$, and a geometric answer to a century-old measurement

1 The expanding universe: a short history

If you had asked an astronomer in 1900 what the universe was, the answer would have been serene: a single island of stars — our Milky Way — floating in an unchanging void, without a beginning, without an end, and without much of a history. The stars wheeled, comets came and went, but the universe as a whole simply *was*. Of all the revolutions the twentieth century inflicted on human self-understanding, none was larger, or stranger, than what happened to that serene answer over the following hundred years.

The first shock was size. Through the 1910s and 1920s astronomers argued over whether the faint spiral “nebulae” were wisps of gas inside our own galaxy or entire island universes of their own. Henrietta Leavitt had found the key without knowing the door it would open: a class of pulsing stars, the Cepheids, whose rhythm reveals their true brightness — and therefore, by how faint they appear, their distance (Leavitt and Pickering, 1912). When Edwin Hubble picked out Cepheids in the Andromeda nebula (Hubble, 1925), the answer was unarguable: Andromeda lay far beyond the Milky Way’s edge. Our island of stars was one among billions, and the universe was suddenly, dizzyingly, larger than anyone had imagined.

The second shock was motion. Vesto Slipher had noticed that the light of the spirals was almost always stretched toward the red — the optical signature of recession (Slipher, 1917) — and by 1929 Hubble, combining those velocities with his Cepheid distances, saw the pattern that now carries his name: the farther away a galaxy lies, the faster it recedes, in direct proportion, $v = H_0 d$ (Hubble, 1929) — a pattern soon confirmed to far greater depths (Hubble and Humason, 1931). Georges Lemaître had already derived that very law two years earlier from Einstein’s equations (Lemaître, 1927). The proper picture is not galaxies fleeing *through* space like shrapnel, but space itself stretching, the way dots on an inflating balloon all recede from one another — every observer sees the same recession, no one is at the center, and the speed grows with distance simply because there is more stretching space in between.

Einstein himself had been told as much by his own equations a decade before, and had refused to listen: general relativity, left alone, insists that a universe of matter cannot stand still — as Alexander Friedmann demonstrated by simply solving the equations honestly (Friedmann, 1922). Einstein had instead inserted an extra term — his cosmological constant — expressly to freeze the cosmos into the eternal stillness everyone then believed in (Einstein, 1917). The expansion’s

discovery turned that term into his famous regret, and taught the century its deepest cosmological lesson: the universe is not a stage; it is a *dynamical object*, with equations of motion of its own.

And a dynamical object has a past. Run the expansion backward and the universe grows denser and hotter without limit: the stretching that carries galaxies apart today, reversed, compresses everything that exists into a hot, bright beginning. What began as an almost disreputable extrapolation — Lemaître’s “primeval atom” — made a concrete prediction: the heat of that beginning should still be out there, stretched by the expansion into a faint microwave glow filling the whole sky, its temperature estimated at a few kelvin as early as 1948 (Alpher and Herman, 1948). In 1965 that glow was found, largely by accident, humming into a New Jersey radio horn (Penzias and Wilson, 1965), and identified for what it was next door (Dicke et al., 1965). With it, the Big Bang ceased to be speculation. The universe now had an age, a childhood photograph, and a biography.

The century had one more surprise to deliver — in 1998, the discovery that the stretching is now *accelerating*, driven by something no one had ordered (Perlmutter et al., 1999; Riess et al., 1998) (Part II tells that story) — but the arc was already complete. In a hundred years the universe went from an eternal, static room of stars to an evolving, expanding, datable object with a violent birth, a measurable middle age, and a calculable future. And every frame of that biography — how old, how large, how bright, how long the lights stay on — is governed by a single quantity: *how fast space stretches*. The story of that one number, and of the moment it split in two, is the subject of this paper.

2 The Hubble parameter, the Hubble constant, and the tension

Return to the balloon. Paint two dots on it a centimeter apart and inflate it so that every distance doubles each minute: the dots separate at one centimeter per minute. Two dots ten centimeters apart separate at ten centimeters per minute — same stretching, more rubber in between. That proportionality between distance and recession speed is exactly what Hubble measured, and its coefficient is the *expansion rate* of the universe. Astronomers quote it in practical units: a galaxy one megaparsec away (about 3.26 million light-years) recedes at roughly 70 kilometers per second; a galaxy twice as far, twice as fast.

To make this precise, cosmology tracks the balloon’s size with the *scale factor* $a(t)$: physical distances between freely-falling observers scale as $a(t)$, normalized so that $a(t_0) = 1$ today. The expansion rate is its fractional rate of growth,

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)}, \quad (1)$$

and this — the *Hubble parameter* — is a *function of time*, not a constant: it was enormous in the hot early universe and has been falling ever since, as the contents of the universe dilute. What astronomers call the *Hubble constant*, H_0 , is simply the value this falling curve takes at the present moment, $H_0 \equiv H(t_0)$: one point on a curve, kept as a shorthand from the century in which only that one point could be measured. The number also sets the scales of everything: its inverse, $1/H_0 \approx 14.7$ billion years, is the natural age-scale of the universe, and its value fixes how large the observable

universe is and how the story ends.

Measuring that one number has been a hundred-year enterprise (Freedman and Madore, 2010), and Figure 4 compresses the whole of it into one chart. Read the top panel first: the measured value *fell by a factor of seven* as calibration errors were found and fixed (Baade, 1956; Sandage, 1958), then stalled for thirty years in a factor-of-two standoff (Sandage and Tammann, 1975), and was finally settled — it seemed — at ten-percent precision at the turn of the millennium (Freedman et al., 2001). Then read the bottom panel: as precision sharpened below three percent, the two great ways of measuring — climbing the ladder of stellar distances outward (Riess et al., 2022), and reading the Big Bang’s afterglow backward (Camphuis et al., 2025) — did not converge on one value. They split, hardened, and now sit seven standard deviations apart: the *Hubble tension*, the sharpest crack in modern cosmology.

A tension between two measurements of “the same” number is a puzzle only if the number is truly a single, well-defined constant. The viewpoint developed in this paper is that it is not: H_0 is one reading of the derived function $H(t)$, and the two routes measure that function through two different operational windows. Whether the disagreement is new physics, a measurement systematic, or a mis-specification of what is being compared is the question the rest of the paper takes up.

3 The GG-Theory approach — a preview

Before this paper can say anything about the tension, the reader deserves a plain account of the theory it speaks for — assuming no prior acquaintance whatsoever.

Start from a deficit that modern physics rarely advertises. Our two standard models — of particles and of cosmology — describe the world with magnificent precision, but they do not *explain* their own settings. The Standard Model of particle physics runs on roughly twenty-eight measured numbers (masses, coupling strengths, mixing angles); the standard cosmology adds six more (the density fractions, the expansion rate among them). Nothing in either theory says why a single one of these dials points where it does. It is as if we owned a magnificent clock that keeps perfect time, with thirty-four hands whose positions must be looked up rather than understood.

GG-Theory — the framework of this twenty-two-paper series — is a wager that every one of those dials is set by geometry. It begins from a short constitution of axioms and postulates (Arisaka, 2026, A1-Constitution) and derives from it a six-dimensional stage (Arisaka, 2026, A2-GG6). The two extra dimensions are the theory’s signature idea, and they are emphatically *not* places one could travel to. One is a *scale*: the depth of magnification, the coordinate along which physics changes as one zooms from everyday lengths toward the smallest distances. The other is an *angle*: a compact phase, a dial that runs from 0 to 2π and closes on itself. A scale and an angle — a zoom and a dial — are pure numbers rather than directions, and on that quiet reinterpretation the whole construction stands.




Figure 4: A century of the Hubble constant. *Top:* the measured value fell by a factor of seven — Lemaître and Hubble near 500–575, Baade’s population split (1952), Sandage’s HII-region correction (1958) — then stalled in the three-decade factor-two war between the Sandage–Tammann (~ 50 – 55) and de Vaucouleurs (~ 90 – 100) camps, ended by the HST Key Project at 72 ± 8 (Freedman et al., 2001). *Bottom:* the precision era. The distance-ladder track (red) and the CMB+ Λ CDM track (orange) sharpened from ten-percent to percent precision and, instead of converging, separated into today’s 7.1σ split between H0DN’s 73.50 ± 0.81 and the combined early-universe 67.24 ± 0.35 . The green dashed line is this paper’s derived background rate 66.8 ± 0.2 (Part III) — printed here so the reader can see, before any argument, where the geometry lands on the century’s chart.

The stage comes with a discipline: *maximum symmetry, minimum freedom*. Quantum consistency on the six-dimensional stage is demanding — most candidate laws simply destroy themselves — and requiring consistency alone fixes the particle content and a ledger of five whole numbers, $(1, 3, 5; 7, 35)$, with not one continuously adjustable dial left anywhere (Arisaka, 2026, A3-BI6). Those five are not chosen but *forced*, each with a fixed meaning — $M = 1$ a single tensor seam, $N = 3$ the color triplicity, $L = 5$ the $SU(5)$ block, $J = 7$ the odd closure index, and $K = 35 = LJ$ the throat depth whose warp factor e^{-K} fixes, among much else, the present value of H_0 . Everything we observe is then read as the *shadow* this stage casts on its four-dimensional boundary — our world. The companion papers work the consequences out sector by sector: the twenty-eight Standard-Model numbers emerge as expressions in the five integers and just two transcendentals — π , the angle of the compact gauge dimension, and e , the base of the exponential scale dimension (Arisaka, 2026, A9-SM); quantum measurement acquires a geometric reading (Arisaka, 2026, A8-QM); and the standard cosmology itself — the subject of the parent paper GG-A10-Cosmos — follows with its density fractions locked to π (Arisaka, 2026, A10-Cosmos).

That last result is the one this paper spends. If the cosmic recipe — how much matter, how much dark energy, how much ordinary atoms — is fixed by geometry, then Einstein’s equations leave nothing further to choose: the *entire expansion history* $H(t)$, from the first second to the endless future, is an output. The Hubble “constant” stops being something one measures and becomes something one *checks*.

Two facts organize everything that follows. *First*, because the densities are derived, the Hubble parameter is too: $H(t)$ is fixed algebraically by the energy content through the Friedmann law, and the present value comes out at $H_0^{(0)} = 66.8 \pm 0.2 \text{ km s}^{-1} \text{ Mpc}^{-1}$ — which lands, to about a percent, on the CMB value. The early side of the tension therefore needs no new physics; it is where the derived rate is read most cleanly. *Second*, and just as honestly, GG-Theory does *not* manufacture the local 73: it computes the two channels it once offered for part of the excess — a local-void lift and a coarse-graining term — and finds, as this paper will show quantitatively, that the first needs an under-density the surveys do not show and the second is three orders too small, leaving the whole excess as a distance-ladder calibration question the framework predicts and the ladder-free gravitational-wave standard siren will settle. The result is not a triumphant closure of the tension but something more useful: a derived expansion history, an honest account of where the local value sits, and a sharp, dated test that will decide the matter.

4 Three reading paths

This is a long report, and not every reader needs all of it — but every kind of reader deserves an itinerary built for them.

The cosmologist in a hurry. You want the claim, the evidence, and the kill switch, and you are entitled to skip everything else. Read Part II for the measurement — especially the statistics of the 7.1σ (§16), which says exactly what the headline number does and does not establish. Then the two hearts: the derivation of the rate (§24) with its over-determined CMB test (§26), and the

reconciliation of the local value — the worked estimate (§32), the void channel priced and not carried (§33), and the systematic priced against your own error budget (§34). Close with the falsifier registry and its decision calendar (§56–§58) and the forecast of the verdict’s arrival (Part VII). If you have one evening, that is the paper; if you have ten minutes, run `python c2_siren_forecast.py` and read the ten-numbers table of the Conclusions.

The program reader. If you arrive from the cohort, you want to know what is new here and what is inherited. Inherited, at its home grades: the inventory, the two-route rate, and the spectrum test (all GG-A10-Cosmos); the seesaw normalization (GG-A5-GDOS); the ledger (GG-A4-Euclid). New in this paper: the theorem-grade statement of the inventory with its Hopf-fibration reading (Theorem 22.1), the landmarks walk (§25), the recombination-switch proposition, the backreaction bound at the exact $1/\pi$ inventory (§35), the void-channel scenario ledger, the systematic-versus-budget confrontation, the siren-forecast Part with its released script, and the derivation courses of the Technical Appendices, which work every headline number to reproduction grade. The registries (App. N.3–N.5) map the two-way flow between this paper and its parent; the naming convention (App. A.1) lets you cite any object from outside.

The skeptic. Start where we are weakest, not where we are strongest. The honesty triage (Table 5) declares what is derived, what is conditional on the one named lemma, and what is open; the percent-versus-sigma objection (O-Hubble-32) is, in our own judgment, the sharpest charge in the paper, and its answer states plainly the standard by which a zero-parameter chart should and should not be judged. Then read the landscape and comparison of Part V — including the admission that for the ten percent this reading is a disciplined member of the unknown-systematics class — and the seventy-nine objections of Part XI, which close with a seven-member hostile panel that returns for a second round against the answers themselves. The nobility clause (§88) states the commitments you may hold this paper to, and the falsifiers are not rhetoric: several can be executed from your desk, starting with the released scripts (App. M.2).

The graduate student. The Technical Appendices were written for you, and they are self-sufficient: start at the welcome guide and the primers (App. B.1), take the Friedmann course (App. C) with a pen, then the ladder and CMB courses, and only then read Part III — it will read like a summary of things you have already derived. The career-honest sentence you deserve is printed in the personal note of Part IX: the *appointments* of this paper are mainstream observations that pay professionally whichever way the framework fares; the framework itself is a proposal under test. Work the first; watch the second.

5 Summary, and the road to Part II

Part I has told the century’s story in miniature, because everything in this paper stands on it. A hundred years ago the universe was an eternal, motionless room of stars; by the century’s end it was an expanding, datable object with a violent birth, a measured biography, and a calculable future. One quantity governs every frame of that biography — the rate at which space stretches —

and the twentieth century’s last gift to cosmology was the discovery that this rate is not a constant at all but a falling curve, $H(t)$, of which the famous “Hubble constant” is merely today’s reading.

Then the reading split. Measured up the ladder of stellar distances, the present rate comes out near 73.5; inferred back from the Big Bang’s afterglow, near 67. The century chart of Figure 4 shows both the humility this subject demands — calibration has moved this number by a factor of seven before — and the seriousness of the present impasse: the split emerged precisely when the errors became too small to hide it, and no single method, team, or instrument can any longer be blamed.

This paper’s stance, previewed in §3, is that the impasse is a diagnosis, not a duel. In GG-Theory the cosmic recipe is fixed by geometry, the whole curve $H(t)$ follows from it with no dial to turn, and its present value lands at 66.8 ± 0.2 — on the afterglow’s side. What remains is to explain the ladder honestly: no derived channel carries any of the excess; it is named as a calibration systematic the framework predicts and priced rather than explained away, and a ladder-free referee — the gravitational-wave standard siren — is scheduled to settle the matter this decade. The roadmap and the cohort map of the closing Section (§6) show where each step of that argument lives, and the three reading paths above offer each kind of reader a route of their own.

Part II now sets GG-Theory aside entirely. Before a single geometric claim is spent, the problem deserves to be stated at full strength: how the Hubble constant is actually measured, what the thirty-eight-author H0DN consensus did that no previous measurement could do, why its 7.1σ really is unimpeachable arithmetic — and what, precisely, that arithmetic does and does not say. It is the tension as any cosmologist would pose it; only afterwards does the derivation begin.

6 The cohort inheritance and the roadmap of the twelve Parts

Before the problem is stated at full strength (Part II), it is useful to lay out two reader-orientation maps in one closing Section: a panorama of the twenty-two-paper cohort in which this paper stands (§6.1, Figure 5), and a panorama of the twelve Parts of the present report (§6.2, Figure 6). 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 twelve-Part arc.

6.1 The cohort inheritance

This report is one paper of twenty-two, and it leans on its siblings at every load-bearing joint: the inventory it divides into the CMB densities is a theorem of the cosmology readout GG-A10-Cosmos; the integer ledger behind the recombination switch belongs to GG-A4-Euclid; the boundary law that makes the two measured values registered observables of one projection is GG-A5-GDOS's. The reader therefore deserves the whole map before the thread begins. The present paper is GG-C2-Hubble, the middle member of the Cosmology & Astrophysics branch — the dedicated Hubble companion of GG-A10-Cosmos, from which it inherits the derived expansion history and to which it returns the H0DN reading.

Figure 5, which closes this subsection, places the present paper among the twenty-two papers of the GG-Theory cohort. The reader meeting the program for the first time can read its architecture off the figure top to bottom — the constitutional root, the bulk trio, the open-system boundary law and its pillars, the three textbook readouts, the capstone, the two phenomenology branches, and the Foundations Branch on the left flank. What follows walks that same descent in words, naming the papers a reader of *this* report will most want on the desk.

The descent begins with a constitution. GG-A1-Constitution ([Arisaka, 2026, A1-Constitution](#)) fixes four axioms and seven postulates — the rules of admissibility every later claim must obey — and the bulk trio then erects the stage: GG-A2-GG6 ([Arisaka, 2026, A2-GG6](#)) derives the six-dimensional geometry itself, GG-A3-BI6 ([Arisaka, 2026, A3-BI6](#)) proves the anomaly freedom that makes the construction quantum-mechanically legal, and GG-A4-Euclid ([Arisaka, 2026, A4-Euclid](#)) traces the lineage back to Euclid and promotes the five-integer ledger (1, 3, 5; 7, 35) to theorem grade. GG-A5-GDOS ([Arisaka, 2026, A5-GDOS](#)) is the hinge of the whole program: the open-system law by which the six-dimensional bulk projects to the four-dimensional boundary we inhabit, organized from its geometric side by GG-A6-GRIP ([Arisaka, 2026, A6-GRIP](#)) and from its quantum-informational side by GG-A7-Qi4 ([Arisaka, 2026, A7-Qi4](#)).

Then come the three readouts, where the machinery pays out in textbook physics — and two of them are this paper's essential companions. GG-A9-SM ([Arisaka, 2026, A9-SM](#)) derives the *Standard Model*: all twenty-eight of its parameters as closed expressions in the five integers, π , and e . A reader who wants to see the same ledger spent far from cosmology — and to judge the program's credibility on particle-physics ground — should read GG-A9-SM alongside this paper. GG-A10-Cosmos ([Arisaka, 2026, A10-Cosmos](#)) does the same for *cosmology*, deriving Λ CDM's inventory and with it the expansion history. GG-A10-Cosmos is this paper's *parent*: the source of

record for everything Part III inherits — the π -locked fractions, the two-route rate, the spectrum test — and the single most valuable companion to the present report. Readers who finish this paper persuaded, or convinced it must be wrong, are warmly encouraged to read GG-A10-Cosmos next; the case for and against both papers is ultimately one case. (The third readout, GG-A8-QM ([Arisaka, 2026, A8-QM](#)), recovers quantum measurement; GG-A11-Origin ([Arisaka, 2026, A11-Origin](#)) closes the A-series as its capstone.)

A word on the Foundations Branch, the column on the map’s left flank, whose subject — the boundary’s complex structure and its mirrors — sits beneath the derived history rather than beside it; this paper’s relations to its two members are light, and they are named here at roster level. GG-P4-Real ([Arisaka, 2026, P4-Real](#)) is the origin-of- i member: from the two real structures of the forced arena — the compact phase circle and the causal clock — it derives the single imaginary unit of boundary physics, the complex structure that every phase and amplitude in this report’s cosmology quietly presupposes. GG-P5-CPT ([Arisaka, 2026, P5-CPT](#)) is the mirrors-and-matter member: C, P, and T as orientation reversals of the three real bulk factors, closing with why the universe is made of matter — and its four-endowment definition of time (time orders, repeats, opens one way, and shelters the quantum world’s mathematics) is kin to the chronology of the derived $H(t)$ read here, one clock beneath every reading. Neither is load-bearing for the derivation; they are named so the panorama is complete.

Table 1: The twenty-two-paper GG-Theory cohort. Each paper’s main achievement and its role in the program. Every paper name links to that paper’s own page at preprints.arisaka-gg.org, where all twenty-two are collected. The paper marked (*this paper*) is the one you are reading.

Paper	Main achievement	Role in the cohort
GG-A1-Constitution	Four axioms and seven postulates fixing what counts as an admissible claim — the constitution of the whole program	The constitutional root; every later paper inherits its four axioms and seven postulates
GG-A2-GG6	Derivation of the six-dimensional GG-6 bulk, BI-6, OGN-5, compact holonomy, and the unification of all six forces	First of the GG-6 bulk trio; fixes the geometry and ledger the boundary law is built on
GG-A3-BI6	Complete proof of BI-6 anomaly freedom, unified gauge consistency, and conservation of the dual current	GG-6 bulk trio; the ultraviolet legality certificate that lets the projected ledger exist
GG-A4-Euclid	The historical lineage from Euclid to GG-6, and the zero-parameter integer OGN-5 ledger behind the constants	GG-6 bulk trio (historical companion); supplies the integer ledger behind the constants
GG-A5-GDOS	OGGEP descent to boundary quantum mechanics, boundary general relativity, and formal 4D QFT: the GDOS boundary law	The open-system boundary law itself — the hinge where the 6D bulk projects to 4D physics

continued on next page

Table 1, continued.

Paper	Main achievement	Role in the cohort
GG-A6-GRIP	The runtime constitution — GRIP, the force law $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$, and the scale-invariance principle SIP	A pillar paper organizing the boundary law from the geometric (Gi-4) side
GG-A7-Qi4	The Qi-4 pillars — SeOS, QCI, the identity $\text{IH} + \text{OH} = \text{EIC}$ — and the five-story informational pagoda	A pillar paper organizing the boundary law from the quantum-information (Qi-4) side
GG-A8-QM	Quantum mechanics derived from GG-Theory — the four-dimensional reformulation of measurement and state	Consequence trio; the member that recovers textbook quantum mechanics on the boundary
GG-A9-SM	The Standard Model as a theorem — the descent of all twenty-eight Standard-Model parameters	Consequence trio; the member that recovers the Standard Model of particle physics
GG-A10-Cosmos	Λ CDM as the cosmological-boundary projection — the π -locked inventory $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$	Consequence trio; the member that recovers Λ CDM cosmology — <i>the parent of this paper</i>
GG-A11-Origin	The closing capstone — the universal identity and the GRIP cascade across cosmos, life, and mind	The capstone that closes the A-series and proves the program’s universal identity
GG-P1-KK	The Kaluza–Klein lock cluster near $M_{\text{lock}} \approx 3.85 \text{ TeV}$, with explicit LHC and future-collider falsifiers	Predictions branch; the collider phenomenology of the Kaluza–Klein lock scale
GG-P2-N3	A three-leg derivation of $N_g = 3$ fermion generations from BI-6, MSMF, and the integer ledger	Predictions branch; the three-leg derivation of the fermion generation count
GG-P3-Electron	The electron as the ground state of nature — the six oldest electron questions, closed	Predictions branch; the electron ground-state and charged-lepton-stratum result
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; the origin-of- <i>i</i> member — the derived complex structure that every phase and amplitude in this report presupposes
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 mirrors-and-matter member, whose four-endowment definition of time is kin to the chronology read here
GG-C1-Kaxion	Cold dark matter as the non-spatial compact-angle pseudoscalar Kaxion at $8.1 \mu\text{eV}$	Cosmology & astrophysics branch; the identification of cosmological dark matter

continued on next page

Table 1, continued.

Paper	Main achievement	Role in the cohort
GG-C2-Hubble (<i>this paper</i>)	The Hubble tension read against one derived expansion history: the background rate from geometry, the ladder as the predicted outlier	Cosmology & astrophysics branch; the resolution of the Hubble tension
GG-C3-BH	Black-hole entropy $A/4$ recovered as a boundary theorem within GG-Theory's holographic ledger	Cosmology & astrophysics branch; horizon thermodynamics and black-hole entropy

Finally, two phenomenology branches and one foundations column carry the program outward. The Predictions Branch stakes the collider and structural bets: the TeV lock cluster of [GG-P1-KK](#) ([Arisaka, 2026, P1-KK](#)), the three-generations theorem of [GG-P2-N3](#) ([Arisaka, 2026, P2-N3](#)), and the electron ground-state result of [GG-P3-Electron](#) ([Arisaka, 2026, P3-Electron](#)). The Cosmology & Astrophysics Branch is this paper's home: [GG-C1-Kaxion](#) ([Arisaka, 2026, C1-Kaxion](#)) identifies the dark matter and names the exact frequency to listen for it at; the present paper reads the Hubble tension; and [GG-C3-BH](#) ([Arisaka, 2026, C3-BH](#)) recovers black-hole entropy from the same boundary law. And the Foundations Branch — [GG-P4-Real](#) ([Arisaka, 2026, P4-Real](#)) on the origin of the boundary's complex structure, [GG-P5-CPT](#) ([Arisaka, 2026, P5-CPT](#)) on nature's three mirrors and the matter excess — sits on the spine's left flank. The roster of Table 1 compresses all twenty-two into one line each.




Figure 5: 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 $N_g=3$ into GG-P3-Electron, and that paper’s ground-attractor method, GGA, on into GG-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 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.

6.2 Roadmap of the twelve Parts

The rest of the report is organized into eleven further Parts, and the walk through them is plotted in Figure 6. The present Part — the on-ramp — is nearly done: it has told the century’s story in miniature, introduced the falling curve $H(t)$ of which the famous constant is one reading, previewed the geometric answer, offered three reading paths, and drawn the two orientation maps of this closing Section.

Part II, *The Hubble Tension*, states the problem at full strength before a single geometric claim is spent: a century of measuring H_0 , the dark-energy surprise, the two definitions of the expansion rate, the 7σ H0DN consensus, and why the tension is genuinely hard to explain. Nothing in it depends on GG-Theory; it is the tension as any cosmologist would pose it.

Parts III and IV are the two hearts, and the roadmap highlights them. Part III, *The GG-Theory of the Hubble Parameter*, reads the expansion rate as a GDOS/GRIP observable and derives the whole history $H(t)$, with today’s value at 66.8 ± 0.2 from the geometry alone — and shows why the microwave background, reading the same geometry through its own sound-horizon calibration, reports the value it does. One rate, read by every ladder-free probe, is this paper’s central claim. Part IV, *Reconciling the Hubble Tension with GG-Theory*, then faces the larger local value honestly — the void channel and its bound, the quantitative estimate, the limit of backreaction as a remedy, and a call to revisit Λ CDM in the light of the derived $H(t)$.

Three Parts then take the claim outdoors. Part V, *The 2026 Measurement Landscape and the Rival Explanations*, draws every reported H_0 of the audited reference base — forty-five determinations, charted and tabulated — sorts the field into three bands, and reads the six rival classes against them, ending with the one measurement that bisects the field. Part VI, *Kill Switches, the Decision Calendar, and Wider Implications*, commits: the eleven named falsifiers with their pre-declared rules, on a dated decision calendar 2026–2035. Part VII, *Forecasting the Verdict*, is the standard-siren decade in full — the physics of the siren, the referee’s report card 2017–2026, and the events-to-verdict forecast on empirical rates.

Part VIII, the *Discussion*, weighs the reading as a scientific claim — significance, uniqueness, innovation, and the open questions stated as open — and Part IX, the *Conclusions*, closes the body with the retirement of a measured constant and a dozen dated ways the universe can answer. The back matter is the reference machine: Part X carries the *Technical Appendices* — a graduate course in the Hubble tension with every step shown; Part XI answers the *Common Objections and Responses*, closing with the hostile panel’s second round; and Part XII closes the paper with the annotated *References*, each entry carrying its Relevance note and its first citation. Read in order, the twelve Parts compose one self-contained argument; read by need, each Part names its imports and can stand a visit on its own.




Figure 6: Roadmap of the twelve Parts of GG-C2-Hubble. The argument develops in eight phases. *On-ramp* (Part I): the expanding universe, the two readings of one rate, and the maps of this closing Section. *The measurement* (Part II): a century of measuring H_0 and the 7σ H0DN consensus, stated at full strength before any geometry is spent. *The derivation* (Part III, highlighted): the π -locked inventory turned into the whole history $H(t)$, today's value $H_0^{(0)} = 66.8 \pm 0.2$, and the reason the CMB reads the same geometry as it does. *The reconciliation* (Part IV, highlighted): the origin of the larger local value, the quantitative estimate, and the limit of backreaction as a remedy. *The tests* (Parts V–VII): the 2026 landscape and its rivals, the eleven kill switches on a dated decision calendar, and the standard-siren forecast of the verdict. *Critique* (Part VIII) and *summit* (Part IX) close the body, and the three back-matter Parts collect the Technical Appendices (Part X), the Common Objections and Responses (Part XI), and the annotated References (Part XII). Parts III and IV are highlighted because together they carry the paper's central claim: one derived expansion history, read by the microwave background and predicted for every ladder-free late probe (Part III), with the local ladder's excess named as the calibration systematic the framework predicts and does not derive (Part IV).

7 Part I in plain English: the two numbers, and why one of them stopped being a number

The universe is getting bigger, and there is a number that says how fast. Write it as a fraction and it is easy to hold: every cosmic distance grows by about seven percent in a billion years. That fraction is the Hubble parameter. Because the growth rate has not been the same throughout cosmic

history, the parameter is a curve and not a constant, and what is called the Hubble constant is one reading off that curve — the value it takes today.

For most of the last century the job was to measure that reading, and the job was hard, and successive answers were far apart because the measurement was hard. That is an ordinary scientific story with an ordinary ending, in which the errors shrink and the answers converge.

Something else happened. Two ways of getting the number were each pushed to about one percent, and they did not converge. One climbs outward from nearby stars, rung by rung, and arrives at about 73.5. The other starts from the oldest light in the sky and computes what today's rate must be, and arrives at about 67.2. The gap is nine percent, which is several times larger than either error bar. Two careful measurements of the same quantity that disagree by that much are not a pair of measurements; they are a problem.

This paper's answer is not another measurement. The claim is that the whole curve — not the reading, the curve — follows from a geometry fixed elsewhere in this cohort for reasons that had nothing to do with cosmic expansion, and that the early-side number is what the curve predicts while the late-side number is what an observer standing where we stand would read off it. Everything that follows is that calculation and the tests it exposes itself to.

Part II

The Hubble Tension

A century of measuring H_0 ; the dark-energy surprise; the two definitions; the 7σ H0DN consensus; and why the tension is hard to explain

8 Defining and measuring H_0 : a century of observation

The two definitions of Part I are worth restating here, so that this Part can be read on its own. The expansion of the universe is described by one function of time — the *Hubble parameter* $H(t)$, the fractional rate at which cosmic distances stretch. The *Hubble constant* H_0 is not a second object but a single reading of the first: the value $H(t)$ takes today. It is the constant of proportionality in Hubble’s law, $v = H_0 d$, and its customary unit spells out the operation: a value of $70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ says that a galaxy one megaparsec away recedes at seventy kilometers per second — equivalently, that every cosmic distance grows by about seven percent per billion years — and its inverse sets the age scale of the universe, $1/H_0 \approx 14$ billion years. Everything in this Part is an attempt to measure this one number, and the tension of the title is the discovery that the two best attempts, each impeccable by its own standards, disagree about it.

Measuring H_0 has been a century-long program built on the *distance ladder*: geometric parallaxes calibrate the period–luminosity relation of Cepheid variables; Cepheids calibrate the peak luminosity of Type-Ia supernovae; and supernovae, visible across cosmological distances, carry the calibration into the smooth Hubble flow. Each rung inherits the uncertainty of the one below, so the enterprise is as much about controlling systematics as about collecting photons.

The cleanest way to tell the century is to walk Figure 4 (Part I, page 23) from left to right. The curve has five features — a start high off the modern chart, two great cliffs, a thirty-year plateau split into two camps, a brief convergence, and a fork at the right edge — and each one is a lesson in how this particular number is measured, and mismeasured.

The start (1927–1929): a factor of seven too high. The measurement began in confusion and error — instructive error. Vesto Slipher’s spectra had shown the spirals receding; Henrietta Leavitt’s period–luminosity law had made distance measurable; and by 1929 Hubble could plot velocity against distance and read off a slope of roughly $500 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Hubble, 1929). Georges Lemaître had derived the expansion law from general relativity two years earlier and extracted essentially the same constant from the same sparse data (Lemaître, 1927). The number was wrong by a factor of seven — not through any failure of the idea, but through calibration: the Cepheids used to gauge distance were a mixture of two stellar populations with different period–luminosity laws, and the “brightest stars” used beyond them were often unresolved HII regions. The very first value of the Hubble constant was, in modern language, a distance-ladder systematic of six hundred percent.

The first cliff (1952): two kinds of Cepheid. Walter Baade separated the two Cepheid populations, and every galaxy calibrated on the wrong one abruptly moved twice as far away: H_0 fell to ~ 250 (Baade, 1956).

The second cliff (1958): stars that were not stars. Allan Sandage identified the HII-region confusion — the “brightest stars” were glowing clouds of ionized gas, brighter than any star could be — and cut the value to ~ 75 (Sandage, 1958). Within three decades the number had fallen sevenfold, and not once because the physics changed: every drop was a calibration error found.

The plateau (1960s–1990s): the factor-two war. Then the fall stopped, and the disagreement began: for three decades Sandage and Tammann defended ~ 50 –55 (Sandage and Tammann, 1975) while de Vaucouleurs defended ~ 90 –100, each side publishing careful, internally consistent analyses that disagreed with the other by a factor the quoted errors could not explain. Every rung was contested — the Cepheid zero-point, the supernova standardization, the flow corrections — and the dispute could not be arbitrated from within, because both camps shared the same kinds of calibration and differed only in their judgment of them.

The convergence (2001): the Key Project. The Hubble Space Telescope was built in part to end that war, and it did: the Key Project’s final value, $72 \pm 8 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Freedman et al., 2001; reviewed in Freedman and Madore, 2010), split the difference at ten-percent precision and made the factor-two era history. For a few years the constant seemed settled — one number, one decimal in doubt.

The fork (2013–today): precision divides. Then precision did something unexpected. As both routes pushed below three percent, they did not converge on the settled value; they moved apart and hardened. The distance ladder, rebuilt around HST and then JWST photometry, climbed through 74.03 ± 1.42 (Riess et al., 2019) to 73.04 ± 1.04 (Riess et al., 2022) and now to the H0DN consensus 73.50 ± 0.81 ; the microwave background, read through Λ CDM, descended to 67.4 ± 0.5 and now 67.24 ± 0.35 (Camphuis et al., 2025). A curve that fell by a factor of seven, paused in a thirty-year stalemate, and converged — split in two precisely when the errors became too small to hide the difference.

The century teaches a double lesson, and this paper will need both halves. The first is humility: calibration systematics have repeatedly moved this particular number by amounts that dwarf today’s tension — by a factor of seven at the start, by a factor of two within living memory. The second is respect: the modern network was engineered, rung by redundant rung, precisely so that history could not repeat — no single method, anchor, or instrument can move the consensus by more than a fraction of its error bar (§12 and below). Whether the last nine percent is new physics or the century’s final systematic is exactly the question this report takes up, and neither half of the lesson may be dropped.

9 The surprise: dark energy and accelerating expansion

The fork at the right edge of the century figure was the second great shock of modern cosmology. The first arrived in 1998, and it is worth telling properly, because it changed what the Hubble “constant” *is*: it turned a number into a curve.

By the mid-1990s observational cosmology was, in Allan Sandage’s famous phrase, a search for two numbers (Sandage, 1970): H_0 , the present expansion rate, and q_0 , the *deceleration* parameter. That the second would be positive was hardly questioned — gravity attracts, so a universe made of matter must expand ever more slowly, and the only debate was whether the braking was strong enough to someday reverse the expansion. Two rival teams — the High- z Supernova Search (Schmidt et al., 1998) and the Supernova Cosmology Project — set out to measure that braking, using Type-Ia supernovae as standardized candles bright enough to be seen more than halfway across the observable universe.

Both teams found the opposite of what they were built to find. The High- z team’s sixteen distant supernovae were ten to fifteen percent *farther* away than even a low-density, gently decelerating universe could place them: the expansion is not slowing down but *accelerating*, with $q_0 < 0$ and a positive cosmological constant preferred at better than 3σ (Riess et al., 1998). The Supernova Cosmology Project’s independent sample of forty-two supernovae, analyzed with a separate pipeline, landed on the same conclusion, with a flat-universe fit of $\Omega_m \approx 0.28$ and the remainder in Λ (Perlmutter et al., 1999). Two competing teams, two data sets, one verdict — recognized by the 2011 Nobel Prize in Physics — and a new component, *dark energy*, carrying roughly seventy percent of the cosmic budget, observationally consistent with a cosmological constant of equation of state $w = -1$.

Within seven years the discovery hardened into a concordance, each confirmation arriving through a channel that owed the supernovae nothing. The Boomerang balloon flight located the first acoustic peak of the microwave background at $\ell \approx 197$, pinning the total density to critical — a flat universe — at the ten-percent level (de Bernardis et al., 2000). WMAP’s first-year data then delivered the six-parameter fit that made Λ CDM the standard model of cosmology: $H_0 = 72 \pm 5$, $\Omega_\Lambda \approx 0.73$, and a cosmic age of 13.7 billion years (Spergel et al., 2003). And in 2005 the baryon acoustic peak — the same sound horizon that patterns the microwave sky — was detected in the correlation function of some forty-six thousand SDSS luminous red galaxies, at the predicted $\sim 100 h^{-1}$ Mpc (Eisenstein et al., 2005). Flat, accelerating, one-third matter, two-thirds Λ : by the mid-2000s the budget was set, and two decades of subsequent data have held w at -1 to within a few percent (Brout et al., 2022; Planck Collaboration, 2020).

What the concordance did not do is explain itself. Λ CDM *fits* the acceleration; it does not say why the dark-energy density has the value it has — an energy scale of a few milli-electron-volts, absurdly small by particle-physics standards — nor why matter and dark energy, whose densities scale differently by factors of billions over cosmic history, happen to be comparable *now*. Those two questions are precisely where this paper’s framework enters: Part III derives the dark-energy fraction ($\Omega_\Lambda = 1 - 1/\pi$) and its milli-electron-volt scale from the same geometry that fixes H_0 , and Part IV prices the why-now coincidence.

The acceleration matters for this paper in a way that is easy to miss: it makes the expansion rate a genuinely *time-dependent* quantity whose late-time behavior is set by Λ . The universe decelerated for its first eight billion years, while matter dominated, and began accelerating at $z \approx 0.6$ — about six billion years ago — when the thinning matter density finally dropped below the constant dark-energy density. Speaking of “the Hubble constant” as though it were a fixed

number is therefore a convenient shorthand for a single reading of the function $H(t)$, taken at the one epoch we happen to occupy — and it is exactly at that shorthand that the two measurement routes part company. Part III makes the function itself — derived, not fitted — the central object.

10 Anatomy of the two routes

The tension is a disagreement between two measurement philosophies. One climbs outward from the Earth, rung by rung, and reads the present-day rate directly from the recession of nearby galaxies; the other starts from the oldest light in the universe and *computes* what the present rate must be, given a model of everything that happened in between. Both are precision instruments of the highest order, and neither is naive: each is the best possible version of its philosophy. Because everything later in this paper — and every proposed resolution in the literature — lives somewhere inside this anatomy, this section builds each route explicitly and then sets the two side by side (Figure 7 and Table 2).

10.1 The late route: climbing the ladder

The late route is *direct and upward*, and it climbs four rungs.

Rung 0 — geometry. At the base sit distances that involve no astrophysical assumption at all. Trigonometric parallaxes from Gaia and HST fix the distances of Milky-Way Cepheids (Lindgren et al., 2021); twenty detached eclipsing binaries fix the distance of the Large Magellanic Cloud to one percent (Pietrzyński et al., 2019); and the Keplerian rotation of water megamasers around the nucleus of NGC 4258 fixes that galaxy’s distance to one and a half percent (Reid et al., 2019). These three anchors are the ladder’s foundation: pure geometry, no candles.

Rung 1 — primary indicators. Within the anchors live stars whose luminosity is predictable: classical Cepheids, whose period–luminosity law has been the workhorse since Leavitt, and the tip of the red giant branch (TRGB), a sharp feature of stellar evolution. Calibrated in the anchors, they measure distances out to tens of megaparsecs — far enough to reach the few dozen galaxies that have hosted a well-observed Type-Ia supernova.

Rung 2 — supernovae. Type-Ia supernovae are standardizable candles: after correction for light-curve shape and color, their peak luminosities agree to better than ten percent, and they outshine their host galaxies. The roughly forty calibrator hosts transfer the geometric zero-point to the supernova scale (Scolnic et al., 2022).

Rung 3 — the Hubble flow. Hundreds of supernovae at $0.023 < z < 0.15$ — far enough that peculiar velocities of a few hundred km s^{-1} are small against expansion velocities of many thousands, near enough that the assumed cosmology barely matters — trace out the Hubble diagram, and its intercept is H_0 .

The structure of the errors follows from the structure of the climb. Each rung transfers a luminosity zero-point to the next, so independent fractional errors add in quadrature,

$$\left(\frac{\sigma_{H_0}}{H_0}\right)^2 \simeq \sigma_{\text{anchor}}^2 + \sigma_{P-L}^2 + \sigma_{\text{SN}}^2 + \sigma_{\text{flow}}^2, \quad (2)$$

and the H0DN’s 1.09% is the sum of dozens of sub-percent terms held down simultaneously — a remarkable feat of metrology. Two structural properties matter for everything that follows. First, the route is *cosmology-free*: nothing in the climb assumes an expansion history (the only cosmological inputs are the kinematic priors $q_0 = -0.55$ and $j_0 = +1$, and varying them moves the result by less than $0.1 \text{ km s}^{-1} \text{ Mpc}^{-1}$). Second, the route is *calibration-inheriting*: a bias at any rung propagates multiplicatively to every rung above it, so the surviving worry is not statistics but a calibration error shared *across* rungs — the one hypothesis the quoted error bar cannot price (§16).

10.2 The early route: computing the rate from the oldest light

The early route is *inferential and downward*, and it proceeds in three steps.

Step 1 — measure the pattern. The microwave sky delivers three exquisitely measured quantities. The angular scale of the acoustic peaks, $\theta_* \approx 0.010411$ radians, is known to three parts in ten thousand — the single most precisely measured quantity in cosmology; and the physical densities of baryons and of matter, ω_b and ω_m , follow from the relative heights of the peaks at the percent level (Planck Collaboration, 2020). None of the three is an expansion rate.

Step 2 — compute the ruler. Before recombination, the baryon–photon plasma carries sound waves; the distance such a wave travels by decoupling — the sound horizon

$$r_s = \int_{z_*}^{\infty} \frac{c_s(z) dz}{H(z)} \approx 144 \text{ Mpc} \quad (3)$$

— is fixed by well-understood plasma physics once ω_b and ω_m are known, and is computed to a fraction of a percent by the standard Boltzmann codes (Blas et al., 2011). The acoustic pattern is this ruler, photographed on the sky at $z_* \approx 1100$.

Step 3 — assume a history, and solve. An angle is a ruler divided by a distance,

$$\theta_* = \frac{r_s}{D_M(z_*)}, \quad D_M(z_*) = \int_0^{z_*} \frac{c dz}{H(z)}, \quad (4)$$

so the measured θ_* and the computed r_s together fix the comoving distance to last scattering, $D_M \approx 13.9$ gigaparsecs. To extract H_0 from it one must assume the *shape* of $H(z)$ across the thirteen-billion-year integral — in Λ CDM, $H(z) = H_0 \sqrt{\Omega_m(1+z)^3 + \Omega_r(1+z)^4 + \Omega_\Lambda}$ — and solve for the normalization H_0 that makes the integral come out right. The answer is 67.36 ± 0.54 from Planck alone (Planck Collaboration, 2020) and 67.24 ± 0.35 from the combined Planck+SPT+ACT analysis (Camphuis et al., 2025).

The structure of the risks again follows from the structure of the argument. The route is superbly precise and, as §15 shows, instrumentally redundant; but it is *model-bound at both ends*. The ruler end depends on pre-recombination physics: add new energy density before decoupling and r_s shrinks. The conversion end depends on the assumed thirteen-billion-year history: bend $H(z)$ between last scattering and today and the inferred H_0 moves. That double model-boundness is the lever arm every early-universe “solution” of the tension pulls (Part IV quantifies it); the same logic run in reverse — calibrating the ruler at high redshift and descending — is the “inverse distance ladder,” and it lands with the CMB (Aubourg et al., 2015).




Figure 7: Anatomy of the two routes. The late route climbs: geometric anchors calibrate primary indicators, which calibrate supernovae, which reach the Hubble flow — each rung inheriting the calibration of the one below, so the result is direct but systematics-compounding. The early route descends: the acoustic pattern fixes θ_* and the physical densities, the sound horizon is computed, and H_0 is *inferred* through an assumed expansion history between $z_* \approx 1100$ and today — precise but model-bound. The two meet, 7.1σ apart, at the present value of one function $H(t)$; Part III replaces the assumed history with a derived one.

10.3 The comparison, and the three possible homes

Table 2 compresses the anatomy into one view. The two routes share *nothing*: no data, no instrument, no calibration, and — most importantly — no philosophy. One measures a velocity field within a few hundred megaparsecs and never asks what the universe is made of; the other measures a pattern imprinted at $z \approx 1100$ and cannot speak of today without a model of everything since. Their disagreement is therefore not a technical squabble between rival teams but a confrontation between the two ways cosmology knows anything at all.

The asymmetry matters more than the disagreement. A discrepancy between a direct-but-calibration-heavy number and a precise-but-model-bound one has exactly three possible homes: a shared calibration bias in the ladder, new physics in the pre-recombination universe (a shorter ruler), or a mis-specified expansion history bridging the two epochs (a bent conversion). Every proposal in the literature — and this paper too — is a bet on one of the three.

Where this paper stands can now be said in one sentence: GG-Theory removes the bridge’s freedom — the expansion history between the early and late epochs is *derived*, with no continuous parameter (Part III) — and with the bridge fixed and the early side shown unbiased, the question is relocated to the ladder’s calibration, where the framework prices the one environmental channel and does not carry it, names the whole excess ($\approx 10\%$) as a predicted calibration systematic (Part IV), and calls for the one channel that owes nothing to either route — the gravitational-wave standard siren — to decide (Part VI).

Table 2: The two routes compared. Each row isolates one structural difference; the last two rows carry the logic of the rest of the paper — what each route must fear, and what it would take to bias each by the full nine percent of the tension. The two columns are not two measurements of the same thing made in different ways. The late route reads the flow it sits in and assumes nothing about the history; the early route computes today’s rate from a model of the whole history and assumes a great deal. A reader who wants one of the two to be simply wrong should say which of the two dominant risks is the one that is realized.

	Late route (ladder)	Early route (CMB)
Observable	recession velocities and distances at $z \lesssim 0.15$	the acoustic pattern at $z_* \approx 1100$
Physics used	geometry plus stellar standard candles	plasma physics plus a Λ CDM history
Character	<i>direct</i> : reads the local flow	<i>inferred</i> : computes today’s rate
Result	73.50 ± 0.81 (1.09%)	67.24 ± 0.35 (0.52%)
Cosmology assumed	none (kinematic priors only)	the shape of $H(z)$ over 13.8 Gyr
Dominant risk	a calibration bias shared across rungs	a mis-modeled ruler or bridge
A 9% shift needs	a coherent 9% zero-point bias	a $\sim 9\%$ shorter r_s , or a bent $H(z)$

11 Why this measurement matters

For more than a decade the *Hubble tension* — the mismatch between the early-universe and late-universe values of H_0 — has been the most scrutinized anomaly in precision cosmology (Di Valentino et al., 2025; Kamionkowski and Riess, 2023; Verde et al., 2024). The early-universe route (cosmic microwave background plus Λ CDM) and the late-universe distance-ladder route were, historically, dominated by a small number of teams and a small number of methods. Critics of the tension often framed it as the disagreement of “two or three measurements” against one very precise CMB inference, and argued that an *unknown unknown* — a silent systematic error in the local distance ladder — might eventually dissolve it.

The H0DN Collaboration set out to retire this framing by constructing a *Local Distance Network* that is neither one ladder nor a handful of parallel ladders but a full covariance-weighted combination of every credible local probe. The paper is the written output of the *What’s under the Hood?* workshop at ISSI-Bern in March 2025 and was co-authored by 38 specialists — including representatives of every major subfield of local distance determination (parallaxes, detached eclipsing binaries, masers, Cepheids, the tip of the red giant branch, Miras, JAGB, Type-Ia supernovae, Type-II supernovae, surface-brightness fluctuations, the fundamental plane, and Tully–Fisher). The core result is a 1.1% consensus value whose central number is insensitive to which subset of methods is used. The point is structural, not numerical: the late-universe measurement of H_0 can no longer be retired by identifying a systematic in one method, because no one method determines the result.

12 The central claim: the Hubble rate as an open-system observable

It is worth stating, in one place, the recognition this paper shares with the whole GG-Theory cohort. No physical system is closed — so the four-dimensional world is a lawful projection (OGGEP,




Figure 8: The central claim, as one deductive funnel, shared across the GG-Theory cohort. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection *is* observation itself; its dynamics is GDOS, forced by the Gi-4 and Qi-4 pillars. The same funnel opens the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs — here, the derived Hubble rate.

AX-3) of the six-dimensional bulk, and that projection *is* observation itself; its dynamics is the open-system law GDOS (GG-A5-GDOS), forced by the geometric Gi-4 (GG-A6-GRIP) and quantum-information Qi-4 (GG-A7-Qi4) pillars. It follows that every measured value of H_0 is a *registered observable* of one projection: one background rate, which the geometry derives, read through the calibration chain of each instrument. That is why this paper can derive the rate, why it predicts the calibrated ladder to be the one instrument displaced from it, and why a single clean ladder-free measurement can refute it. The funnel of Figure 8 is shared, in structure, with the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs.

13 Architecture of the Local Distance Network

13.1 Definitions

The H0DN paper re-codifies the vocabulary of distance-ladder work into a consistent six-layer structure (paraphrasing Section 2.2 of [Casertano et al. 2026](#); we do not reproduce their prose):

1. **Anchor:** an object whose distance is geometric, e.g. via parallax, detached eclipsing binaries, or accretion-disk megamasers. The H0DN baseline uses three anchors: the Milky-Way Cepheid field (Gaia EDR3 + HST parallaxes), the Large Magellanic Cloud (DEBs; [Pietrzyński et al. 2019](#)), and NGC 4258 (megamasers; [Reid et al. 2019](#)).

2. **Primary distance indicator:** a stellar feature directly calibratable in an anchor (Cepheids, TRGB, Miras, JAGB).
3. **Host:** typically a galaxy whose distance is set by primary indicators and which contains a secondary indicator.
4. **Secondary distance indicator:** a feature measurable at larger distances, typically Type-Ia supernovae, Type-II supernovae, SBFs, the fundamental plane, or Tully–Fisher.
5. **Calibrator:** an object whose distance constrains the zero-point of a secondary indicator.
6. **Tracer (Hubble-flow object):** an object distant enough that its redshift reliably reflects cosmological expansion, yielding a constraint on H_0 itself.

Above this scaffolding the H0DN places a genuinely new object: a *covariance matrix* that explicitly tracks which uncertainties are shared between paths. It is the structured treatment of off-diagonal covariance that separates the H0DN from earlier “whisker-diagram” summaries.

13.2 The baseline pre-registration

A critical methodological feature of the H0DN is its *pre-registration*. Before any calculations yielding H_0 , the workshop voted, by anonymous ballot, on which methods were sufficiently technically mature to constitute the baseline versus which belonged in variants. The results (Tables 2 and 3 of Casertano et al. 2026) placed Cepheids, TRGB, SNe Ia, SBF, megamasers in the Hubble flow, Gaia parallaxes, DEBs, and NGC 4258 as baseline; Miras, JAGB, SNe II (both SCM and EPM; Vogl et al. 2025), DESI FP (Said et al., 2025), Tully–Fisher, and the SMC Cepheid anchor (Breuval et al., 2024) as variants. These pre-registered choices fix the central value before the data are combined; the variants then explore robustness. This is how a 38-person collaboration avoids post-hoc cherry-picking.

14 The baseline result and its structural features

The H0DN baseline value is

$$H_0^{\text{H0DN}} = 73.50 \pm 0.81 \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad \chi_{\text{red}}^2 = 0.9879, \quad \text{PTE} = 0.49, \quad (5)$$

corresponding to a relative uncertainty of 1.09% (Casertano et al., 2026). Several features of this result, all of which the H0DN paper verifies with its variant catalog, are worth isolating.

Cross-method stability. The dispersion of central values across the full list of thirty structured variants is $0.40 \text{ km s}^{-1} \text{ Mpc}^{-1}$, about half of the 1σ uncertainty of the baseline itself. No single method, when removed, moves the central value by more than $\sim 1 \text{ km s}^{-1} \text{ Mpc}^{-1}$; no single method, when added, improves the central value significantly either. This is the hallmark of a system limited by *consensus rather than by systematics*.

SNe Ia are not the answer or the problem. Variant V13 removes *all* SNe Ia from the network. The central value changes by $-0.07 \text{ km s}^{-1} \text{ Mpc}^{-1}$. The uncertainty roughly doubles (because SNe Ia extend the ladder deepest into the Hubble flow), but the central value is essentially unchanged. Whatever one’s view of SN Ia calibration debates (Brout et al., 2022; Murakami et al., 2023; Scolnic et al., 2022), this variant retires the hypothesis that the local-universe value of H_0 is a SN Ia artifact.

Neither Cepheids nor TRGB alone drives the result. Variant V08 removes all Cepheids; V09 removes all TRGB; V10 removes all Gaia parallaxes; V11 removes the LMC as an anchor; V12 removes NGC 4258. None of these removals shifts the central value by more than $\sim 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ once the corresponding custom baseline is used to isolate the effect of losing SNe Ia calibrators. The anchor and the primary-indicator axes of the network are *redundant*, not hierarchical.

Instrument swaps agree. Variant V16 excludes all HST-based observations (forcing LMC and Milky Way to drop as anchors). Variant V17 excludes all JWST-based observations. These two variants are *not* exact complements, because both share NGC 4258, but they represent the hardest practical instrument-systematics hypothesis one can formulate. Their values bracket the baseline: $73.75 \pm 1.33 \text{ km s}^{-1} \text{ Mpc}^{-1}$ without HST, and $73.00 \pm 0.86 \text{ km s}^{-1} \text{ Mpc}^{-1}$ without JWST. The baseline sits cleanly between them.

The “everything” solution. Variant V99 turns on every method simultaneously and gives $73.99 \pm 0.70 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Variant V99A excludes Tully–Fisher, which generates a χ^2 excess, and gives $73.66 \pm 0.71 \text{ km s}^{-1} \text{ Mpc}^{-1}$ at 0.96% relative precision. Both of these sit within $0.5\text{--}1\sigma$ of the baseline and increase the tension with Planck+SPT+ACT to $\sim 8\sigma$.

Orthogonal paths agree. The H0DN performs a crucial consistency check by splitting the network into two statistically *fully independent* paths (Section 5.3.1 of Casertano et al. 2026):

$$\text{O1} = \text{MW+LMC/SMC} \rightarrow \text{Cepheids} \rightarrow \text{SNe Ia} + \text{FP},$$

$$\text{O2} = \text{NGC 4258} \rightarrow \text{TRGB} \rightarrow \text{SBF} + \text{megamasers}.$$

O1 yields $73.11 \pm 0.92 \text{ km s}^{-1} \text{ Mpc}^{-1}$; O2 yields $73.43 \pm 1.80 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Two statistically independent local probes of H_0 agree to better than 0.2σ . This is the strongest single discriminator against the “unknown unknowns” narrative.

15 The early-universe network: Planck, ACT, SPT, and DESI

Everything so far could leave the impression that network-grade redundancy is the late route’s monopoly — one consensus band standing against one satellite’s model fit. The impression would be wrong, and honesty requires stating the early side’s case with the same care, because it is structurally the same case.

The microwave background has now been measured by three independent experiments with separate telescopes, detectors, observing sites, sky coverage, and analysis pipelines. Planck’s final

release gives $67.36 \pm 0.54 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Planck Collaboration, 2020). The Atacama Cosmology Telescope’s sixth data release — ground-based, higher in resolution, analyzed end to end by a different collaboration — lands on a consistent six-parameter model (Louis et al., 2025). SPT-3G at the South Pole is the third instrument, and the joint analysis of all three yields $67.24 \pm 0.35 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Camphuis et al., 2025) — a half-percent measurement in which no single experiment is indispensable. The redundancy is internal as well: temperature, polarization, and gravitational lensing each carry the acoustic scale, and the Λ CDM parameters inferred from each sector separately agree (Planck Collaboration, 2020).

More telling still, the early value survives with the microwave background removed entirely. The same sound horizon that patterns the CMB is imprinted on the clustering of galaxies, and DESI’s second data release measures that baryon-acoustic ruler across multiple tracer classes out to redshift two and beyond. Calibrating it with the primordial deuterium abundance — big-bang nucleosynthesis fixes ω_b with no CMB input at all — yields $H_0 = 68.51 \pm 0.58 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Abdul Karim et al., 2025): 5.0σ from the H0DN, through a channel that shares no instrument, no sky, and no epoch with Planck. The “inverse distance ladder,” which calibrates supernovae against the BAO ruler and descends, lands in the same place (Aubourg et al., 2015); and cosmic chronometers, which read $H(z)$ directly from the differential ages of passively evolving galaxies, trace an expansion history consistent with the early-anchored curve at intermediate redshift (Jimenez and Loeb, 2002; Moresco et al., 2022).

The balance sheet is therefore symmetric. The late side is a network: no single method, anchor, or instrument moves the consensus (§14). The early side is a network too: three CMB experiments, a CMB-free acoustic ruler, and a direct $H(z)$ probe agree among themselves in the 67–68.5 band. A proposed systematic-error resolution now faces the same structural objection in either direction: blaming Planck must also explain ACT, SPT, and DESI+BBN; blaming the Cepheids must also explain TRGB, the megamasers, and the orthogonal paths.

And the significance only grows when channels are combined: the H0DN baseline sits 7.1σ from the joint CMB value and 5.0σ from the CMB-free one, while remaining consistent with independent late-universe probes such as the TDCOSMO lensing time-delay program (Birrer et al., 2025). The tension is not one measurement against another; it is network against network — which is precisely what makes it hard, and precisely what gives Part III its opening. The H0DN paper is the empirical basis for retiring the single-systematic hypothesis on the late side; this section is the parallel basis on the early side.

16 The statistics of 7.1σ : what the number does and does not say

The headline significance is elementary arithmetic: the two central values differ by $73.50 - 67.24 = 6.26 \text{ km s}^{-1} \text{ Mpc}^{-1}$, the quadrature error is $\sqrt{0.81^2 + 0.35^2} = 0.88$, and the ratio is 7.1. If the quoted uncertainties are complete, Gaussian, and independent, the chance of a discrepancy this large is about one part in 10^{12} — far beyond any conventional discovery threshold. One bookkeeping sentence, so the cohort’s two headline significances are never mistaken for a disagreement: the 7.1σ of this paper is the H0DN 73.50 ± 0.81 against the 2025 three-instrument combination 67.24 ± 0.35 ;

the $\sim 5\sigma$ the parent GG-A10-Cosmos prints is Riess’s 73.0 ± 1.0 against Planck’s 67.4 ± 0.5 . Same tension, two pairs. It is worth being precise about what the number does and does not establish.

Each of the three assumptions has a different standing. *Independence* is solid: the two routes share no data, no instrument, and no calibration. *Gaussianity* is supported on the local side by the network’s own goodness-of-fit ($\chi^2_{\text{red}} = 0.99$, PTE = 0.49) and on the early side by the likelihood machinery of the CMB analyses. *Completeness of the error budgets* is the entire debate. A significance is a statement of incompatibility *under the stated errors*: it measures how loudly the two numbers disagree, not why. No amount of σ can arbitrate between “new physics” and “a systematic the budget does not contain,” because both produce exactly the same arithmetic.

The sensitivity to that last assumption is worth quantifying, because it is the pivot of Part IV. The H0DN covariance already prices every *modeled* shared systematic — that is the network’s great advance. But a common-mode calibration term of fractional size s , coherent across the network and absent from the model, adds $s \times 73.5$ in quadrature to the local error, and the significance collapses quickly: $s = 2\%$ brings 7.1σ down to 3.7σ ; $s = 4\%$ to 2.0σ ; $s = 6\%$ to 1.4σ ; and $s = 10\%$ — the size of the systematic this paper will name in Part IV — to 0.8σ . The variant catalog cannot exclude such a term, because variants that share calibration DNA move together by construction. The 7.1σ is therefore both real and conditional: unimpeachable as arithmetic, and silent on the one hypothesis that matters most.

History suggests how such stand-offs end. The factor-two war was not settled by either camp accumulating significance — both sides had plenty — but by a new instrument that made the contested calibrations measurable. The modern analogue is a channel whose systematics are disjoint from both routes: the gravitational-wave standard siren, which needs neither a sound horizon nor a ladder. It is the referee this paper commits to in advance (§57), and the reason the reader will find no attempt in what follows to litigate the tension by significance alone.

17 What the H0DN paper does *not* do

It is equally important to be honest about what the H0DN does *not* claim, because several readings of the tension have to be carefully distinguished.

No theoretical resolution. The H0DN offers no model to explain *why* local and CMB-inferred values disagree. It only demonstrates that the disagreement cannot be retired by dropping any single method on the local side.

No ruling on early-universe systematics. The H0DN is silent on whether the CMB inference, recombination history, or BBN+BAO inference is itself modified by new physics. A large literature now canvasses early dark energy, extra radiation, interacting dark sectors, and modified recombination as options; (Kamionkowski and Riess, 2023; Poulin et al., 2023); the H0DN does not arbitrate among these.

No anisotropic claim. The paper does *not* claim that H_0 varies across the sky, nor does it claim a large local void. Variant V19 tests CMB-frame velocities (no peculiar corrections) and shifts the

result by only $-0.44 \text{ km s}^{-1} \text{ Mpc}^{-1}$; the various redshift-cut variants V20–V21 shift the result by at most $\pm 0.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Large-scale isotropy and the Carrick et al. (2015) peculiar-velocity model are approximately consistent.

No closed book on $w(z)$ or Ω_m . The H0DN fixes $q_0 = -0.55$ and $j_0 = +1$ (from Riess et al. 2022) as input priors. Varying these by reasonable amounts changes H_0 by less than $0.1 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Since this report is centrally concerned with cosmological theory, it is important to recognize that H0DN is only mildly sensitive to the background expansion model.

18 Summary: the target for any theory, and the road to Part III

Part II set out to pose the problem the rest of this paper must solve, and it can now be posed exactly. This closing section gathers the whole Part into a few paragraphs, ending with the single object any theory of the tension must reproduce.

The Hubble parameter $H(t)$ is the fractional rate at which cosmic distances stretch; the Hubble constant H_0 is its present-day reading. Measuring H_0 consumed a century (§8): the value started a factor of seven too high, fell down two great calibration cliffs, stalled in a thirty-year factor-two war, and converged at 72 ± 8 when the Key Project closed the old dispute. The century’s lesson cuts both ways — calibration systematics have repeatedly moved this number by amounts that dwarf today’s tension, and the modern network was engineered precisely so that this could not happen again.

Meanwhile the discovery of cosmic acceleration (§9) turned the number into a curve. A universe that decelerated for eight billion years and then began to accelerate has a genuinely time-dependent $H(t)$, and “the Hubble constant” is one reading of that function at the single epoch we occupy. The concordance model that absorbed the discovery — Λ CDM: flat, one-third matter, two-thirds dark energy — fits the curve but explains neither the milli-electron-volt scale of the dark energy nor the why-now coincidence; both stand as open invitations to a deeper theory.

The tension itself is a disagreement between the two ways cosmology measures the curve (§10). The late route climbs a ladder — geometry, Cepheids and TRGB, supernovae, the Hubble flow — and reads the local rate directly, assuming no cosmology and inheriting every calibration: 73.50 ± 0.81 , a 1.09% measurement. The early route photographs the sound horizon at $z_* \approx 1100$, computes the ruler from plasma physics, and infers today’s rate through an assumed expansion history: 67.24 ± 0.35 , a half-percent inference. A discrepancy between numbers of these two characters has exactly three possible homes — a shared calibration bias in the ladder, a shorter ruler from new pre-recombination physics, or a bent expansion history in between — and every proposal in the literature is a bet on one of the three.

Both sides are now networks. The H0DN consensus is not one team’s ladder but a covariance-weighted combination of every credible local probe — pre-registered before combination, stable under the removal of any single method, and confirmed by two statistically independent orthogonal paths (§14). The early side answers with three independent CMB experiments, a CMB-free acoustic ruler calibrated by nucleosynthesis (68.51 ± 0.58), and direct $H(z)$ chronometers (§15). Neither side can be retired by finding a bug in one instrument.

The arithmetic of the disagreement is unimpeachable — a gap of $6.26 \pm 0.88 \text{ km s}^{-1} \text{ Mpc}^{-1}$, a significance of 7.1σ , odds of one part in 10^{12} under the stated errors — and it is conditional (§16): a coherent two-percent calibration term absent from the local error model would already collapse it to 3.7σ . No accumulation of significance can arbitrate between “new physics” and “a systematic the budget does not contain,” because both produce identical arithmetic; and the H0DN itself claims neither, offering no theoretical resolution, no ruling on early-universe physics, and no anisotropy (§17).

The object a theory must reproduce is therefore not an individual measurement but the *network-level consensus band*:

$$72.7 \text{ km s}^{-1} \text{ Mpc}^{-1} \lesssim H_0^{\text{H0DN}} \lesssim 74.3 \text{ km s}^{-1} \text{ Mpc}^{-1} \quad (6)$$

(baseline $\pm 1\sigma$), standing 7.1σ from the early-universe network. Any account of the tension must meet this band at the network level: either produce it through structurally inherited mechanisms, or locate, explicitly and falsifiably, the bias that creates it. That is the criterion we apply in the rest of this report.

Nothing in this Part has been **GG-Theory**. It is the problem as any cosmologist would pose it: a background rate near 67 inferred from the early universe by a redundant network, a local rate near 73 measured by an equally redundant ladder, and a nine-percent gap that has survived a decade of cross-checks — sharpened into one band a theory must account for and one split a theory must explain. The Parts that follow do not adjust this target; they derive the rate and predict which instrument is displaced from it. Part III derives the expansion history $H(t)$ — with no continuous parameter — and lands, without being asked, on the early side of the split; Part IV must then explain the late side, and does so by pricing it; and Part VI names the referee — a channel that owes nothing to either network — whose verdict this paper agrees in advance to accept.

19 Part II in plain English: the problem as any cosmologist would state it, before any geometry is spent

There is no geometry anywhere in this Part, and the omission is deliberate. Before a framework is allowed to explain a discrepancy, the discrepancy should be stated by somebody with no stake in the explanation.

The late route begins with distances that can be paced off — geometric parallaxes to nearby stars — and works outward in steps, each step calibrating the next, until it reaches galaxies far enough away that the expansion dominates their motion. It assumes nothing whatever about the history of the universe. Its risk is structural: an error made at the bottom is carried, unchanged and unnoticed, all the way to the top.

The early route begins with the pattern of hot and cold spots in the microwave background, which is a photograph of sound waves frozen into the plasma at the moment the universe turned transparent. From the size of those waves, and a model of everything that has happened since, one computes what the expansion rate must be today. It assumes a great deal about the history. Its

risk is that the model, or the ruler the pattern is measured against, is wrong.

Both have been redone by independent teams using independent instruments, and both have stayed where they were. The local network now sits in a band from 72.7 to 74.3; the early network sits near 67.2 with a smaller error bar. The separation is 7.1 standard deviations, and a decade of looking has not turned up a single mistake that would account for it.

So the answer can only be in one of three places: a bias shared by every rung of the ladder, new physics before the universe turned transparent, or an expansion history that bends between then and now in a way nobody has modeled. Every proposal in the literature is one of those three, this paper's included. The Parts that follow say which one this framework is forced into, and why it has no choice in the matter.

Part III

The GG-Theory of the Hubble Parameter

The Hubble parameter as a GDOS/GRIP observable; the direct derivation of $H(t)$ and H_0 ;
and why it reproduces the CMB value

20 The GG-Theory framework in one view

This report does not re-derive GG-Theory; for the full account we refer to the synthesis paper GG-A11-Origin and to the component preprints (Arisaka, 2026, A1-Constitution; Arisaka, 2026, A9-SM; Arisaka, 2026, A5-GDOS; Arisaka, 2026, A10-Cosmos; Arisaka, 2026, A7-Qi4). We summarize only the pieces that enter the H_0 derivation. Everything used below is inherited from the constitution layer (the axioms and postulates of Arisaka 2026, A1-Constitution) and from the theorem ledger of the component papers; this report introduces no new postulate and no new continuous parameter, and where an entry is not yet theorem-level (the selection lemma; the measure lemma) it is flagged explicitly.

GG-Theory treats what is usually a stack of independent postulates (geometry, gauge, measurement, openness) as forced consequences of a single constitution built from four axioms (AX-1 LC, AX-2 GGEP, AX-3 OGGE, AX-4 MSMF) and seven postulates (POS-1 GI, POS-2 UEP, POS-3 AF, POS-4 CP, POS-5 RG, POS-6 OGN, POS-7 MF). The stage on which everything plays out is the six-dimensional bulk¹ of GG-A2-GG6 (Arisaka, 2026, A2-GG6), and it is best met through its metric, which is definite and parameter-free:

$$ds_6^2 = e^{-2r} \eta_{\mu\nu} dx^\mu dx^\nu + dr^2 + d\theta^2, \quad M_6 = M_4 \times I_r \times S_\theta^1, \quad (7)$$

with $\eta_{\mu\nu}$ the Minkowski metric of the four-dimensional boundary, $r \in [0, K]$, and $\theta \in [0, 2\pi)$.

Three pieces make up the bulk. The four dimensions M_4 are the spacetime we live in — more precisely, the boundary on which every observation is registered. The radial direction I_r is the energy axis: the coordinate is literally the renormalization-group logarithm, $r = \ln(E_{\text{GUT}}/\mu)$, running from the ultraviolet mouth at $r = 0$ ($E_{\text{GUT}} \simeq 2.3 \times 10^{16}$ GeV) down to the integer endpoint $r_{\text{IR}} = K = 35$; “deeper into the throat” simply means “lower in energy,” and the Planck–electroweak hierarchy is then the topological number e^{-K} , not a fine-tuned ratio. One coordinate, two anchors — worth one explicit sentence, because the same integer appears in both roles. The GUT mouth sits 2π e-folds below the Planck scale ($E_{\text{GUT}} = M_{\text{Pl}} e^{-2\pi}$: $\ln(1.22 \times 10^{19}/2.28 \times 10^{16}) = 6.28$), so the lock mass, whose formula $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-K}$ counts its K e-folds from the *Planck* anchor, sits *inside* the

¹Throughout the cohort, the “warped six-dimensional bulk” is read on the field-theory side of the holographic dictionary: the fifth (scale) dimension is the renormalization-group axis $r = \ln(E_{\text{GUT}}/\mu)$, not a Randall–Sundrum spatial dimension, and “warped” denotes RG running; the explicit holographic-RG construction is given in Arisaka (2026, A2-GG6).

throat at $r_{\text{lock}} = K - 2\pi + \ln 2 = 29.4$, not at the endpoint; and on the boundary side the same RG coordinate simply continues past the bulk interval as physics cools below the mouth’s image — the cohort’s descent charts run it to the terminus $r \approx 64.5$ (the dark-energy floor) and, tracked by the relic temperature, to today’s $r_{\text{now}} = \ln(E_{\text{GUT}}/T_0) \approx 66.7$ (the $H(t)$ chart’s radial axis carries 66.6, an anchor-rounding offset of 0.1 e-fold noted here so no reader hunts for physics in it). The interval $[0, K]$ is the bulk’s radial extent; the coordinate is the boundary’s energy logbook, and the logbook keeps running. And the small compact circle S^1_θ is the phase fiber: it carries the integer-quantized Wilson-line holonomy, and its winding is where the dark sector — and, one story up, the π -locked inventory of §22 — will be read off. The warp profile in the exponent, $A(r) = r$, is not chosen but forced: demanding constant bulk curvature with a flat boundary gives $A''(r) = 0$, so there is no curvature scale, no compactification radius, and no brane tension to tune.

On this bulk the gauge and gravitational structure is a single group (Arisaka, 2026, A2-GG6),

$$G_{\text{Bulk}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}, \quad (8)$$

with the Lorentz/spin factor $\text{Spin}(1, 5)$ — the gravity sector — on an equal footing with the internal factors: $\text{SU}(5)$ is the minimal simple unification containing the Standard Model; $\text{U}(1)_B$ is gauged baryon number, which turns proton stability into a theorem; and $\text{U}(1)_{Q_i}$ is the informational gauge force the Qi-4 pillars organize. Projected onto the four-dimensional boundary, this one group becomes the six forces an observer actually measures — strong, weak, electromagnetic, baryonic, informational, and gravitational. The projection’s legality certificate is the modified Bianchi identity proved anomaly-free in GG-A3-BI6 (Arisaka, 2026, A3-BI6),

$$\text{BI-6 :} \quad dH_3 = \text{tr } R^2 - \text{Tr } F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Q_i}^2, \quad (9)$$

and its arithmetic is the Odd-integer Gauge Normalization ledger proved in GG-A4-Euclid (Arisaka, 2026, A4-Euclid),

$$\text{OGN-5 :} \quad (M, N, L; J, K) = (1, 3, 5; 7, 35), \quad K = LJ = M^2 + N^2 + L^2.^2 \quad (10)$$

Equations (9) and (10) together are the “integer ledger” that will fix the corridor-switch redshift $z_* = e^7$ and the seesaw coefficient ζ_Λ entering the derived rate.

From this geometry the Standard-Model closure GG-A9-SM (Arisaka, 2026, A9-SM) derives all 28 Standard-Model parameters as algebraic expressions in $(1, 3, 5; 7, 35)$ and the transcendentals $\{\pi, e\}$ plus the dimensionful $\{c, \hbar, G_N\}$. The key exports are collected in Table 3 — the ledger’s already-spent predictions, each parameter-free, spanning the electroweak sector this paper inherits and the cosmological shares it will re-derive. Two further exports matter here and do not fit a numerical column: strong-CP protection, $\bar{\theta} \lesssim 4.7 \times 10^{-18}$, and exact proton stability, $\tau_p = \infty$ (exact local $\text{U}(1)_B$).

²Notation, kept uniform across the cohort: N is the OGN-5 ledger’s second entry, and every ledger identity is written with it. $N_c = 3$ is the color triplicity, the three colors of $\text{SU}(3)_C$ — what that entry counts, named where the color reading is the point. $N_g = 3$ is the generation (family) count; it is *not* a ledger entry and is not fixed by the ledger (GG-P2-N3). $N_f = 2N_g = 6$ is the quark-flavor count of a beta function at high scale. Three of these are three and one is six, and they are three different quantities.

Table 3: What the same five integers have already bought (parent papers GG-A9-SM and GG-A10-Cosmos; the format and values of the companion GG-C1-Kaxion’s track-record table). Every entry is parameter-free; deviations are against PDG/CODATA/Planck values. Deviations, not pulls, are the honest column here, and we say why: these are *first-order* closed forms whose known second-order (QD) completions shift them at the 0.1–1% level (GG-A9-SM); in raw measurement units the rigid entries sit 8–9 σ out (+0.05% on $\hat{\alpha}^{-1}$ is +8.8 σ), and closing that residual is GG-A9-SM’s audit, not this table’s. The claim here is parameter-free landing at the permille-to-percent level — exactly that, and nothing sharper.

Quantity	OGN-5 formula	Predicted	Measured	Dev.
$\hat{\alpha}^{-1}(M_Z)$	2^J	128	127.930 ± 0.008	+0.05%
$\alpha^{-1}(0)$	$2^J + N^2$	137.0	137.036	−0.03%
$\sin^2 \theta_W(M_Z)$	$L^2/(4N^3)$	$25/108 = 0.23148$	0.23122 ± 0.00003	+0.11%
$\alpha_s(M_Z)$	$1/(\pi e)$	0.1171	0.1180 ± 0.0009	−0.8%
Higgs vev v	$M_{\text{lock}}(2/L)^N$	246.33 GeV	246.22 GeV	+0.05%
Decoupling z_{dec}	atomic chain	1091.5	1089.92 ± 0.25	+0.15% ³
Matter share Ω_m	$1/\pi$	0.3183	0.3153 ± 0.0073	+1.0%
Dark-matter share Ω_{dm}	$1/\pi - 1/(2\pi^2)$	0.2676	0.2645 ± 0.0070	+1.2%
Dark-energy share Ω_Λ	$1 - 1/\pi$	0.6817	0.6847 ± 0.0073	−0.4%

These are the *same* integer package that will be asked to produce the expansion rate. The logical challenge is therefore: can the ledger that pins Standard-Model parameters to $\lesssim 1.5\sigma$ also derive the rate the sky reads, with no new tuning?

20.1 What a theory of geometry may output: the GU-5 dimensional bridge

Before the rate is derived, one objection should be met head-on, because it decides how the whole of §24 must be read: *can a theory of pure geometry output a dimensionful number at all?* The Hubble rate carries units — inverse time, dressed by astronomers as $\text{km s}^{-1} \text{Mpc}^{-1}$ — and a geometry, it would seem, can only deliver pure numbers like $1/\pi$. The answer is the founding move of GG-A4-Euclid (Arisaka, 2026, A4-Euclid): GG-Theory declares its dimensional inputs once, as the five-element set $\text{GU-5} = \{c, G_N, \hbar, k_B, Q_B\}$, and GG-A4-Euclid proves that the first three are not physics parameters but the *conversion factors between the dimensions of the six-dimensional geometry* — length to time (c), geometry to energy (G_N), energy to frequency ($\hbar$) — which is exactly why the natural-unit choice $\{c, G_N, \hbar\} = \{1, 1, 1\}$ costs nothing. Once that bridge is declared, “dimensionful” and “dimensionless” cease to be different kinds of prediction: every dimensionful observable is a derived pure number times the appropriate combination of $E_{\text{Pl}} = \sqrt{\hbar c^5/G_N}$ and its siblings.

This is not a promissory note; it is how GG-A9-SM already operates. The electron’s charge is the derived pure number $e = \sqrt{4\pi\alpha}$ with $\alpha^{-1}(0) = 2^J + N^2 = 137$ — a dimensionful charge, delivered by integers and the GU-5 bridge. The fermion masses are derived ratios times the scale chain, whose top rung is pure GU-5. The only empirical content left in a unit like “kilogram” or “ $\text{km s}^{-1} \text{Mpc}^{-1}$ ” is the SI definition of the meter and the second — bookkeeping, not physics.

³The corridor switch $z_* = e^7 = 1096.6$ is a topological marker of the integer $J = 7$, 0.38% above the atomic-chain decoupling in $1 + z$; it is annotated and not scored.

Applied to the present problem, the bridge turns the question “what is H_0 ?” into a pure-number assignment the geometry either meets or fails: in Planck units the observed rate is

$$H_0 t_{\text{Pl}} \approx 1.16 \times 10^{-61}, \quad \text{equivalently} \quad \rho_\Lambda / E_{\text{Pl}}^4 \sim 10^{-123}, \quad (11)$$

and *that* — sixty-one orders of magnitude, as a pure number — is what §24 owes the reader from geometry alone. No astronomical observation is entitled to appear anywhere in the delivery; and none will. The sky enters afterwards, in its proper role: as the measurement the derived number is tested against.

20.2 The recast cosmology GG-A10-Cosmos

GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) is the parent paper of the present report: the volume in which this machinery is run on the whole of cosmology, and the source of every background quantity Part III uses. Its central move is to change the status of the concordance model. Λ CDM is ordinarily specified by six fitted numbers; GG-A10-Cosmos first derives general relativity itself and then the Friedmann system as *theorems* of the boundary projection (its Part III; the derivation is transcribed in §21 below), and then shows that the integer ledger fixes the background half of the six parameters in closed form — $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/2\pi^2$, $w = -1$, and the corridor-switch redshift $z_* = e^7$ — so that 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.

The discipline that keeps this honest is *joint* solution. The guiding principle of GG-A10-Cosmos: inflation, dark matter, baryogenesis, BBN, the CMB, H_0 , S_8 , and JWST-era early structure must be solved jointly through one inherited source vector, or fail jointly. The admissibility criterion forbids the inclusion of any dark sector, early-universe fluid, or late-time component without geometric inheritance — nothing may be added to fix one observable, because nothing may be added at all. Under this discipline the dark matter is the Kaxion of the compact angle (identified, and staked to an exact search frequency, in GG-C1-Kaxion (Arisaka, 2026, C1-Kaxion)); the launch is an inflationary descent whose tilt is computed at $n_s = 0.9634$ — 1.5σ below the 2025 CMB-only value and 3.2σ below the DESI-combined one, the framework’s one currently failing test, printed as such in GG-A10-Cosmos and again in §56; and the terminus is dark energy at exactly $w = -1$, the ground geometric attractor whose constancy Part V’s kill switches guard.

What Part III imports from GG-A10-Cosmos is exactly four objects, each carried with its provenance grade. First, the inventory theorem — the π -locked fractions with their Hopf-fibration proof and the one named lemma still owed (§22). Second, the derived expansion history $H(t)$ that those fractions force through the Friedmann theorem (§21 and §23). Third, the two-route determination of the present rate, $H_0^{(0)} = 66.8 \pm 0.2$ — the vacuum-energy seesaw and the density anchors that confirm it (§24). Fourth, the over-determined microwave-background confrontation, ending in the coherent π -offset (§26). What this paper returns to its parent is the tension itself, read at full depth: the early/late split priced in Part IV, the rival-class landscape audit and comparison of Part V, and the standard-siren stake of Part VI — the sharpest dated falsifier the parent framework carries.

A reader who wants the derivations behind the imports should read GG-A10-Cosmos’s middle Parts alongside this one; a reader who wants the particle-physics floor under both — the same ledger pinning all twenty-eight Standard-Model parameters to $\lesssim 1.5\sigma$ — should read GG-A9-SM (Arisaka, 2026, A9-SM). Nothing below hides its grade: the fractions and everything downstream of them carry D* (derived, conditional on one named lemma), the lock scale is a theorem, the seesaw coefficient is Derived conditional on the codebook-selection lemma (§24.2), and the provenance table of §24 keeps the whole ledger in one place.

21 The Hubble parameter as a GDOS/GRIP observable

Before deriving a number, it is worth being clear about *what* is being derived. The object is the Hubble parameter $H(t) = \dot{a}/a$, the fractional stretching rate of the universe, of which H_0 is the present-day reading. This section establishes the load-bearing rung: the Friedmann law which turns an energy inventory into an expansion rate is itself a *theorem* of the framework, not an imported assumption. The next section then derives the inventory the law needs, and §23 feeds the one into the other, producing the whole curve $H(t)$ with no further freedom anywhere. The derivation is transcribed from GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos); we carry it out in full because this paper is about nothing else.

21.1 The Friedmann equations as a theorem

A cosmology is the special solution of the Einstein equation in which the universe looks the same at every place and in every direction — the cosmological principle. 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 $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ 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 the attractor is unique is a D* result, resting on the attractor-uniqueness statement of GG-A6-GRIP and the inflationary smoothing of GG-A10-Cosmos; we state the condition rather than bury it.) The Einstein equation itself is likewise inherited, not assumed: GG-A10-Cosmos derives it from the bulk action through the boundary projection, with G_N the throat overlap integral. Granting the attractor, the road from symmetry to the Friedmann system is the standard graduate computation (Dodelson and Schmidt, 2020; Weinberg, 2008), and we walk it, because it is the load-bearing rung.

Step 1 — the metric. Homogeneity and isotropy fix the metric up to one function of time and a discrete curvature sign (Friedmann, 1922; Lemaître, 1927). In comoving coordinates,

$$ds^2 = -dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2 d\Omega_2^2 \right], \quad (12)$$

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. In GG-Theory the discrete choice is not free either: the

boundary Friedmann export of GDOS carries $k = 0$, so flatness is inherited, not measured in.

Step 2 — the curvature. With an overdot for d/dt , the non-vanishing Christoffel symbols of the metric are

$$\Gamma_{ij}^0 = a\dot{a}\gamma_{ij}, \quad \Gamma_{0j}^i = \frac{\dot{a}}{a}\delta_j^i, \quad \Gamma_{jk}^i = \hat{\Gamma}_{jk}^i(\gamma),$$

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

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

and contracting gives the Ricci scalar $R_4 = 6(\ddot{a}/a + \dot{a}^2/a^2 + k/a^2)$. Assembling $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R_4 g_{\mu\nu}$, the independent components are

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

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.

Step 3 — the source. The same symmetry forces the boundary stress-energy to be that of a perfect fluid, $T^\mu{}_\nu = \text{diag}(-\rho, p, p, p)$: 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. The density and pressure are not free fluids to be fitted but the coarse-grained boundary projection of the bulk source ledger,

$$\rho = \rho_r + \rho_b + \rho_{\text{DM}} + \rho_\Lambda, \quad (13)$$

each sector's identity, equation of state $w_i = p_i/\rho_i$, and normalization fixed by the integer ledger of §20 rather than chosen.

Step 4 — the equations. Inserting the Einstein tensor and the perfect-fluid source into $G_{\mu\nu} = 8\pi G_N T_{\mu\nu}$ (the vacuum term carried inside ρ_Λ), the time–time component yields the first Friedmann equation,

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G_N}{3}\rho - \frac{k}{a^2} \quad (14)$$

and combining it with the trace of the space–space component gives the acceleration equation,

$$\frac{\ddot{a}}{a} = -\frac{4\pi G_N}{3}(\rho + 3p) \quad (15)$$

while covariant conservation, $\nabla_\mu T^{\mu 0} = 0$, gives the continuity equation $\dot{\rho} + 3H(\rho + p) = 0$, whose solution sector by sector is the dilution law $\rho_i \propto a^{-3(1+w_i)}$. *Physical picture:* the first equation is a first integral of the motion — energy density sets the expansion rate; the second says pressure gravitates, so a sector with $\rho + 3p > 0$ brakes the universe and a vacuum with $p = -\rho$ accelerates it; the third is simply dilution — matter thins as a^{-3} , radiation as a^{-4} (one extra power for the redshifting of each quantum), and a $w = -1$ vacuum not at all.

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 — which sectors exist, with which equations of state, in which proportions, and with which curvature. The Friedmann system is therefore a theorem of the framework: the boundary Einstein equation, restricted to the runtime-selected symmetric attractor, sourced by a ledger with no free dials. The law is now ready; what it still needs is its right-hand side, and deriving those densities is the next section's work.

22 The π -locked inventory: derived, not assumed

The rate $H(t)$ of §21 is only as derived as the densities that feed it, so the derivation must begin one story down, with the inventory itself. It is the central claim of the parent paper GG-A10-Cosmos that the present cosmic budget is not a set of measured fractions but a *theorem*:

$$\Omega_m = \frac{1}{\pi} = 0.3183, \quad \Omega_\Lambda = 1 - \frac{1}{\pi} = 0.6817, \quad \Omega_b = \frac{1}{2\pi^2} = 0.05066, \quad (16)$$

with the dark-matter share the difference $\Omega_{\text{dm}} = 1/\pi - 1/2\pi^2 = 0.2676$. These are not fitted; they agree with the Planck 2018 concordance values to better than three percent (Ω_Λ and Ω_{dm} within 1σ , Ω_b at 2.3σ), and their status is the subject of a named theorem. We state it, sketch its proof, and name precisely what it still owes, because everything downstream — $H(t)$, H_0 , the tension — inherits exactly this grade.

Theorem 22.1 ($\Omega_m = 1/\pi$ as the Hopf fiber-volume ratio, conditional (T-Hubble-1)). *Assume (I) the cosmological-boundary stratum is a connected, orientable four-manifold $\mathcal{M}_4$ obtained as the OGGEF projection of the GG-6 bulk; (II) the modified Bianchi identity BI-6 holds throughout cosmological history, $dX_4 = 0$ on $\mathcal{M}_4$ with coefficient vector $(1, -1, -2/5, -2/5)$; and (III) the OGN-5 integers $N = N_c = 3$ (colors) and $L = 5$ ($\text{SU}(5)$) fix the boundary holonomy manifold as the electroweak $\text{SU}(2) \cong S^3$, and the matter sector's charge as its Hopf fiber S^1 . Then, conditional on the measure-normalization lemma VIII-b'' (GG-A10-Cosmos's Conjecture VIII),*

$$\Omega_m = \frac{\text{Vol}(S^1)}{\text{Vol}(S^3)} = \frac{2\pi}{2\pi^2} = \frac{1}{\pi},$$

with no continuous adjustable parameter and no reference to the measured inventory.

Proof outline (five steps). Step 1 (closure). By hypothesis (II) the four-form X_4 is closed, $dX_4 = 0$, so it represents a class in $H^4(\mathcal{M}_4, \mathbb{R})$; the class is non-trivial because $\text{tr } R^2$ generates a non-trivial Pontryagin class on the curved boundary slice. *Step 2 (integer skeleton).* The OGN-5 tuple $(1, 3, 5; 7, 35)$ fixes the integer content: $N = 3$ and $L = 5$ identify the holonomy as electroweak $\text{SU}(2)$, and the packet is one canonical unit — a rational arithmetic weight $M/(LN) = 1/15$ is a provable dead end. *Step 3 (Pontryagin normalization).* $\int_{\mathcal{M}_4} \text{tr } R^2 = 8\pi^2 p_1$, and for the OGN-5 slice $p_1 = N = 3$, so the integral is $24\pi^2$; crucially the $8\pi^2$ is *consumed upstream* (the geometric free-energy functional 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**. By the Hirzebruch signature theorem ($p_1 = 3\sigma$),

$p_1 = 3$ forces $\sigma = 1$: the minimal closed gravitational topology, hence one canonical packet. *Step 4 (the Hopf ratio — where the π is born)*. The matter charge lives on the compact phase circle S^1_θ , which sits inside $SU(2) \cong S^3$ as its Hopf fiber $S^1 \hookrightarrow S^3 \rightarrow S^2$ (a canonical fibration, no freedom of choice). Under the frozen stationary GDOS measure each sector's present-epoch share is its fiber volume divided by the total holonomy volume, and both are canonical unit spheres, $\text{Vol}(S^1) = 2\pi$, $\text{Vol}(S^3) = 2\pi^2$, so $\Omega_m = 2\pi/2\pi^2 = 1/\pi$ exactly, with the π descending from the *denominator volume* and nothing read off the sky. *Step 5 (uniqueness — competitors die by geometry, not by data)*. $1/(3\pi)$ would require a total holonomy volume $6\pi^2$ that no manifold in the construction carries; $1/(2\pi)$ is the *baryonic* projection (one unit per fiber revolution), not the matter/total ratio; $1/\pi^2$ would need an S^5 that the four-dimensional boundary does not carry. The measured value of the inventory enters the selection at no step. ■

Corollary 22.2 (the complement and the baryon line). *The dark-energy share is the part of the S^3 holonomy volume the matter Hopf circle does not cover, a topological complement $\Omega_\Lambda = 1 - 1/\pi$; and baryons are one integer unit of $U(1)_B$ charge threaded once around the fiber, $\Omega_b = \Omega_m/\text{Vol}(S^1) = (1/\pi)/2\pi = 1/2\pi^2$. One fibration fixes three coupled numbers; none can be moved without the others — the structural signature that separates a derivation from a per-observable fit.*

What the theorem still owes. The one step not yet closed is the measure-normalization lemma: that the frozen GDOS measure weights exactly the Hopf circle of the electroweak $SU(2)$ as the matter fiber and $\text{Vol}(S^3)$ as the total (GG-A10-Cosmos's Conjecture VIII). Its structure is now well understood and partly discharged. It decomposes into *VIII-a* (delocalization and share-robustness of the stationary measure — proved in consumer form, with the gravity-mediated floor $\sim e^{-K} \approx 10^{-15}$); *VIII-b'* (the real-sector half-measure — *proved*: exact mode counting, $2\Lambda + 1$ real degrees of freedom against $2(2\Lambda + 1)$ charged, a ratio of exactly $1/2$, which forces $\Omega_b/\Omega_m = 1/2\pi$ with no residual freedom); and *VIII-b''* (the unit-packet normalization — the remaining open core), now traced to the crossing-epoch identity $u = 1$, i.e. the declared choice that the present epoch is read at the $\Omega_m = 1/\pi$ crossing. Until *VIII-b''* is exhibited the fractions carry the grade D^* (derived, conditional on one named lemma). This is a strict improvement over “an unexplained numerical coincidence”: no empirical selector enters (the fraction is not chosen for its agreement with the measured $\Omega_m/\Omega_\Lambda \approx 0.46$), and the residual is one clean geometric lemma, provable or falsifiable on the bulk construction.

The geometry the theorem trades in is drawn in Figure 9. What the proof does, in pictures: Steps 1–3 build one canonical packet of gravitational topology — closure makes the four-form a cohomology class, the integer skeleton picks the electroweak $SU(2) \cong S^3$ as the holonomy, and the Pontryagin normalization consumes its own $8\pi^2$ so that no stray transcendental leaks downstream. Step 4 is where the π is born: the matter charge lives on the Hopf circle threading the 3-sphere, and under the frozen stationary measure each sector's share is a ratio of volumes — 2π of fiber against $2\pi^2$ of holonomy. Step 5 closes the exits: the rival fractions die of geometry (no manifold in the construction carries their volumes), not of disagreement with data. The one thing the figure cannot draw is the measure lemma itself — *that* the frozen measure weights exactly fiber against holonomy is the conditional the D^* grade rests on, stated above and tracked in the promotion roadmap (App. N.5).

What the picture actually shows deserves a paragraph of its own, because the fibration is the one geometric fact this paper 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_{eta} of the bulk metric (7) embedded canonically in the $SU(2) \cong S^3$ the integer ledger selects.

The intuition for the ratio 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. And the three numbers are rigid together: a larger matter share would need a longer fiber, but a Hopf circle's length is fixed by the fibration; a smaller dark-energy share would need the fiber to cover more of the sphere, but a circle covers exactly what a circle covers; and the baryon line threads one unit of charge per revolution, no more. What in Λ CDM are three independent dials is here one shape, read three ways.

23 From the law to the curve: $H(t)$ in six steps

The two halves are now in hand: the Friedmann law of §21, a theorem with an empty right-hand side, and the π -locked inventory of §22, the derived densities that fill it. This section feeds the one into the other and builds the expansion history explicitly, in six steps.

Step 1 — normalize. Define the critical density $\rho_c = 3H_0^2/8\pi G_N$ and the fractions $\Omega_i = \rho_i(t_0)/\rho_c$. Dividing Eq. (14) by H_0^2 , with $k = 0$ inherited and the dilution law for each sector, the Friedmann equation becomes the working form used throughout this paper:

$$H(t)^2 = \frac{8\pi G_N}{3} \sum_i \rho_i(t) = H_0^2 [\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_\Lambda]. \quad (17)$$

Step 2 — insert the inventory. Here the derived fractions enter: $\Omega_m = 1/\pi$ and $\Omega_\Lambda = 1 - 1/\pi$ (§22), with $w = -1$ — itself derived, not assumed — keeping ρ_Λ constant. The radiation term is not free either: it is fixed by the measured photon temperature $T_\gamma = 2.7255$ K (Fixsen, 2009) plus the three neutrino species, giving $\Omega_r = 9.32 \times 10^{-5}$ at the carried rate. After this step *nothing* on the right-hand side of Eq. (17) is adjustable.

Aside — which Ω is which. A notational subtlety deserves one explicit paragraph here, because it is the source of a natural confusion. The symbols Ω_r , Ω_m , Ω_Λ in Eq. (17) are *constants by definition*: each is the fraction evaluated once, at the present epoch, $\Omega_i \equiv \rho_i(t_0)/\rho_c(t_0)$. The time dependence of the budget is carried entirely by the dilution factors standing beside them — a^{-4} ,




Figure 9: The Hopf fibration behind Theorem 22.1. *Left:* the 3-sphere of the electroweak holonomy, drawn by stereographic projection: the Hopf map $S^1 \hookrightarrow S^3 \rightarrow S^2$ 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 $U(1)_B$ charge per fiber revolution, $1/2\pi^2$. One fibration fixes three coupled numbers; none can be moved without the others.

a^{-3} , and 1 are precisely the solutions $\rho_i(t)/\rho_i(t_0)$ of the continuity equation for $w = 1/3, 0, -1$. The *running* fractions are a different object, $\Omega_i(t) \equiv \rho_i(t)/\rho_c(t)$, and they evolve dramatically: $\Omega_r(t) \rightarrow 1$ deep in the radiation era, $\Omega_m(t) \rightarrow 1$ in the matter era, $\Omega_\Lambda(t) \rightarrow 1$ on the de Sitter floor. Indeed Eq. (17) itself says the budget at the corridor switch was $\Omega_m(z_*) \approx 0.76$, $\Omega_r(z_*) \approx 0.24$, and $\Omega_\Lambda(z_*) \approx 2 \times 10^{-9}$ — matter and radiation nearly balanced, dark energy utterly negligible. Nobody inserts $1/\pi$ at early times; Figure 10 draws the running budget across the whole history, and it is nothing like today's.

The π -locked values are therefore a *boundary condition at the present epoch*, inserted once and propagated by the law to every other time — not a claim that matter holds a $1/\pi$ share always, which would contradict the continuity equation. The theorem of §22 is calibrated to exactly this structure: the frozen measure fixes the share *at the crossing epoch*, and the declared observing epoch is the moment the running $\Omega_m(t)$ crosses $1/\pi$ — 75 Myr before the present (Landmark 7 of §25); the why-now window of Part IV (§36) is the price of that declaration, stated rather than hidden.

Two corollaries keep the later chapters honest. First, early-universe physics never consults the present-epoch fractions: everything at $z \approx 1100$ — the baryon loading, the equality epoch, the sound horizon — depends on the *physical* densities ω_b and ω_m , which redshift deterministically as a^{-3} , and the Λ term is one part in 10^9 there; the CMB determination of §24 uses the fractions only in the present-day bookkeeping identity $\omega_i = \Omega_i h^2$. Second, the ratio $\Omega_b/\Omega_m = 1/2\pi$ is *epoch-independent* — both sectors dilute as a^{-3} , so the ratio at recombination equals the ratio today — which is why

the baryon-loading test of §26.1 may be applied directly to the $z \approx 1100$ sky, while $1/\pi$ and $1 - 1/\pi$ remain strictly present-epoch statements. (Objection O-Hubble-48 restates this aside against its sharpest form.)

Step 3 — read the eras. Each term dominates in turn, and each dominance is a power law: radiation gives $a \propto t^{1/2}$ and $H = 1/2t$; matter gives $a \propto t^{2/3}$ and $H = 2/3t$; the vacuum gives H constant and a exponential. *Intuition:* $1/H$ is the instantaneous e-folding time of the universe; in the decelerating eras that time keeps lengthening as the sources dilute, and in the vacuum era it locks.

Step 4 — the closed form. From matter–radiation equality onward the two-component equation integrates in closed form. With $\tau \equiv \frac{3}{2}\sqrt{\Omega_\Lambda} H_0 t$,

$$a(t) = \left(\frac{\Omega_m}{\Omega_\Lambda}\right)^{1/3} \sinh^{2/3} \tau, \quad \boxed{H(t) = \sqrt{\Omega_\Lambda} H_0 \coth \tau} \quad (18)$$

— one formula carrying the whole late history. Its two limits are the two eras: for small τ , $\coth \tau \rightarrow 1/\tau$ and $H \rightarrow 2/3t$, the matter era; for large τ , $\coth \tau \rightarrow 1$ and H settles on the de Sitter floor

$$H_\infty = \sqrt{\Omega_\Lambda} H_0 = \sqrt{1 - 1/\pi} H_0 = 55.2 \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (19)$$

The curve itself — the time dependence beneath every dilution factor of Eq. (17) — is drawn as its own object in Figure 11, with the three era slopes visible and the inflection marked.


Step 5 — the reading today. The present is one point on the curve: $H_0 \equiv H(t_0)$, the number the geometry fixes at 66.8 ± 0.2 by Route I of §24 and confirms at 66.6 ± 0.2 by Route II. The same closed form gives the age relation, $H_0 t_0 = \frac{2}{3\sqrt{\Omega_\Lambda}} \operatorname{arcsinh} \sqrt{\Omega_\Lambda/\Omega_m} = 0.948$, hence $t_0 = 13.88$ Gyr (App. C.8) — an output, not an input.

Step 6 — monotonicity, and what acceleration is not. Differentiating Eq. (17) gives $\dot{H} = -\frac{3}{2}(1 + w_{\text{eff}})H^2 \leq 0$ for any mixture with $w_{\text{eff}} \geq -1$: the rate falls monotonically for all time and *never rises* — not even after the expansion begins to accelerate, because $H = \dot{a}/a$ keeps decreasing while $\ddot{a} > 0$. Dark energy stops the fall at the floor H_∞ ; it does not reverse it. The familiar intuition that dark energy “speeds the expansion up” is true of the scale factor and false of the rate, and the difference is exactly what the phrase “Hubble constant” obscures.

The result of the six steps is a single curve — the entire expansion history, from the radiation era through the matter era to the de Sitter terminus, with no continuous parameter adjusted. It is drawn in Figure 13 (placed with its landmark walk in §25, page 74), and it is the central object of this paper. In the language of GDOS, $H(t)$ is a *registered observable* of the open-system projection, and every determination is an operational reading of it; the “tension” is a statement about the calibration of one of those readings, not about the derived history, which is single-valued. This is the viewpoint the rest of Part III makes quantitative and Part IV reconciles.

Scope — where this curve begins. Equations (12)–(19) describe the classical eras — radiation, matter, vacuum — and that is everything the argument of this paper needs: nothing here reaches earlier than the radiation era except through two measured survivors, the relic temperature fixing Ω_r and the CMB’s physical densities. The history before that — the inflationary launch, reheating, and the lock crossing that populates the ledger — is derived and drawn in the parent GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos), and deliberately not reproduced: this paper is self-sufficient for what it claims.

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 $\rightarrow \Lambda$ handover, drawn gradually
scale factor $a = 1/(1+z)$




Figure 10: The running budget $\Omega_i(t)$, shown twice — the evolution the present-epoch constants seed. The four fractions always sum to one; only the matter: Λ seam moves, the internal matter ratios frozen for all time ($\Omega_b/\Omega_m = 1/2\pi$, §26.1). *Top:* logarithmic cosmic time, radiation handing over to matter at $z = 3400$. *Bottom:* linear time, with the π -locked inventory as the snapshot at the $\Omega_m = 1/\pi$ crossing, 75 Myr before the present (inset; today $\Omega_m = 0.315$, inside the ± 160 Myr blur). Adapted from the parent GG-A10-Cosmos at the derived rate; generated by `c2_fig_omega_budget.py`, released with the paper.

24 Geometrical determination of H_0 : geometry first, the sky second

The background expansion rate is not an input to be bracketed against the data; it is an *output* of the same six-dimensional geometry that fixes the Standard-Model couplings. This section derives it twice, and the order of the two derivations is the point. *Route I* is geometry alone: the GU-5 conversion factors (§20.1), the theorem-grade scale chain, the vacuum-energy seesaw, and one codebook-selected ledger monomial — no telescope, no microwave sky, no ladder, nothing measured off any astronomical object. *Route II* is the precision confrontation: the sky’s two h -independent physical densities, divided by the derived fractions, land on the same rate — and a variant of Route II using only Big-Bang nucleosynthesis removes the microwave sky from the confirmation as well. The derivation is imported from the cosmology paper GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) and sharpened here; it is stated self-contained and in the register of a working cosmologist. Everything below is GG-6/GDOS/GRIP output — the stage of §20, the ledger, and the bridge — and nothing else.

24.1 Route I — the rate from the geometry alone: the vacuum-energy seesaw

The dark-energy sector delivers H_0 first, and delivers it with no astronomical input at all: every symbol in the route below is a theorem of the scale chain, a GU-5 conversion, a derived fraction, or a codebook-selected ledger monomial. It is worth walking the route factor by factor, because every step can be checked on a calculator. The first Friedmann equation at the present epoch, with the flat late-time budget, reads $H_0^2 = \rho_\Lambda / (3 \Omega_\Lambda M_{\text{Pl,red}}^2)$; and the geometric seesaw of the cohort fixes the vacuum density itself, $\rho_\Lambda = \zeta_\Lambda M_{\text{lock}}^8 / M_{\text{Pl,red}}^4$, with ζ_Λ a single dimensionless coefficient of the boundary projection. Substituting one into the other gives the closed form

$$H_0 = \frac{1}{M_{\text{Pl,red}}} \left(\frac{M_{\text{lock}}^2}{M_{\text{Pl,red}}} \right)^2 \sqrt{\frac{\zeta_\Lambda}{3 \Omega_\Lambda}} = 498.3 \sqrt{\zeta_\Lambda} \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (20)$$

The chain, number by number. The lock scale is a theorem of the scale chain (GG-A5-GDOS T8.7): $M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K} = \frac{1}{2} \times 1.221 \times 10^{19} \text{ GeV} \times e^{-35} = 3.85 \text{ TeV}$ — the same number on which the collider companion GG-P1-KK stakes its falsifiers. The first seesaw rung is $M_{\text{lock}}^2 / M_{\text{Pl,red}} = (3.85 \text{ TeV})^2 / 2.435 \times 10^{18} \text{ GeV} = 6.08 \times 10^{-12} \text{ GeV} = 6.08 \text{ meV}$ — already the right neighborhood for a vacuum-energy scale, out of two inputs fixed nineteen and three decades up. The second rung divides by the Planck mass once more: $M_{\text{lock}}^4 / M_{\text{Pl,red}}^3 = 1.52 \times 10^{-41} \text{ GeV} = 712 \text{ km s}^{-1} \text{ Mpc}^{-1}$ (converting by E/h and the megaparsec — the GU-5 bridge at work, and the only place any unit enters). The inventory supplies the last factor, $\sqrt{1/(3\Omega_\Lambda)} = 0.699$ at $\Omega_\Lambda = 1 - 1/\pi$, giving the coefficient carried cohort-wide as 498.3 (our re-evaluation with PDG constants gives 498.1, a 0.04% convention difference in the Planck-mass digits, invisible against every error bar in play). Note what the chain has already done: it has produced the 10^{-61} of Eq. (11) from $\frac{1}{2}e^{-35}$, one seesaw, and π — one hundred and nineteen of the famous one hundred and twenty orders of magnitude, before the last coefficient is even discussed.

The lock scale is pinned in the laboratory. Because $H_0 \propto M_{\text{lock}}^4$, the route’s skeptic will ask what a percent-level threshold correction to the lock scale would do: move the rate by several percent.




Figure 11: The derived scale factor $a(t)$ — the time dependence beneath every dilution factor. The solution of Eq. (17) with the π -locked inventory, drawn as its own object. *Top:* nineteen decades on log-log axes: the radiation-era slope $a \propto t^{1/2}$, the matter-era slope $a \propto t^{2/3}$, and the de Sitter upturn $a \propto e^{H_\infty t}$, with equality ($z = 3416$), the corridor switch ($z_* = e^7$), and today ($a = 1$, $t_0 = 13.88$ Gyr) marked. *Bottom:* the late universe in linear time: the closed form of Eq. (18), the inflection where deceleration turns to acceleration ($z = 0.62$) — the one moment $\ddot{a}$ changes sign, while H keeps falling — and the approach to the pure exponential, doubling every $\ln 2/H_\infty = 12.3$ Gyr. Generated by `c2_fig_a_of_t.py`, released with the paper.

But the fear is bounded *without cosmology*: the same lock scale predicts the electroweak vacuum expectation value, $v = M_{\text{lock}} (2/5)^3 = 246.33$ GeV against the measured 246.22 — a laboratory pin (through G_F) at +0.045%, which $H_0 \propto M_{\text{lock}}^4$ maps to $\pm 0.18\%$ on the rate, $\pm 0.12 \text{ km s}^{-1} \text{ Mpc}^{-1}$. The chain is therefore stiff at every rung but one, and everything now rests on the single dimensionless coefficient ζ_Λ .

What ζ_Λ is allowed to be. The coefficient is not 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_θ^1 , on which exactly three objects carry names: the seam factor 2, the holonomy angle $2\pi/L$, and the dual-lock count K . The enumeration is carried out on that alphabet, under a residue rule stated before the table is read (GG-A10-Cosmos, its released `zeta_table_v2.py`): the angle enters as an amplitude at a non-negative power, the seam factor only as the orbifold halving, and K only as the closure suppression. Fourteen monomials at cost ≤ 5 land in $40\text{--}100 \text{ km s}^{-1} \text{ Mpc}^{-1}$; five are admitted by the rule and nine excluded with their reason printed. The admitted cluster, through Eq. (20), is

$$H_0|_{\text{admitted, cost} \leq 5} \in \{59.6, 66.8, 74.8, 84.2, 94.4\} \text{ km s}^{-1} \text{ Mpc}^{-1}, \quad (21)$$

and read that way the result is two-sided and weaker than an earlier version of this paper printed: exactly one admitted monomial at cost ≤ 4 lands in the observed decade, $\pi/(LK) \rightarrow 66.8$, with zero observational input of any kind anywhere in the chain; the first admitted alternative, $(2\pi/L)^2/(2K)$, costs one unit more 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 ($\zeta_\Lambda = 0.0218$). The earlier four-row cluster “66.8–71.2”, drawn from the unrestricted alphabet $\{2, N, L, J, \pi, e, K\}$, did not show what it was said to show — on that alphabet forty-nine monomials at cost ≤ 4 land between 60 and 80, including $e/L^3 \rightarrow 73.5$ — and the sentence “pure geometry refuses the high reading” is withdrawn with it. What survives is true and is enough: the geometry names one number in the decade the sky reads, and the ladder’s number is not on the list. Which member the codebook’s naming rules select, as the second filter, is the subject of the next subsection.

The meV form — a unit identity, not a second test. The same chain can be read at the quarter power: $\rho_\Lambda^{1/4} = \zeta_\Lambda^{1/4} M_{\text{lock}}^2/M_{\text{Pl,red}} = 0.366 \times 6.08 \text{ meV} = 2.23 \text{ meV}$, against the familiar $(2.24 \pm 0.02) \text{ meV}$. We bill this honestly: the “observed” meV scale is $(3\Omega_\Lambda)^{1/4} \sqrt{\bar{M}_{\text{Pl}} H_0}$ — arithmetic on the same H_0 and Ω_Λ the seesaw already targets — so this paragraph is the rate agreement converted to different units, printed as bookkeeping, and counted *once*. What the quarter-power form does add is the framing of the “ 10^{120} problem” as two powers of a derived hierarchy — and one silent premise, now given its own open-question entry (OQ-Hubble-9): that the Standard-Model vacuum loops below M_{lock} do not re-source ρ_Λ on the boundary. The seesaw states what the vacuum energy *is*; the statement of why the loop contributions do not gravitate is a projection-sourcing property the GDOS/OGGEP layer owes and has not yet delivered as a theorem.

24.2 The selection of ζ_Λ : by the codebook, not by the data

A hostile reader will now expect the trick: with five admitted cluster members spanning 59.6–94.4, and the cheapest candidates of the unrestricted alphabet crowding the observed decade, surely the paper picks the one the sky prefers and calls it a derivation. That is precisely what this subsection does *not* do, and the discipline is not invented here: it is the published selection machinery of GG-A9-SM, applied without modification.

GG-A9-SM selects its twenty-eight Standard-Model parameters by a four-step procedure (its MSMF/OGN selection principle): a candidate space of rational monomials over a *structural base set on the parameter's boundary stratum*; an admissibility filter — every constituent must have a *named structural role on that stratum*, with π available only on strata containing the compact circle S_θ^1 and e only on radial strata; an empirical-precision filter; and the cost minimum. The decisive discipline is the admissibility filter, and GG-A9-SM has already used it against its own arithmetic twice, in exactly our situation. For $\sin^2 \theta_{12}^{\text{PMNS}}$ the bounded-cost search returned the *cheaper* orphan $1/\pi$ — and GG-A9-SM rejected it as “geometrically isolated,” adopting the costlier structural form $L^2/3N^3$. For $\sin^2 \theta_{23}^{\text{PMNS}}$ a numerical tie was broken by the *structural-naming rule*: the winner must be nameable as a mechanism. Cost never selects; structure does.

Apply that rulebook here. The stratum is fixed by what ζ_Λ *is*: the residue of the Wilson-line (Hosotani) zero mode on S_θ^1 — GG-A10-Cosmos identifies it as “the same kind of object as the strong- CP residue $\bar{\theta}$ of GG-A9-SM, a geometric closure-index suppression, not a radiative integral.” So ζ_Λ lives on the Wilson-line stratum, the same stratum as $\bar{\theta}$, whose established grammar prices exactly three kinds of object: the L -holonomy angle $2\pi/L$ (the sine argument of $\bar{\theta}$ itself — “the smallest compact-angle phase”), the orbifold/dyadic seam factor 2 (the $\frac{1}{2}$ of $M_{\text{lock}} = \frac{1}{2}E_{\text{Pl}}e^{-K}$, the $\frac{1}{2}$ of $\Omega_b = 1/2\pi^2$), and the dual-lock count K (the e^{-K} suppression; the Kaxion’s cap-loading ladder $K + \Sigma$). Now walk the cluster:

The surviving monomial is not merely admissible; it has the mechanism name the structural-naming rule demands, composed of three objects that are each already load-bearing elsewhere in the cohort:

$$\zeta_\Lambda = \frac{\pi}{LK} = \underbrace{\frac{1}{2}}_{\text{orbifold half } (M_{\text{lock}}, \Omega_b)} \times \underbrace{\frac{2\pi}{L}}_{\text{holonomy defect } (\bar{\theta}'\text{'s angle)}} \times \underbrace{\frac{1}{K}}_{\text{dual-lock share } (\text{Kaxion's cap ladder})} = \frac{\pi}{175} = 0.017952 : \quad (22)$$

the undrained vacuum residue of the lock episode is one orbifold-halved L -holonomy defect, shared over the K winding cells of the dual lock. Nothing new is introduced — the $\frac{1}{2}$ that prices the lock scale, the angle that prices $\bar{\theta}$, the cell count that prices the Kaxion’s cap ladder. Two clauses complete the rulebook, both implicit in GG-A9-SM’s structural-naming requirement and made explicit here because the audit asked. *Single occupancy*: a mechanism name attaches to one geometric feature, so each named object enters a candidate at most once — a squared holonomy defect or a tripled seam ($\frac{1}{2}(2\pi/L)^2/K$, $2^{-3}/J$) is a different physical claim with no name on the stratum, and is inadmissible however cheap. *No empirical filter*: of GG-A9-SM’s four selection steps, the empirical-precision filter was *not consulted* for ζ_Λ — the cluster and the survivor were fixed by admissibility and naming alone, before any comparison with the sky, which is what “pre-registered” means at F-Hubble-11

Table 4: The codebook selection of ζ_Λ , applied *before* any comparison with data: the second filter, GG-A9-SM’s admissibility and structural-naming rules on the Wilson-line stratum, applied to the cheapest candidates of the *unrestricted* alphabet (the first filter — the stratum alphabet and the residue rule of GG-A10-Cosmos — excludes four of the six before this table is drawn, and the two filters agree on every row). Exactly one candidate survives. The rejections are made by rules written down elsewhere and for other purposes, not by whether the survivor lands on a useful number. Two of the six candidates were raised by the 2026 audit rather than by me, and one of those two lands on the right value and is rejected anyway. That is the selection working blind, and it is the only reason the survivor counts for anything.

Candidate	Verdict	Reason (the GG-A9-SM rulebook)
$2/(NK)$	inadmissible	N ’s named roles (the triplicity of the fermionic boundary replication) exist only on matter strata; N has no named role on the Wilson-line stratum.
$1/J^2$	orphan	J ’s named roles are radial ($z_* = e^J$) and dyadic-depth (the 2^{-J} seam price inside $\bar{\theta}$); a <i>bare</i> J^{-2} has no angular-stratum name — J enters this stratum only through the composite $K = LJ$. The same class as the rejected $1/\pi$ orphan of θ_{12} , and likewise cheaper than the winner.
$2/(K\pi)$	orphan	π ’s only named role on this stratum is inside the composite holonomy angle $2\pi/L$ — exactly as J enters only through $K = LJ$. Decomposed over the priced objects $\{\frac{1}{2}, 2\pi/L, 1/K\}$, the monomial $2/(K\pi)$ demands an <i>inverse</i> defect $(2\pi/L)^{-1}$, which names no mechanism; a bare π downstairs is the same class as the rejected $1/\pi$ orphan of θ_{12} . The physical reading agrees: a residue holds its angle <i>upstairs</i> ($\bar{\theta}$ ’s $\sin(2\pi/L)$), while π downstairs is the cohort’s <i>share</i> form ($\Omega_m = 1/\pi$), and a vacuum residue is not a share of a budget.
$\pi/(LK)$	selected	Decomposes entirely into named stratum objects; see Eq. (22).
$1/(2L^2)$	inadmissible	Bare L has no named role on the Wilson-line stratum outside the composite holonomy angle $2\pi/L$ (the same rule that orphans $2/(K\pi)$); raised by the 2026 audit and priced here.
$2^{-3}/J = 1/56$	inadmissible	J ’s named roles are radial, not Wilson-line (as for $1/J^2$); and a <i>repeated</i> dyadic seam (2^{-3}) has no single mechanism name — see the occupancy rule below. Raised by the 2026 audit; lands on 66.6 and is <i>still</i> rejected, which is the selection working blind.

(and why the table above can afford to reject a candidate that lands exactly on 66.6). And one commitment, so that the grammar is law rather than jurisprudence: the rulebook as codified here — admissibility by named stratum role, single occupancy, no empirical filter — is *frozen*, and any future selection anywhere in the cohort must use it unchanged; a clause added after the next inconvenient candidate would itself be a falsification of the method. (Single occupancy forbids *unnamed* repetition; ledger-indexed towers such as 2^J or e^K are one named object each — a depth- J or depth- K mechanism — which is why the frozen rule convicts $2^{-3}/J$, whose exponent has no stratum role, and acquits every standing selection of GG-A9-SM.) With the selection made, Route I closes:

$$H_0^{\text{geom}} = 498.3 \sqrt{\pi/175} = 66.8 \pm 0.2 \text{ (structural) km s}^{-1} \text{ Mpc}^{-1}, \quad (23)$$

where the structural error carries the laboratory pin on the lock scale and the conditional status of the selection lemma. Every factor is GG-6/GDOS/GRIP; the sky has not yet been consulted.

Three honesty clauses, stated before the data are allowed in. *First, the grade.* This selection yields Derived (cond., MSMF/OGN) — conditional on one named lemma, the *stratum assignment* (ζ_Λ is the Wilson-line zero-mode residue), which GG-A10-Cosmos carries as its flagged physical reading of the coefficient. That is exactly the grade structure of GG-A9-SM’s PMNS entries; Derived (theorem) waits on the Hosotani computation itself (GG-A5-GDOS O6). *Second, the cost accounting.* Under GG-A9-SM’s closure rule for K ($K = LJ$ is not an independent atom, so a K -factor costs two), the “cost-3” cluster entries are cost four and the orphan $1/J^2$ stays cost two — i.e. cost favors the orphan *more* strongly, and the selection above leans on admissibility and naming alone, exactly as the θ_{12} precedent requires. *Third, the interpretive step, named.* The three rejections are one rule applied three times — a constituent with no bare name on the stratum: N nowhere, J only through the composite $K = LJ$, π only through the composite $2\pi/L$ — so no candidate needed a rule invented for it. What remains interpretive is the enumeration of the stratum’s priced objects $\{\frac{1}{2}, 2\pi/L, 1/K\}$, read off its three cohort precedents ($M_{\text{lock}}, \bar{\theta}$, the Kaxion splitting); that reading is carried as the second clause of the selection lemma, not as a theorem. All three clauses are cashed by one pre-registered number: the Hosotani computation, when completed, must return $\zeta_\Lambda = \pi/175$ to within $\pm 0.6\%$, the tolerance the rate’s own ± 0.2 implies — registered as falsifier F-Hubble-11 (§56), two-sided by construction. The target is printed before the lemma; it cannot be renegotiated after it.

24.3 The closed form: the expansion rate from $\{c, G_N, \hbar\}$ and the ledger alone

Route I was walked factor by factor above, and its coefficient was selected by the codebook; it is worth collapsing the whole route into a single analytic expression, because the claim of this section — geometry first, the sky second — is nowhere more visible than in one equation whose every symbol can be pointed at. Four rungs enter. The only dimensionful objects anywhere are the first three GU-5 constants, through the Planck mass $M_{\text{Pl}} = \sqrt{\hbar c/G_N}$ and its reduced companion $M_{\text{Pl,red}} = M_{\text{Pl}}/\sqrt{8\pi}$. The lock scale is the theorem $M_{\text{lock}} = \frac{1}{2}M_{\text{Pl}} e^{-K}$ (GG-A5-GDOS T8.7), which in reduced units is the pure statement $M_{\text{lock}} = \sqrt{2\pi} M_{\text{Pl,red}} e^{-K}$. The vacuum residue is the seesaw dressed by the selected monomial, $\rho_\Lambda = \zeta_\Lambda (M_{\text{lock}}^2/M_{\text{Pl,red}})^4$ with $\zeta_\Lambda = \pi/(LK)$. And the present-epoch Friedmann

equation on the derived inventory divides by $3\Omega_\Lambda = 3(1 - 1/\pi)$. Substituting each rung into the next, every power of the Planck mass cancels except a single overall rate, the transcendentals collect into $2\pi^5/(\pi - 1)$, and Eq. (20) becomes

$$H_0 = e^{-4K} \sqrt{\frac{2\pi^5 c^5}{3LK(\pi - 1)\hbar G_N}} = 66.7 \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (24)$$

Evaluated with CODATA constants the right-hand side gives 66.74 — the $498.1\sqrt{\zeta_\Lambda}$ of the re-evaluated chain, now with no intermediate scale quoted at all. The provenance of every symbol can be stated in one sentence: c , G_N , and $\hbar$ carry the dimensions (they are the GU-5 bridge, and $\sqrt{c^5/\hbar G_N}$ is nothing but the inverse Planck time); 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 $v = M_{\text{lock}}(2/5)^3$ that stiffens the chain; π enters three ways (Einstein’s 8π , the monomial’s numerator, and the inventory’s $1 - 1/\pi$); e is the scale chain’s exponential; and the structural rationals $\frac{1}{2}$ (the lock theorem) and 3 (Friedmann) are geometry’s bookkeeping, not ledger entries. There is nothing else in the formula. In particular there is no astronomical quantity of any kind: 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 of the closed form as corollaries, and they are worth one paragraph each of plain statement. First, the age of the universe in Planck units. The formula gives directly $H_0 t_{\text{Pl}} = \sqrt{2\pi^5/[3LK(\pi - 1)]} e^{-4K} = 0.738 \times e^{-140} = 10^{-60.9}$: the famous sixty-one orders of magnitude between the Planck rate and today’s expansion — the dimensionless task of Eq. (11) — are $4K = 140$ e-folds of the derived descent ($4K/\ln 10 = 60.80$ orders), dressed by a geometric factor of order unity. The question “why is the universe so old?” receives, on this reading, a one-integer answer: because $K = 35$, entering the vacuum sector at fourth power. Second, the cosmological-constant problem itself. The same rungs give $\rho_\Lambda/\rho_{\text{Pl,red}} = \zeta_\Lambda (2\pi)^4 e^{-8K} = 10^{-120.2}$: the most notorious mismatch in physics — one hundred and twenty orders of magnitude — is $8K$ e-folds (121.60 orders) corrected by 1.4 orders of pure geometry. What the literature treats as a catastrophic cancellation to be explained is, in the closed form, two powers of one derived hierarchy — the statement made in prose below Eq. (20), here exhibited as arithmetic.

One honesty note completes the subsection, in the register the codebook requires. Of the five OGN-5 integers, the closed form contains K and L explicitly; $M = 1$, $N = 3$, and $J = 7$ do not appear in the final expression, and no claim is made that they do. They act upstream — selecting the fibration and the anomaly closure that make T8.7 and the π -locked inventory theorems in the first place — and J surfaces elsewhere in the history, in the switch redshift $z_* = e^7$ and in the rejected cluster member $1/J^2$ (§24.2). And the equation is exactly as conditional as its rungs: the inventory carries its D^* grade, and the monomial carries the selection lemma, cashed by the pre-registered Hosotani target of F-Hubble-11. Nothing about writing the route as one line strengthens its grade; what it strengthens is its *exposure* — a reader with a calculator, Eq. (24), and five constants can now check the central claim of this paper in under a minute, and a referee who wishes to attack it knows exactly which five symbols to aim at.

24.4 Route II — the precision confirmation: the microwave sky reads the derived rate off two densities

The acoustic peaks of the CMB do not measure H_0 . What they measure, exquisitely, is two *physical* densities (Planck Collaboration, 2020), each independent of the expansion rate:

$$\omega_b \equiv \Omega_b h^2 = 0.02237 \pm 0.00015, \quad \omega_m \equiv \Omega_m h^2 = 0.1430 \pm 0.0011. \quad (25)$$

The physics behind each is worth one sentence, because the two are read off *different* features of the sky. The baryon density ω_b comes from *baryon loading*: baryons add inertia to the oscillating photon–baryon plasma without adding pressure, so the compression peaks (first, third) are enhanced relative to the rarefaction peaks (second) — the second-to-first peak-height ratio is a clean baryometer, and the primordial deuterium abundance reads the same number through entirely different physics. The matter density ω_m comes from the *epoch of matter–radiation equality*: modes that enter the horizon during radiation domination have their gravitational potentials decay under them, which drives the oscillations and boosts the peak-height envelope; the shape of that envelope reads z_{eq} , and hence ω_m .⁴

Why neither number carries h is the pivot of the confirmation. The convention $\omega_i \equiv \Omega_i h^2$ encodes a physical fact: ω_i is an *absolute* density (an energy per volume, in fixed units), while Ω_i is a fraction of the critical density — and the critical density itself scales as h^2 . The plasma physics at $z \approx 1100$ responds to the actual densities present, not to fractions defined against today’s expansion rate; that is why the microwave sky can pin ω_b and ω_m to a fraction of a percent while knowing nothing of h . This is the celebrated *geometric degeneracy*: the primary anisotropies alone cannot measure the present rate at all. Standard Λ CDM breaks the degeneracy with a fitted late-time model; GG-Theory breaks it with theorems — and, after Route I, does not even need to break it to *know* the rate. What follows is a confirmation, not a derivation.

The inversion is then one line. Standard practice runs the six-parameter fit: vary all six dials, including h , until the whole pattern matches — H_0 emerges as an output of the fit, with the fractions along for the ride. GG-Theory runs the logic in the opposite direction: the fractions are theorems (§22), so dividing a *measured* physical density by a *predicted* fraction returns h with no freedom left,

$$\omega_i = \Omega_i h^2 \implies h = \sqrt{\frac{\omega_i}{\Omega_i}} = \begin{cases} \sqrt{0.02237/0.05066} = 0.665 & \text{(baryon anchor, } H_0 = 66.5), \\ \sqrt{0.1430/0.31831} = 0.670 & \text{(matter anchor, } H_0 = 67.0). \end{cases} \quad (26)$$

Both anchors land inside Route I’s structural band: what the geometry predicted, the sky now confirms with its full precision.

The agreement of the two lines is not one test written twice. The two densities are measured from different features of the spectrum — the peak ratio for ω_b , the envelope for ω_m — and the two predicted fractions come from one fibration (§22), rigidly coupled. If the fibration were wrong — if, say, the true baryon fraction sat ten percent above $1/2\pi^2$ — the baryon anchor would return an h

⁴ 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.

five percent below the matter anchor, and the two lines of Eq. (26) would disagree with *each other*, before any comparison with Planck’s fitted rate. Their agreement at 0.7% is therefore a consistency test of the inventory’s coupled structure, passed internally.

One variant of the inversion deserves its own paragraph, because it removes the microwave sky from the confirmation entirely. The baryon density is measured a second way — the primordial deuterium abundance, with the LUNA reaction rates — giving $\omega_b^{\text{BBN}} = 0.02233 \pm 0.00036$ with no CMB anisotropies anywhere in the chain; divided by the derived $1/2\pi^2$, it returns

$$H_0^{\text{BBN}} = 66.4 \pm 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (27)$$

Route II therefore does not even need the sky it confirms against: any h -independent absolute density will do, and every one available — two from the microwave sky, one from the first three minutes — lands on the derived rate.

The error budget is elementary and worth displaying once. Since $h \propto \sqrt{\omega_i}$, the fractional error halves: $\sigma_h/h = \frac{1}{2} \sigma_\omega/\omega$, giving 0.34% (± 0.23) on the baryon anchor and 0.38% (± 0.26) on the matter anchor; the conditional systematic of the shared measure lemma (§22) is common to both and does not enter their difference. The two-anchor consistency yields the canonical value — the adopted center of the route span, not a formal weighting; App. E.7 prints the full bookkeeping:

$$\boxed{H_0^{(0)}|_{\text{two-road}} = 66.6 \pm 0.2 \approx 67 \text{ km s}^{-1} \text{ Mpc}^{-1}.} \quad (28)$$

This confirmation is itself parameter-free — an inversion, not a fit. Standard Λ CDM *fits* Ω_m and Ω_b and reads off $H_0 = 67.4$; GG *predicts* the fractions, so h is forced — and it lands on the independently derived Route I value. A full Boltzmann solution completes the confrontation: fed GG’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 fit’s 220/536/813), the second-to-first peak-height ratio to 0.8%, the diffusion-damping scale to 0.4%. The inventory does not merely imply $H_0 \approx 67$ arithmetically; it fits the microwave sky at that value. The residual $\sim 1\%$ below Planck’s 67.4 is exactly GG’s exact π -values sitting 1–3% above Planck’s best-fit fractions — the *coherent π -offset*, not a new tension.

24.5 The concordance, and the canonical rate

The two routes now stand side by side, and Figure 12 draws them on one axis. Route I — geometry alone — gives 66.8 ± 0.2 (structural), the one admitted reading at cost ≤ 4 in the observed decade, with the first admitted alternative at 74.8. Route II — the sky’s h -independent densities against the derived fractions — gives 66.6 ± 0.2 , with its CMB-free variant at 66.4 ± 0.5 . The two central values differ by 0.25%, 0.6σ of their combined errors: the routes share no ingredient except the geometry itself — one runs through TeV-scale physics and the vacuum, the other through the plasma of the young universe — and they meet inside each other’s error bars. This paper carries the structural 66.8 ± 0.2 as its rate, as the parent GG-A10-Cosmos does; its companions carry the concordance value 66.6 ± 0.2 , the adopted center of the route span, and the 0.25% between the two is inside every comparison either makes. Where this paper prints 66.6 it is the two-road reading

The two routes to the present rate: geometry determines, the sky confirms
 Route I is pure GG-6/GDOS/GRIP (GU-5 conversion factors + the OGN-5 ledger); Route II divides the sky's h -independent densities by the derived fractions.




Figure 12: The two-route determination of the present rate. *Top lane:* Route I, geometry alone — the GU-5 bridge, the theorem-grade lock scale (laboratory-pinned through the electroweak vev at 0.045%), and the codebook-selected $\zeta_\Lambda = \pi/175$ (Eq. (22)) — with the admitted stratum monomials of GG-A10-Cosmos's residue rule as open diamonds (59.6, 66.8, 74.8, 84.2, 94.4), the selected $\pi/(LK)$ filled, and the ladder's 73.50 marked between the two admitted neighbors; no admitted monomial returns it. *Middle lane:* Route II, the confirmation — the baryon and matter anchors of Eq. (26), the CMB-free BBN anchor of Eq. (27), and the two-road concordance 66.6 ± 0.2 (the rate this paper carries is Route I's 66.8). *Bottom lane:* the two famous readings, for comparison. No astronomical observation enters the top lane; the sky enters only as the confirmation and the comparison. Generated by `c2_fig_two_routes.py`, released with the paper.

of the same background — the derived fractions read against the measured densities — and it is labeled so. Agreement of this kind cannot be arranged after the fact without breaking one of the routes' other commitments — the acoustic spectrum on one side, or the electroweak vev and the measured dark-energy scale on the other.

What would break the concordance is equally explicit: any of the three anchors of Route II leaving the band; the Hosotani computation returning a value off the pre-registered $\pi/175$ (F-Hubble-11); or a future inventory determination driving Ω_m firmly below $1/\pi$ (F-Hubble-4 territory). The rate of this paper is not one number defended once; it is one number exposed three ways.

24.6 Provenance: what is derived, what is conditional, what is input

Honest bookkeeping, in the grading of the cohort codebook, is collected in Table 5: the GU-5 bridge is the declared dimensional input; the lock scale is a theorem, pinned in the laboratory; the seesaw coefficient is Derived conditional on the selection lemma, its closure pre-registered; the π -inventory and the rates it forces are Derived-conditional under the Hopf-fiber measure lemma; the two CMB densities are a measurement that enters only the confirmation; and of the local excess nothing is derived — the void channel is priced and not carried, the systematic named and left open. The robust content is the concordance of the two routes at $H_0^{(0)} \approx 67$, with the geometric route standing on its own. The early-versus-late *split* of this one background rate — the Hubble tension itself — is the subject of the next section.

25 Landmarks on the derived $H(t)$ curve

A derived expansion history is not one prediction; it is a continuum of them. Because Eq. (17) fixes H at *every* epoch once the inventory is fixed, each physically meaningful moment on the curve — primordial nucleosynthesis, matter–radiation equality, recombination, the onset of acceleration, the present, the far future — is a place where the geometry has already committed to a number that observation can check. This section walks the curve from the first second to the infinite future, treating each landmark as a checkpoint: what happens there physically, what the derived chart says its epoch and value must be, and what is measured. Figure 13 draws the derived curve itself — the central object of this paper — and Figure 14 annotates the same curve with the walk: nineteen decades of cosmic time on one zero-parameter history. The two figures are a pair, and the pair is the heart of the paper: the first shows *what* the geometry forces, the second shows *where* observation has checked it. Table 6 compresses the walk; the derivation courses of the Technical Appendices (App. K.5–E.2) carry the arithmetic in full.

Landmark 1: primordial nucleosynthesis ($t \approx 1$ s–3 min). The earliest checkpoint the chart can face is the oldest quantitative test in cosmology. During the first minutes the expansion rate — here the radiation-era branch $H = 1/2t$ of the derived curve, Eq. (87) — races the weak interactions, and the surviving light-element abundances depend on one cosmological number, the baryon-to-photon ratio η_b . The derived chart fixes it: $\Omega_b = 1/2\pi^2$ at the carried rate gives $\omega_b = 0.02258$ and hence $\eta_b = 6.18 \times 10^{-10}$, against 6.13×10^{-10} from the measured ω_b (confirmed independently by the deuterium abundance): a +0.9% confrontation, carrying the by-now-familiar sign. The chart’s first checkpoint is passed at half a percent, thirteen decades of time before the epoch whose calibration the tension disputes.

Landmark 2: matter–radiation equality ($z_{\text{eq}} = 3416$, $t \approx 51$ kyr). Radiation dilutes as a^{-4} , matter as a^{-3} , so their crossing is fixed by the ratio of two shares, $1 + z_{\text{eq}} = \Omega_m/\Omega_r$. With the derived $\Omega_m = 1/\pi$ and the measured radiation share at the carried rate ($\Omega_r = 9.32 \times 10^{-5}$), the crossing lands at $z_{\text{eq}} = 3416$; Planck’s fitted value is 3387 ± 21 — a +0.9% agreement, same sign again. Equality matters twice over: it is where structure growth begins in earnest, and its epoch

Table 5: The honesty triage for the derived Hubble rate: every object in the two-route derivation, its grade in the cohort codebook, and its basis. This is the single canonical triage referenced throughout the paper; the grades match GG-A10-Cosmos’s, with the ζ_Λ grade carrying the codebook-selection lemma of §24.2, and the two rows GG-A10-Cosmos does not carry — the void channel and the ladder systematic — graded “priced, not carried” and “named, not derived” here.

Object	Grade	Basis
GU-5 bridge $\{c, G_N, \hbar\}$	Declared input	Unit-conversion factors of the 6D geometry (GG-A4-Euclid); natural units $\{1, 1, 1\}$; the only dimensional content anywhere.
Lock scale $M_{\text{lock}} = 3.85 \text{ TeV}$	Theorem	The scale chain, GG-A5-GDOS T8.7; laboratory-pinned by the electroweak vev at +0.045%.
Seesaw coefficient $\zeta_\Lambda = \pi/175$	Derived (cond., MSMF/OGN)	Selected by stratum availability and structural naming among the bounded-cost cluster (§24.2); conditional on the stratum-assignment lemma; the Hosotani closure (GG-A5-GDOS O6) pre-registered as F-Hubble-11.
π -inventory $\Omega_m, \Omega_b, \Omega_\Lambda$	D*	The Hopf-fiber identification, conditional on the measure lemma (Conjecture VIII) and the declared present epoch.
Geometric rate $H_0^{\text{geom}} = 66.8 \pm 0.2$	Derived (cond., MSMF/OGN)	Route I, Eq. (23): no astronomical input.
Two CMB densities ω_b, ω_m	Measurement	$\hbar$ -independent inputs to Route II (the confirmation) only; no role in Route I; the BBN anchor replaces even these.
Two-road rate (concordance) 66.6 ± 0.2	D*	Route II read against Route I; the same background through the measured densities; inherits the inventory grade; no additional freedom. The rate this paper carries is Route I’s 66.8.
Local-void channel ϵ_{Hole}	Priced, not carried	Linear response computed (sign unconditional, coefficient inventory-fixed); the required depth is not observed (§33; GG-A10-Cosmos).
Ladder systematic ϵ_{syst}	Named, not derived	The whole $\approx 10\%$ excess, read as a distance-ladder calibration systematic the framework predicts; bounded and left open.




Figure 13: The Hubble parameter $H(t)$ is a derived geometric function; H_0 is one reading of it. The GG-Theory background expansion rate computed from Eq. (17) 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 input is the GU-5 conversion G_N ; the rate itself is derived (§24). *Top:* the full history on log-log axes, radiation ($H \propto a^{-2} \propto t^{-1}$) $\rightarrow$ matter ($H \propto a^{-3/2} \propto t^{-1}$) $\rightarrow$ dark energy. *Bottom:* the late universe in linear time. The rate falls monotonically and asymptotes to the de Sitter floor $H_\infty = H_0 \sqrt{\Omega_\Lambda} = 55.2 \text{ km s}^{-1} \text{Mpc}^{-1}$ — never rising, even once expansion accelerates. The familiar $H_0 \equiv H(t_0) = 66.8$ is one point on this curve, at $t_0 = 13.88 \text{ Gyr}$; a single number does not capture the dynamical change of the Universe's size. The whole history is a GG-6/GDOS/GRIP output: fix the densities and their time dependence, and $H(t)$ follows.




Figure 14: Landmarks on the derived $H(t)$. Nineteen decades of cosmic time — from primordial nucleosynthesis to the de Sitter future — on the one zero-parameter curve the π -locked inventory forces. The numbered checkpoints are walked in §25 and tabulated in Table 6: BBN (1), matter–radiation equality (2), the corridor switch at $z_* = e^7$ where the sound horizon freezes (3), the acceleration onset (4), matter–vacuum equality (5), the $\Omega_m = 1/\pi$ crossing and today (6), and the de Sitter floor $H_\infty = 55.2$ (7). The rate falls monotonically through all of it; the eras change only the slope. Seven measured checkpoints, none deviating by more than a percent, all with the coherent sign of §26.4.

shapes the peak-height envelope through which ω_m is measured — the reason it is deliberately *not* double-counted as an independent test in Table 7.

Landmark 3: the corridor switch ($z_* = e^7 = 1096.6$, $t \approx 0.37$ Myr). The integer-ledger landmark: the GSSB transition of Prop. 31.1 (Part IV), whose redshift the OGN-5 integer $J = 7$ fixes with no continuous freedom, lands at 0.6% from the measured decoupling 1089.92 ± 0.25 — a distinct quantity: GG-A10-Cosmos’s atomic chain, evaluated at the framework’s own $h = 0.6676$, derives the decoupling itself at $z_{\text{dec}} = 1091.5$, $+0.15\%$ from the measured value and inside the 0.18% modeling systematic of the recombination calculation; the switch sits $\delta = 0.0038$ e-folds above it, 0.38% in $1 + z$. At the switch the derived curve reads $H(z_*) = 1.57 \times 10^6$ km s⁻¹ Mpc⁻¹ — the expansion was twenty-four thousand times faster than today — and everything to its past is read through the acoustic dictionary, everything to its future by the late probes on the present attractor.

It is also the moment the fourth landmark freezes.

Landmark 4: the sound horizon freezes ($r_s = 144.7$ Mpc). The pressure wave that has been crossing the plasma since the first instant stops at the switch, leaving a fossil ruler whose length the derived chart computes with no free parameter (App. E.2): 144.7 Mpc against the measured 144.43 ± 0.26 (+0.2%). Every early-universe “solution” of the tension must shorten this ruler; the derived reading, landing on it, is structurally forbidden from doing so — which is why its account of the tension lives on the late side (Part IV).

Landmark 5: the acceleration onset ($z = 0.62$, $t \approx 7.8$ Gyr). The deceleration parameter changes sign where $\rho_m = 2\rho_\Lambda$, i.e. at $z_{\text{acc}} = (2\Omega_\Lambda/\Omega_m)^{1/3} - 1 = 0.62$ — a pure function of π . Supernova reconstructions of the deceleration–acceleration transition place it at $z_t \approx 0.6$ – 0.7 , consistent. Note what the curve does *not* do there: $H(t)$ keeps falling. Acceleration is a statement about $\ddot{a}$, not about H ; the familiar intuition that dark energy “speeds the expansion up” is true of the scale factor and false of the rate, and the difference is exactly what the phrase “Hubble constant” obscures.

Landmark 6: matter–vacuum equality ($z = 0.29$, $t \approx 10.4$ Gyr). The vacuum share overtakes matter at $z_{\Lambda m} = (\Omega_\Lambda/\Omega_m)^{1/3} - 1 = 0.29$, against the concordance 0.30 ± 0.02 . From here the curve begins its long turn onto the floor.

Landmark 7: the crossing, and today (13.80 and 13.88 Gyr). The declared observing epoch — the one input the geometry does not remove (§22) — is the moment the matter share crosses exactly $1/\pi$: $t = 13.80$ Gyr, a mere 75 Myr before the present. One bookkeeping sentence, so no table can be accused of mixing epochs: the inventory rows of this paper are evaluated at the declared crossing epoch, and over the 75 Myr to the astronomers’ present the fractions evolve by -1.0% ($\Omega_m : 0.3183 \rightarrow 0.315$) while the rate falls by 0.24% — both inside the ± 160 Myr calibration blur already priced. We record, without yet claiming anything, an arithmetic coincidence: the evolved 0.315 is numerically the Planck-fitted matter fraction, i.e. the crossing-epoch offset has the same size and sign as the fitted-versus-exact Ω_m displacement — while the rate’s 0.24% drift is far too small to account for the -0.9% h offset, so the epoch declaration can at most explain one entry of the coherent pattern, not the pattern (OQ-Hubble-10). Today the curve reads $H_0 = 66.8$, and the age integral gives $t_0 = 0.948/H_0 = 13.88$ Gyr (App. C.8), against the Planck- Λ CDM 13.80 ± 0.02 : $+0.6\%$, one more shadow of the same displacement. The oldest stellar populations, at 13–13.5 Gyr, fit comfortably beneath it — as they would not beneath the age a local 73.5 implies (≈ 12.6 Gyr, §30).

Landmark 8: the de Sitter terminus ($H \rightarrow 55.2$). The curve’s future is as fixed as its past: with $w = -1$ exactly, $H(t)$ settles onto the floor $H_\infty = H_0\sqrt{\Omega_\Lambda} = 55.2$ km s $^{-1}$ Mpc $^{-1}$ and stays there forever, the universe thereafter doubling in size every $H_\infty^{-1} \ln 2 = 12.3$ Gyr. This is not a checkpoint that can be measured; it is the one the equation of state guards (F-Hubble-5), and the

quiet end of the story the rest of the paper argues about: one clock, still running, at a rate the geometry named before anyone could read it.

Read as a whole, the walk makes the Part's central point more vividly than any single number: seven of the eight landmarks carry a measured comparison, every one lands within a percent, and every deviation carries the sign the single +1% inventory displacement propagates (§26.4). A fitted model could match any one landmark by turning a dial; matching all of them with no dial is what it means for the history to be derived.

Table 6: The eight landmarks of the derived expansion history: epoch, the derived value and its origin, the measurement, and the deviation. Every quantitative deviation is below one percent and carries the coherent sign of §26.4. The eight are not eight independent successes. They share the inventory and they share the sign, so a fault in the inventory would move all of them together and in the same direction. What the table does show is that one set of integers carries a history from the first three minutes to the present with no free parameter anywhere along it, and that where it misses, it misses small and it misses the same way.

Landmark	Epoch	Derived	Measured	Dev
1. BBN	$t \approx 1 \text{ s} - 3 \text{ min}$	$\eta_b = 6.18 \times 10^{-10}$ $(\Omega_b = 1/2\pi^2)$	6.13×10^{-10}	+0.9%
2. Equality	$z = 3416, t \approx 51 \text{ kyr}$	$1 + z_{\text{eq}} = \Omega_m/\Omega_r$	3387 ± 21	+0.9%
3. Corridor switch	$z = 1096.6, t \approx 0.37 \text{ Myr}$	$z_* = e^J, J = 7$	$1089.92 \pm 0.25 \text{ (dec., two-qty)}$	+0.6%
4. Sound horizon	frozen at z_*	$r_s = 144.7 \text{ Mpc}$	$144.43 \pm 0.26 \text{ Mpc}$	+0.2%
5. Accel. onset	$z = 0.62, t \approx 7.8 \text{ Gyr}$	$(2\Omega_\Lambda/\Omega_m)^{1/3} \text{ SN: } z_t \approx 0.6-0.7$		consist.
6. Matter- Λ eq.	$z = 0.29, t \approx 10.4 \text{ Gyr}$	$(\Omega_\Lambda/\Omega_m)^{1/3} - 1$	-0.30 ± 0.02	consist.
7. Crossing / today	13.80 / 13.88 Gyr	$t_0 = 0.948/H_0$	13.80 ± 0.02	+0.6%
8. de Sitter floor	$t \rightarrow \infty$	$H_\infty = 55.2; \text{ doubling } 12.3 \text{ Gyr}$	guarded by $w = -1$ (F-Hubble-5)	—

26 The microwave background as an over-determined test

With the rate derived upstream — by geometry alone, then confirmed by the sky's own densities (§24) — the microwave background changes role: from the source of an input to the field of an over-determined test. Every observable in this section is a GG-6/GDOS/GRIP prediction, computed before the sky is consulted; the sky enters only in the comparison column. And the test is the strongest single argument of this Part. Standard Λ CDM runs on six free parameters $\{\omega_b, \omega_c, \theta_*, \tau, A_s, n_s\}$; GG-Theory makes the background half of them theorems — $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/2\pi^2$,

$w = -1$, and the corridor-switch redshift $z_* = e^7$ — so 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. What remains free is only the seed (A_s, n_s) and the optical depth τ . Crucially, the derived fractions do not sit at a random point in parameter space: they sit *largely along* Λ CDM’s celebrated geometric degeneracy direction (the Ω_m – H_0 line at fixed θ_*), which is why the same chart that returns H_0 also fixes the whole acoustic pattern — and the residue transverse to that direction is precisely why the chart is falsifiable rather than degeneracy-safe. The transverse residue deserves its sharpest lens, so we apply it to ourselves: the sky’s best-measured combination is $\Omega_m h^3$, and the derived chart gives $\Omega_m h^3 = (1/\pi)(0.6676)^3 = 0.0947$ against the published 0.09633 ± 0.00029 — -1.7% , a pull of ≈ -6 , the single sharpest miss in this paper. It is not a new failure mode: it is the -0.9% rate offset entering cubed, net of the $+1.0\%$ fraction offset — the same two displacements the chart already carries, compounded through the sky’s most precise projection. We print it here, in the section that owns the degeneracy argument, because a chart that prints its sharpest miss cannot be accused of curating its comparisons; and its break condition is the same one displacement carries everywhere else.

26.1 The h -free ratio: the baryon-loading test

The sharpest zero-parameter test needs no expansion rate at all. The relative heights of the acoustic peaks are set by the baryon-to-matter ratio, and in GG-Theory that ratio is a pure number, independent even of h :

$$\frac{\omega_b}{\omega_m} = \frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} = 0.15915, \quad \text{observed } 0.15643 \pm 0.0016 \quad (+1.7\%, 1.7\sigma). \quad (29)$$

There is no freedom here whatsoever — not the Hubble constant, not the seed amplitude — and the prediction lands within 1.7σ of the measured baryon loading, the observed value being the ratio of the two measured densities the core table prints ($0.02237/0.1430$); against the neutrino-free baryon-plus-cold-matter sum 0.1424 the same ratio reads 0.15709 and the deviation $+1.3\%$. (The matter–radiation equality redshift is deliberately not counted as a second, independent check, since through ω_m it would double-count the same density.)

26.2 The core observables, computed and confronted

Table 7 collects the observables the derived inventory forces — including, now, the two physical densities themselves, which the old six-parameter fit consumed as inputs and the derived rate returns as outputs — each with its GG closed form, its value, the Planck 2018 measurement, and the fractional deviation. Every entry lands within 1.7% , and — the point of the next subsection — every deviation carries the sign the one displacement propagates.

26.3 The full spectrum and the χ^2

The table is a set of summary statistics; the underlying claim is that the entire temperature power spectrum follows from the derived chart. Fed GG-Theory’s exact fractions, the carried rate

Table 7: The cosmic-microwave-background observables predicted by GG-Theory (grade D* throughout), confronted with Planck 2018. *Zero CMB input:* every row is forced by the geometry — the fractions by the Hopf fibration, the rate by the two routes of §24, the densities by their product — and the sky appears only in the measurement column. Every computed row is evaluated at the carried rate 66.8 ± 0.2 ($h = 0.6676$) in the released pipeline of GG-A10-Cosmos (`gg_cmb_tt_v4.py`, `peaks_v4.py`); at the two-road 66.6 the density rows read 0.02247 (+0.5%) and 0.14119 (−1.3%). The observed baryon-to-matter ratio is the ratio of the two observed rows above (the neutrino-free $0.02237/0.1424 = 0.15709$ would read +1.3%). The z_{dec} row is GG-A10-Cosmos’s atomic-chain derivation of the decoupling itself, evaluated at the framework’s own h with zero fitted parameters, +0.15% from the Planck value and inside the 0.18% modeling systematic of the recombination calculation; the corridor-switch redshift $z_* = e^7 = 1096.6$ is a topological marker of the OGN-5 integer $J = 7$, 0.38% above it in $1 + z$, annotated in that row and not scored. The pull column follows the GG-A10-Cosmos convention, $(\text{GG} - \text{obs})/\sigma$ with $\sigma = \sqrt{\sigma_{\text{obs}}^2 + \sigma_{\text{th}}^2}$ where a theory tolerance is declared and audited there (θ_* , ℓ_D , H_0 , ω_b , ω_m : the propagated inventory tolerance; z_{dec} : the 0.18% modeling systematic); where a tolerance is declared the pull column prints *both* readings side by side, raw / tolerance-adjusted, so that no row’s convention can flatter it; rows without a declared tolerance (r_s) are raw by construction.

Observable	GG forced form	GG	Planck 2018	Dev	Pull
H_0 (Route I, carried)	$498.3\sqrt{\pi/175}$	66.8 ± 0.2	67.4 ± 0.5	−0.9%	−1.1
H_0 (two-road, $I \cap II$)	one reading	66.6 ± 0.2	67.4 ± 0.5	−1.2%	−1.5
ω_b (predicted)	$\Omega_b h^2 = h^2/2\pi^2$	0.02258	0.02237 ± 0.00015	+0.9%	+1.0
ω_m (predicted)	$\Omega_m h^2 = h^2/\pi$	0.14187	0.1430 ± 0.0011	−0.8%	−0.8
ω_b/ω_m	$1/2\pi$ (h -free)	0.15915	0.15643 ± 0.0016	+1.7%	+1.7
z_{dec}	atomic chain, $h = 0.6676$	1091.5	1089.92 ± 0.25	+0.15%	+6.3 / +0.8
$r_s(z_{\text{rec}})$ [Mpc]	acoustic scale	144.7	144.43 ± 0.26	+0.2%	+1.0
$100\theta_*$	r_s/D_M	1.0381	1.04109 ± 0.0003	−0.3%	−10 / −2.5
peaks $\ell_1/\ell_2/\ell_3$	acoustic projection	221/539/816	220/536/813	+0.4–0.6%	—
height ratio 2/1	image of $1/2\pi$	0.449	0.453 ± 0.008	−0.8%	−0.5
height ratio 3/1	image of $1/2\pi$	0.439	0.443 ± 0.009	−0.9%	−0.4
ℓ_D (damping)	diffusion integral	1350 ± 3	1345 ± 15	+0.4%	+0.3 / +0.3

$H_0 = 66.8$ ($h = 0.6676$), and the derived tilt $n_s = 0.9634$, exact on the plateau at $N^* = 52.6$ ($1 - 2/N^* = 0.9620$ at leading order) — with only the seed amplitude A_s and the optical depth τ inherited — a Boltzmann solve (CLASS; Blas et al., 2011) returns a *zero-adjustable-parameter* temperature spectrum. Against the 83 binned Planck TT points it scores

$$\chi_{\text{GG}}^2 = 187 \quad (\text{zero free background parameters}) \quad \text{vs} \quad \chi_{\Lambda\text{CDM}}^2 = 84 \quad (\text{six free parameters}), \quad (30)$$

i.e. $\chi^2/\text{dof} \approx 2.3$ for a curve with no background parameter tuned to the sky. The 187 is not 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* — the same correlated offset the next subsection names. 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. The tilt is not tested here; it is tested in the exact-plateau section of GG-A10-Cosmos, where it is the framework’s one currently failing test (§56, F-Hubble-6 and the tilt beside it). An earlier version of this paper printed a $\Delta\chi^2 = -22$ preference for the derived tilt from a run at the leading-order value; that number is withdrawn.

What the χ^2 means, said first by us. Let us state the exclusion as bluntly as a referee would, because it is ours to say first: taken as an *exact* hypothesis, $\chi^2 = 187$ on 83 bins — $\Delta\chi^2 = +103$ over the six-parameter fit, roughly 10σ Gaussianized as $\sqrt{\Delta\chi^2}$ — rejects the exact chart outright, and a diagonal-covariance protocol can only understate that. This paper’s claim was never that the exact chart is the sky’s last word. The claim is three clauses, each at its own strength: *consistent at the percent level* (a chart with no background parameter tuned to the sky lands within about a percent of the entire acoustic pattern); *excluded as exact* (the sentence above, printed by us, not extracted from us); and *falsifiable in one move* (the miss is one coherent, named displacement carrying its break condition). Every headline in this paper is meant to be read at exactly that strength — no more, and no less.

The spectrum itself is shown in Figure 15: the derived chart’s temperature spectrum overlaid on the eighty-three binned Planck points, with the residual panel making the single coherent oscillation visible to the eye. The figure is produced by the parent GG-A10-Cosmos’s Boltzmann pipeline — this paper adds no spectrum computation of its own — and is reproduced here because Part III’s argument should not have to be taken on a number alone.

26.4 The coherent π -offset: one systematic, visible everywhere

Figure 16 collects every confrontation of this Part — and the landmark checks of §25 — on a single axis. The pattern a fitted model would show is scatter straddling zero, sized by the error bars. The pattern the chart shows is different in kind: seventeen deviations — all but the baryon fraction within 1.7%, none beyond 2.8% — each with the sign the one $+1\%$ displacement of the exact fraction above Planck’s fitted value propagates — positive where the displaced inventory pushes a scale up (the loading, the predicted baryon density, z_{dec} , r_s , the peak positions, the damping scale, equality,

the age), negative where it pushes one down (the derived h , the predicted matter density — where the doubled rate deficit outweighs the fraction excess — the projected angle, the height ratios). This is also why the derived point is testable at all rather than hidden: the exact fractions sit *along* the geometric degeneracy direction of Λ CDM (§26) with a printed transverse residue (θ_* , $\Omega_m h^3$), so the whole pattern rides one lever, and a future determination that drives Ω_m firmly below $1/\pi$ breaks every entry at once.

Let us also quantify the coherence rather than gesture at it. All seventeen displacements carry the sign the single +1% displacement propagates; were the seventeen independent coin flips, that pattern would have probability $2^{-17} \approx 8 \times 10^{-6}$. They are *not* independent — every acoustic entry flows through two or three underlying quantities, so the effective number of one-directional facts is closer to two or three than to seventeen — and the honest reading is therefore this: the sign pattern is the *necessary* signature of a single displaced parameter, cheap for any such parameter to produce, and it distinguishes a displaced derivation from a scattered fit — not a right derivation from a wrong one. The sufficient test is the global likelihood at the fixed chart, registered as owed at OQ-Hubble-4; until it is run, the coherence claim is carried at exactly this weight.

The reason the fit is not perfect is a single, coherent displacement. The exact fraction sits just above Planck’s fitted one, $1/\pi = 0.3183$ versus 0.3153 ± 0.0073 (+1.0%), and that *one* displacement propagates, with the expected sign, through every observable in Table 7: the derived h comes out 0.9% low, θ_* projects 0.3% small, the peaks slide 0.4–0.6% rightward, the damping scale follows at +0.4%, and the atomic z_{dec} and the loading ratio $1/2\pi$ sit +0.15% and +1.7% high. This is the diagnostic difference between a fit and a derivation. *A model fitted to the data scatters its residuals on both sides of zero; a chart forced by geometry and displaced by one small systematic lands coherently on one side.* The offset is therefore either the last small truth about the boundary measure still to be absorbed, or the falsification-in-waiting: a future determination that drives Ω_m firmly below $1/\pi$ would break the chart.

27 Summary: the derived history, and the road to Part IV

Part III has done the constructive work of the paper, and this closing section gathers it in one place, ending where Part IV must begin.

The framework enters through one stage and one ledger (§20): the six-dimensional bulk with its parameter-free metric (7), the single gauge group $\text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}$, the BI-6 legality certificate, and the OGN-5 integers (1, 3, 5; 7, 35). Nothing new was introduced for cosmology: the same package that pins the twenty-eight Standard-Model parameters in GG-A9-SM is the whole of what this Part was allowed to use, with the parent GG-A10-Cosmos supplying the cosmological transcription.

From that stage the expansion law is a theorem, not an assumption (§21). The cosmological principle is the runtime’s symmetric attractor; the Einstein equation is inherited from the boundary projection; flatness is exported, not measured in; and the Friedmann system follows by the standard computation, walked step by step.

The inventory that fills the law is itself a theorem, and the deepest result of the Part (§22): one




Figure 15: The zero-parameter temperature spectrum against the microwave sky (produced by the Boltzmann pipeline of the parent GG-A10-Cosmos; reproduced here as the visual form of Eq. (30)). *Top:* the derived chart (red dashed: every background input derived — $\Omega_m = 1/\pi$, $\Omega_b = 1/2\pi^2$, the derived rate $H_0 = 66.8$ (Route I, $h = 0.6676$), the derived tilt $n_s = 0.9634$ exact on the plateau; only the seed amplitude A_s , the sector's one declared input, and the astrophysical τ inherited — no background parameter is read from the sky the curve is tested against, though A_s and τ are inherited from the sky's own fit) against the Planck 2018 best-fit Λ CDM (black) and the eighty-three binned TT points. *Bottom:* residuals against the best fit, in percent: the derived curve's residual is not scatter but a single coherent oscillation between -1 and $+3.5\%$ across the acoustic peaks — the spectrum-level image of the coherent π -offset of §26.4 — bracketed by the derived rate's uncertainty (dotted, $H_0 = 66.6/67.0$). The scores are printed on the panel: $\chi^2 = 84$ for six fitted parameters, 187 for zero ($\chi^2/\text{dof} \approx 2.3$); on this background χ^2 is monotone in the tilt and does not select the computed value. Reproduction: `gg_cmb_tt_v4.py` of GG-A10-Cosmos.

Hopf fibration — the matter charge as the fiber of the electroweak 3-sphere — fixes $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, and $\Omega_b = 1/2\pi^2$ as three readings of one shape, the π descending from a sphere volume and nothing read off the sky. The grade is stated, not hidden: D^* , derived conditional on one named measure lemma, whose promotion path is tracked in the appendices.

Feeding that inventory into the law (§23) turns it into a single curve, in closed form: $H(t) = \sqrt{\Omega_\Lambda} H_0 \coth \tau$ from equality onward, falling monotonically for all time — never rising, even as the expansion accelerates — onto the de Sitter floor $H_\infty = 55.2 \text{ km s}^{-1} \text{ Mpc}^{-1}$, with the present age $t_0 = 0.948/H_0 = 13.88 \text{ Gyr}$ an output along the way, and the scale factor and the running budget drawn as their own objects (Figures 11 and 10).

The present rate is then derived before the sky is consulted (§24). Route I is geometry alone: the GU-5 bridge turns the theorem-grade lock scale $M_{\text{lock}} = 3.85 \text{ TeV}$ (laboratory-pinned by the electroweak vev at 0.045%) and the codebook-selected vacuum coefficient $\zeta_\Lambda = \pi/175$ into $H_0^{\text{geom}} = 498.3\sqrt{\pi/175} = 66.8 \pm 0.2$ — a chain with no astronomical observation in it, on whose admitted stratum cluster no monomial returns the ladder’s 73.5, and which lands, at the quarter power, as the familiar meV dark-energy scale (2.23 against $2.24 \pm 0.02 \text{ meV}$ — a unit identity of the same agreement, counted once; §24.2). The sky then confirms it with precision (Route II): dividing its two h -independent densities by the two predicted fractions returns 66.5 and 67.0 — two anchors on different observables, agreeing internally at 0.7% — and BBN’s deuterium, with no CMB anywhere, returns 66.4 ± 0.5 ; the concordance reads $66.6 \pm 0.2 \text{ km s}^{-1} \text{ Mpc}^{-1}$ on the two-road side, 0.25% from the carried 66.8 (Figure 12). Agreement of this kind cannot be arranged after the fact without breaking one of the two routes’ other commitments — or the electroweak vev, or the acoustic spectrum.

The derived history then survives every checkpoint it faces. Walked from the first second to the infinite future (§25, Figures 13 and 14 — the pair at the heart of this paper), seven measured landmarks land inside one percent, every deviation carrying the same propagated sign. And the microwave background over-determines it (§26): the ledger deletes the background half of ΛCDM ’s six parameters, the h -free baryon loading lands at 1.7σ , the two physical densities themselves become predictions rather than inputs, the full spectrum returns $\chi^2 = 187$ over the 83 binned TT points with no background parameter tuned (Figure 15), and the residual is not scatter but a single coherent π -offset carried by all seventeen observables with their propagated signs (Figure 16).

The bottom line for the tension is one sentence long: the derived rate lands on the early-universe value — 66.8 against 67.4 ± 0.5 , the gap being the coherent offset itself — so *the early side of the tension needs no new physics*. What remains is the late side: why does the local ladder read 73.5? That question is not a loose end but the paper’s second half. Part IV reads the local excess whole as a calibration systematic the framework predicts, prices the one environmental channel it could have credited and does not carry it, confronts the systematic with the H0DN error budget, and states what would falsify the reading; Part V ranks the rival explanations against the same derived rate; and Part VI hands the verdict to the one channel that owes nothing to either side. The derivation is done; the reconciliation begins.

28 Part III in plain English: how a handful of integers becomes an expansion history with nothing left to adjust

This is the derivation. The thing to watch is not the arithmetic but the absence of dials.

The framework's cosmology fixes what the universe is made of by geometry, not by fitting. The matter fraction is $1/\pi$, the baryon fraction is $1/2\pi^2$, and dark energy is whatever remains. They are not tidy numbers noticed after the fact: the same integer ledger fixes the particle sector in the companion papers, and it was written down before this reading was attempted.

Given the contents, the expansion history follows from the Friedmann equations, which this framework derives rather than assumes. Today's value then falls out of $H(t)$ in two separate ways. The first never looks at the sky: a vacuum-energy relation, its one coefficient chosen by a rulebook written for other purposes, returns the rate from the geometry alone. The second divides the geometric fractions into the two physical densities the microwave sky measures. The two agree, and they land on 66.8 and 66.6, a quarter of a percent apart.

The curve is then asked for more than one number: the sound horizon, matter–radiation equality, the age of the universe, the acoustic spectrum. It supplies all of them, and every one misses in the same direction by less than a percent — which is interesting, and is not proof, because one unaccounted systematic is exactly what a coherent offset would look like. The Part names it rather than absorbing it.

What this settles is half the problem. The derived rate sits on the early-universe side of the split, so everything remaining is a question about the local ladder, and that question is the next Part.

One Displacement, Seventeen Shadows: the Coherent π -Offset

Zero CMB input: every entry predicted by GG-6/GDOS/GRIP; the sky enters only as the comparison column.




Figure 16: One displacement, seventeen shadows — the parameter-audit Dev chart. Fractional deviation (derived – measured)/|measured| for every confrontation in this Part, in the audit format of the parent GG-A10-Cosmos: nested precision bands at $|\text{Dev}| < 10^{-3}$ and 10^{-2} ; inner thick bar the measurement error $\sigma_{obs}/|obs|$; outer thin bar the combined $\sqrt{\sigma_{obs}^2 + \sigma_{th}^2}/|obs|$ where a theory tolerance is declared (H_0 , ω_b , ω_m , z_{dec} , θ_* , ℓ_D); open circles the peak positions, quoted without uncertainties. Every computed row is evaluated at the carried rate 66.8 ($h = 0.6676$) in GG-A10-Cosmos’s released pipeline; the decoupling row is the atomic chain’s 1091.5, the switch e^7 annotated and not scored. *Zero CMB input*: every entry is a GG-6/GDOS/GRIP prediction — including the two physical densities, which the old inversion consumed and the derived rate returns as outputs — and the sky enters only as the comparison. Every entry is D*, and every sign is the one the single +1% displacement of $1/\pi$ above Planck’s fitted Ω_m propagates. A fit scatters its residuals on both sides of zero; a derivation displaced by one small systematic lands coherently — the diagnostic difference the text describes. Generated by `c2_showcase_offsets.py`, released with the paper.

Part IV

Reconciling the Hubble Tension with GG-Theory

The origin of the larger local value; the quantitative estimate; the limit of backreaction as a remedy; and a call to revisit Λ CDM in the light of the derived $H(t)$

29 The task of this Part, and the strategy

Part III ended with an asymmetry. The derived expansion history lands, without being asked, on the early-universe value: 66.8 against 67.4 ± 0.5 , the gap being the coherent π -offset the chart itself predicts. The early side of the tension is therefore *finished business* — no early dark energy, no extra radiation, no shortened ruler. What is not finished is the late side: the local distance ladder reads 73.50 ± 0.81 , ten percent above the derived background, and a derivation that stopped here would have explained the easy half of the tension and gone silent on the hard half. This Part is the hard half.

Recall the three possible homes of Part II (§10): a shared calibration bias in the ladder, new physics before recombination, or a bent expansion history in between. The derivation has already closed the second and third doors from the inside — the ruler is right ($r_s = 144.7$ against 144.43 ± 0.26) and the bridge is a theorem with no continuous freedom — so the framework is *forced* onto the first door. That is worth stating plainly, because it is the opposite of a free choice: most proposed resolutions pick the door they can decorate, while a zero-parameter derivation must walk through the only door its own theorems leave open. Everything in this Part is the disciplined consequence of that forcing.

The strategy is three verbs, applied in order. *Bound* what the framework might have been tempted to invoke: the environmental lift of a local under-density, whose linear response is computed and whose required depth the surveys do not show; the $U(1)_B$ channel, which the boundary-condition termination of GG-A3-B16 closes structurally and the equivalence-principle bounds close empirically; and the coarse-graining (backreaction) term, three orders too small. Every internal escape route is closed. *Name* what remains, which is all of it: a $\sim 10\%$ common-mode calibration systematic in the ladder’s rungs, *predicted* by the framework — on one background every probe without a calibrated rung reads the same rate, so the displaced instrument must be the one with rungs — and not derived. *Confront* it, channel by channel, with the ladder’s own error budget, priced from both directions, and hand the verdict to the one probe that owes nothing to either side. An earlier version of this Part had a fourth verb, *derive*, and credited two to three percent of the excess to the neighborhood; the parent GG-A10-Cosmos has withdrawn that credit, for a reason this Part now prints in full, and there is nothing left to split. The honesty geometry of the whole Part is drawn in

Figure 18: hatched is the predicted systematic, violet is external, faded is the void channel priced and not carried.

One thing has changed since this strategy was first drawn, and it changes the flavor of the second verb. A residual can be named in two ways: as a blank — “some shared calibration error, somewhere” — or as an anatomy, a list of specific physical mechanisms with published sizes and dated tests. The 2025–2026 literature has moved the residual from the first kind to the second: two medium-size, physically motivated, common-mode channels — a dust-law mismatch between the ladder’s calibrator galaxies and its Hubble-flow training set, and a statistical bias built into the ladder’s own Bayesian distance inference — have been proposed, sized, and disputed in the field’s own journals (§34.3). Neither is a GG-Theory result, and neither is settled; but “name what remains” now means pricing *named candidates* rather than confessing a blank, and the arithmetic of §34 changes shape accordingly.

The Part walks that strategy in sequence: the reading and its referee (§30), the switch that separates the two dictionaries (§31), the worked estimate (§32), the void channel priced in full and not carried (§33), the systematic against the error budget (§34), the backreaction bound that closes the last internal channel (§35), and then the wider audit — the why-now window, the not-a-fit argument, cross-ledger coherence, and the S_8 companion — before the Part closes with an invitation and a summary.

30 The Hubble tension as a question of which instruments read the derived rate

With the background rate now derived, $H_0^{(0)} = 66.8 \pm 0.2$ (Section 24), the Hubble tension changes character. It is no longer a contradiction between two measurements of an unknown number; it is one derived number and the question of which instruments read it. The boundary Friedmann equation has one background solution on the π -locked inventory, with one rate — the content of GG-A10-Cosmos’s Theorem T-Cosmos-16 — and there is no second branch. The switch at the *corridor-switch redshift* $z_* = e^7 \approx 1097$ of GG-A10-Cosmos — a topological quantity fixed by the OGN-5 integer $J = 7$, distinct from the operationally measured decoupling redshift $z_{\text{dec}} = 1089.92 \pm 0.25$, with which it coincides to 0.6% in z — separates the two dictionaries through which that one background is read; nothing about the rate changes across it. An earlier version of this section read the early and late values as two branches of one solution meeting at the switch; the parent has withdrawn that reading, and this section now agrees.⁵

30.1 The early side needs no new physics

The CMB-inferred value 67.4 ± 0.5 lands directly on the derived background 66.8; the -0.9% gap is the *coherent π -offset* — GG’s exact fractions sit 1–3% above Planck’s best-fit fractions, and H_0 inherits the same small displacement. Crucially, the derived rate needs nothing from the sky

⁵Where this paper speaks loosely of “recombination” as the switch, the precise object is always the corridor-switch redshift $z_* = e^7$, the topological last-scattering landmark of GG-A10-Cosmos, not a prediction of the measured z_{dec} .

(Route I, §24); the confirmation used only the CMB’s own model-independent densities (25) and reproduced the acoustic peaks with CLASS at ≈ 67 . *The early side of the tension is therefore not biased at all*; it is where the derived rate is read most cleanly, and it requires no pre-recombination modification — no early dark energy, no extra radiation. This is a decisive simplification: the CMB endpoint is not a displaced projection to be explained, but the anchor itself. The intuition is a surveyor’s: most resolutions treat the CMB value as a signpost knocked askew that must be straightened before the map can be trusted, and then invent the physics that knocked it. Here the geometry itself printed the signpost — the derived rate and the CMB reading are one number seen twice — so the map needs no correction, and the whole burden of the tension moves to the other end of the road. And the signpost has grown independent legs since the tension was first framed: the derived rate stands on Route I with zero sky input; the sky’s own reading of it now stands on three separate instruments — Planck, ACT, and SPT-3G, with different sites, detectors, and calibration chains — agreeing at a fraction of the offset in play, the newest of them, SPT-3G alone, reading 66.66 ± 0.60 : directly on the derived rate; and the sound-horizon-*free* early determinations span 66.7–70.2 (with errors of ± 2), every one of them well below 73, so even if the acoustic ruler itself were struck from the record, the early side would not move toward the ladder (Part V returns to all three).

30.2 The late side: one predicted systematic

The local distance-ladder value 73.50 ± 0.81 sits $\sim 10\%$ above the derived background. The framework reads the whole of it as a calibration systematic in the ladder’s rungs,

$$H_0^{\text{late, ladder}} = H_0^{(0)} (1 + \epsilon_{\text{syst}}), \quad \epsilon_{\text{syst}} \simeq +0.10, \quad (31)$$

predicted, not derived. It offers no physical channel for any part of it, and the two it once offered are withdrawn below: the local-void lift, because the under-density it would need is not observed (§33), and the $U(1)_B$ channel, because under the boundary-condition termination of GG-A3-B16 no vector survives below the lock scale to mediate it (§32, App. G.5). An earlier version of this subsection factorized the excess into a derived environmental piece and a named residual and printed, as the reading a ladder inside the Local Hole should return,

$$H_0^{\text{late}}|_{\text{background+lift}} \simeq 68.5 \pm 0.6 \quad (\text{withdrawn: the lift is not carried}), \quad (32)$$

which is kept here only so the record of what changed is visible. GG does *not* derive ϵ_{syst} . It predicts where it lives — in the measurement chain, not in cosmology — and states its size as an open problem, not a result. The honest bottom line: the early side is derived and unbiased; every ladder-free late probe is predicted to read the same rate; the ladder alone is displaced, by a systematic the framework predicts and the ladder-free probes will test. What the framework does *not* do is tune anything to close it; it is reported, bounded, and left open.

30.3 The age of the universe: the time-axis image of the rate

The same derived rate fixes the present age. The dimensionless product $H_0 t_0 = 0.948$ is set by $\Omega_m = 1/\pi$ alone, so at the derived background $H_0^{(0)} = 66.8$

$$t_0 = \frac{0.948}{H_0^{(0)}} \approx 13.9 \text{ Gyr (today)} \quad (13.8 \text{ Gyr at the } \Omega_m=1/\pi \text{ crossing, 75 Myr before}). \quad (33)$$

The familiar 13.8 billion years is recovered exactly as the age at the inventory crossing; the local-ladder reading 73.5 would instead give a younger ≈ 12.6 Gyr. The present age is therefore the time-axis image of the rate — one derived value, 13.9 Gyr; the 12.6 Gyr the ladder’s 73.5 would imply is the same displacement seen in H_0 , read on the time axis, and the oldest stellar populations at 13–13.5 Gyr sit beneath the derived age and not beneath the ladder’s (GG-A10-Cosmos).

30.4 The model-independent referee: standard sirens

Because the derived rate is the true background expansion rate, it must be recovered by any probe that avoids both the sound horizon and the distance ladder. The gravitational-wave standard siren is exactly such a probe: it reads the luminosity distance directly from the waveform. GG-Theory therefore makes one sharp, dated, model-independent prediction:

$$H_0^{\text{GW}} \longrightarrow \approx 67 \text{ km s}^{-1} \text{ Mpc}^{-1} \text{ (not 73)}. \quad (34)$$

A siren determination that converged near 73 instead of near 67 would refute this reading outright: it would mean the true background rate is high, the CMB densities cannot be divided by GG’s fractions to give it, and the derivation of Section 24 is wrong. This is the cleanest single test of the whole reading, and the coming decade of binary-neutron-star and dark-siren catalogs will settle it.

31 The recombination switch and the two dictionaries

The early and late measurements are separated by a real geometric event, and naming it sharpens the whole discussion. In the GRIP cascade, recombination is not merely the moment the plasma turns transparent; it is the GSSB transition from the tightly coupled DGI plasma to the frozen geometric attractor — the defining switch of the boundary Friedmann history — and its redshift is fixed by the integer ledger with no continuous parameter free.

Proposition 31.1 (The recombination-era cosmic GRIP-Out switch (T-Hubble-2)). *Under the GRIP cascade of GG-A6-GRIP, the recombination era hosts the GSSB transition from the DGI plasma to the geometric attractor; the defining switch event of the history within that era is the geometric attractor change, and the OGN-5 integer ledger of GG-A4-Euclid fixes its redshift to*

$$z_* = e^J = e^7 \approx 1096.6 \quad (J = 7),$$

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; nothing about the rate changes across the switch. The atomic decoupling surface — the visibility peak Planck measures — is derived from the framework’s own atomic chain in *GG-A10-Cosmos* and lies 0.4% (in *e*-folds) below the switch, at the measured value.

The switch $z_* = e^7 \approx 1096.6$ is a topological marker, distinct from the Planck decoupling redshift $z_{\text{rec}} = 1089.92 \pm 0.25$, with which it coincides to 0.6% in z -space (0.4% in *e*-folds). The decoupling itself is not left to measurement: *GG-A10-Cosmos* derives it from the framework’s own atomic chain — the Rydberg from the particle sector, the baryon inventory $\Omega_b = 1/2\pi^2$, the derived $h = 0.6676$ — with zero fitted parameters, landing at $z_{\text{dec}} = 1091.5$, +0.15% from the Planck value, and computes the switch–decoupling offset outright ($\delta = 0.0038$ *e*-folds, 0.38% in $1+z$). Scored against the Planck centroid the retired identification “ $e^7 = z_{\text{rec}}$ ” would miss by 27σ — which is exactly why neither paper makes that identification: the precision comparison belongs to the atomic quantity, and Table 7 carries the row in the two-quantity form. The point for the tension is structural, and it cuts the other way from the one an earlier version of this paper drew: reading one background through two operational dictionaries — early inference through a sound horizon calibrated at the switch, late inference through a distance ladder read on the present attractor — is not a license for two answers. The audit must ask, instrument by instrument, whether the chain reads the background or reads a calibrated rung. The sound-horizon dictionary reads the background, 0.9% from the derived value, the difference being the inherited calibration this paper does not recompute; the ladder-free late probes — acoustic oscillations in the galaxy distribution, sirens, time delays — read it directly; the calibrated ladder reads a rung. In *GDOS* language a finite observer never samples the closed six-dimensional bulk, and every reading is a coarse-graining of one boundary history; the switch marks where the dictionaries change, and it does not move the rate.

The intuition deserves its own sentence, because “two dictionaries” sounds like philosophy and is actually bookkeeping. One text — the single background history $H(t)$ — is read through two translation tables: the early dictionary converts an acoustic pattern into a rate via plasma physics and a ruler frozen at the switch; the late dictionary converts photometry into a rate via stellar calibration read on the present attractor. Two careful translators of one known text should agree, and where they do not, the text decides which dictionary carries the error. The framework’s claim is exactly that: the text — the background rate — is known, so the discrepancy diagnoses the dictionaries, and it is the ladder’s dictionary, with its calibrated rungs, that the framework predicts to be the one in error. (The earlier version of this paragraph said that two translators “can render one text differently without either lying”; the parent’s one-background theorem says the opposite, and it is the falsifiable statement.)

The lever by which an early-universe dictionary *could* be made to differ is the acoustic angle. The observed $\theta_* = r_s/D_A$ is fixed to four digits, so a downward shift in the sound horizon must be compensated by an upward shift in the inferred late-time rate:

$$\theta_* = \frac{r_s}{D_A} \implies \frac{\delta\theta_*}{\theta_*} = \frac{\delta r_s}{r_s} - \frac{\delta D_A}{D_A}, \quad \frac{\delta H_0}{H_0} \approx -\frac{\delta r_s}{r_s} \quad (\text{fixed } \theta_*). \quad (35)$$

This is the standard lever by which any early-universe modification trades sound horizon for inferred

H_0 (Bernal et al., 2016; Kamionkowski and Riess, 2023; Poulin et al., 2019). The intuition: θ_* is a protractor reading that cannot be argued with, so if a theory shortens the stick, the only way to keep the measured angle is to hold the stick closer — a smaller distance to last scattering, hence a faster late-time expansion, hence a higher inferred H_0 . Every early-universe “solution” buys its higher rate by shaving the stick. The reading of this paper does *not* pull that lever: the derived background already sits at the CMB value (§26), so $\delta r_s = 0$ and H_0^{CMB} needs no early modification. The whole of the excess is on the late side, and it is to a quantitative estimate of that excess that we now turn.

32 A worked estimate of the local excess

The observed local rate is higher than the background by a definite amount,

$$\epsilon_{\text{local}} \equiv \frac{H_0^{\text{late}}}{H_0^{(0)}} - 1 = \frac{73.50}{66.8} - 1 = 0.100, \quad (36)$$

and the framework carries it whole. An earlier version of this section factorized it into a piece the framework derived and a piece it named,

$$\epsilon_{\text{local}} = \epsilon_{\text{Hole}} + \epsilon_{\text{syst}}, \quad H_0^{\text{late}} = H_0^{(0)} (1 + \epsilon_{\text{Hole}})(1 + \epsilon_{\text{syst}}), \quad (37)$$

and the split is withdrawn — ϵ_{Hole} is not carried; the form is kept here because the bookkeeping below still uses it to say what a void of a given depth *would* contribute. The physical content that survives is simpler than a factorization: one piece, and it lives in the measurement chain.

The one environmental channel, priced. The nearby universe is a mild under-density (the observed “Local Hole”), and a distance ladder calibrated inside an under-density reads the local expansion rate high. In linear theory the monopole of the peculiar-velocity field gives, for a top-hat under-density $\delta < 0$ on the ladder scale,

$$\epsilon_{\text{Hole}} = \frac{\delta H_0}{H_0} = -\frac{1}{3} f \delta, \quad f = \Omega_m^{0.55} = 0.53, \quad (38)$$

so a mild $\delta \approx -0.15$ would yield $\epsilon_{\text{Hole}} \approx +2.6\%$, and the whole excess would need $\bar{\delta} \approx -0.56$ over the flow volume. Neither is observed. The flow data prefer voids smaller than ~ 70 Mpc, the supernova Hubble diagram leaves room for at most a sub-percent shift, the network’s own no-correction variant moves the consensus by 0.6%, and the counts-literature depth that would give even a few percent is an $\approx 11\sigma$ fluctuation on the framework’s own power spectrum (§33). GG-Theory supplies the *response law* as a GRIP-In structure-formation feature — the sign is unconditional and the coefficient is fixed by the inventory — and it credits *no share* of the excess to the neighborhood, because the depth that share needs is not there. §33 computes the channel in full and records why it is not carried.

The excess, named. $\epsilon_{\text{syst}} \approx +10\%$ is a distance-ladder calibration systematic the framework *predicts* and does not derive, and it predicts it for a concrete reason: on one background every probe

without a calibrated rung reads the same rate, so the instrument that disagrees must be the one with rungs. The prediction is structural, not a preference. And the one first-principles channel the framework owns is closed — twice over. Structurally: 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, so there is nothing to mediate a lift. Empirically, and independently: a baryon-coupled fifth force would need coupling $\alpha_B \sim 3 \times 10^{-2}$ to give a 3% lift — the scaling is essentially linear, because a fifth force of relative strength α_B rescales the effective attraction the calibrators and the flow feel, and percent-level rate shifts need percent-level couplings (App. G.5 carries the factors) — and equivalence-principle and lunar-laser-ranging bounds cap it at $\alpha_B \lesssim 10^{-9}$, seven to eight orders too small. (The first-pass estimate an earlier version carried, $\epsilon_{\text{local}}^{(B)} \sim 10^{-3}$, is recorded in App. G.5 and is moot.) The cosmological-averaging (backreaction) channel is bounded below 10^{-3} by the computation of §35. So the excess is not a cosmological effect the framework is hiding; it is a genuine open problem, located — honestly — in the measurement chain, where the ladder-free standard siren (§30) can decide it. And “located in the measurement chain” is no longer a gesture at a blank: the chain’s own literature has begun to name the candidate mechanisms, with published sizes and dated kill tests (§34.3) — and, as §34 prints against interest, they do not yet reach ten percent.

Numerically, then: $H_0^{(0)} = 66.8$ (derived; the microwave background reads it at 67.4 through the inherited sound-horizon calibration), and $\epsilon_{\text{sys}} \approx +0.10$ (a predicted ladder systematic, not derived) $\rightarrow 73.5$. Everything to the left of the last step is a theorem or a graded derivation; the last step is labeled, priced, and left open. The two sections that follow substantiate the reading: §33 prices the void channel in full — evidence, linear theory, dilution, and budget — and records why it is not carried, and §34 confronts the predicted systematic with the ladder’s own error budget, term by term.

33 The local void, priced and not carried: evidence, linear theory, and budget

The environmental lift is the one piece of the local excess the framework could have credited, so it deserves better than a formula and a range: this section assembles the observational evidence for the Local Hole, derives the response from linear theory in three steps, works through the two effects that dilute it — the ladder’s own averaging volume and the network’s flow corrections — and closes with a scenario ledger and a systematic budget in the cohort’s reproducibility format. The honest summary, stated in advance and agreeing with the parent GG-A10-Cosmos: the *sign* of a void’s lift is unconditional and its *response* is computed; the *depth* that would let it carry even a few percent is an $\approx 11\sigma$ fluctuation on the framework’s own power spectrum and is not observed in the velocity data — which is why the paper adopts no lift and records the scenarios as what the surveys can measure. An earlier version adopted +2–3% here; the arithmetic that follows is the same, and the conclusion drawn from it is not.

33.1 The evidence: how deep, how wide, how contested

That the nearby universe is under-dense is not in serious dispute; how deep and how far out is. Three families of readings span the field. At the deep end, near-infrared galaxy counts find a hole of striking size: a $\sim 30\%$ deficit extending to roughly 300 Mpc (Keenan et al., 2013) — the “KBC void” — with regional count-and-redshift surveys finding $\delta \approx -0.15$ to -0.24 over 150–250 Mpc in the survey directions (Whitbourn and Shanks, 2014). At the shallow end, two independent brakes push back. First, statistics: in standard structure formation the rms density contrast on a 300 Mpc sphere is $\sigma_\delta \approx 0.015$ (the same linear machinery as Appendix H.4), so a -0.3 contrast at that scale would be a $\sim 20\sigma$ fluctuation — taken at face value it would itself be evidence against Λ CDM Gaussianity, which is precisely why the deep reading is contested rather than adopted. And the same arithmetic must be charged to this paper’s own ledger: the canonical row’s $\bar{\delta} = -0.16$ over the flow volume is itself a $\approx 11\sigma$ Gaussian fluctuation on the same σ_δ . We print that number rather than hide it: the counts branch is recorded, not carried, as the *counts literature’s claim* — a conditional input, not a GG-6 output — and a neighborhood that deep would itself be a discovery about growth statistics. What GG-6 derives is the *response law* (the master formula below); the depth is the disputed quantity, and the shell test (F-Hubble-4) is the instrument that reads it directly. Second, dynamics: the supernova Hubble diagram is smooth across the relevant redshifts, and dedicated analyses find it leaves room for at most a sub-percent void-induced shift in the SN-anchored rate (Kenworthy et al., 2019); the H0DN’s own no-flow-correction variant (V19) moves the consensus by only $-0.44 \text{ km s}^{-1} \text{ Mpc}^{-1}$. A third brake has since been applied from *within* the void camp itself: field-level forward modeling of the CosmicFlows-4 peculiar-velocity data finds that the velocities prefer voids smaller than ~ 70 Mpc — under a tenth of the KBC extent — and that the deep profiles which would help H_0 most are precisely the ones that contradict the measured bulk-flow curve (Stiskalek et al., 2025). The physics of the brake is worth one sentence: a void lifts the local rate through its outflow, and the same outflow *is* the bulk flow the velocity surveys measure — one cannot order the lift without the flow, and the flow data decline the order. The deep reading is therefore no longer merely contested; it is disfavored by the void camp’s own velocity data. (The CosmicFlows-4 bulk-flow excess itself — $419 \pm 36 \text{ km s}^{-1}$ at $200 h^{-1} \text{ Mpc}$ (Watkins et al., 2023) — is recorded as contested companion evidence, with selection and estimator critiques still in play.) The truth the framework needs lies in the tension between these readings, and the derivation below is built to expose exactly where each reading lands.

33.2 Linear theory: three steps to the response

Step 1 — continuity. In linear perturbation theory the growing mode ties the peculiar-velocity field to the density field: mass conservation gives $\dot{\delta} = -\nabla \cdot \mathbf{v}/a$, and for the growing mode $\dot{\delta} = fH\delta$ with $f \equiv d \ln \delta / d \ln a$, so

$$\nabla \cdot \mathbf{v} = -a f H \delta. \quad (39)$$

An under-density ($\delta < 0$) is a region the matter is flowing out of: positive divergence, outflow.

Step 2 — the observer’s rate. A local observer infers the expansion rate from the recession of tracers inside a volume of radius r : the measured rate is the background plus one third of the mean

velocity divergence inside the volume. By Gauss’s theorem that mean divergence depends only on the *mean interior contrast* $\bar{\delta}(< r)$, giving

$$\epsilon_{\text{Hole}}(< r) = \frac{\delta H}{H} = -\frac{f}{3} \bar{\delta}(< r). \quad (40)$$

This is the master formula, and its structure carries two lessons. The sign is unconditional: an under-dense neighborhood always reads high. And the lever is the *volume-averaged* contrast, not the depth at the center — the origin of the dilution worked in the next subsection.

The picture behind the algebra is worth holding. Gravity always pulls matter toward where matter already is, so an under-dense neighborhood is a valley everything is slowly leaving: the outflow of Eq. (39) is the emptying. An observer inside sees every tracer carried away by that outflow *on top of* the universal expansion and — having no local way to tell the two apart — books the sum as “the” expansion rate. The dilution is equally physical: the rate is read from tracers spread through the whole flow volume, and tracers beyond the valley’s rim feel none of its outflow, so a deep but small valley barely moves a measurement that averages far past it. Hence the two lessons above, and hence also the sharpest future test: cut away the low-redshift shells and the measured rate must walk back down to the background along the valley’s own profile.

Step 3 — the coefficient. The growth rate is itself an inventory output: $f = \Omega_m^{0.55} = (1/\pi)^{0.55} = 0.533$, so $f/3 = 0.178$. Even the lift’s coefficient carries no freedom: a mean contrast of -0.10 gives $+1.8\%$, -0.16 gives $+2.8\%$, and the full KBC depth -0.30 would give $+5.3\%$. (Linearity is the one approximation: at $|\delta| \approx 0.3$ the nonlinear correction is of order ten percent and acts to deepen the outflow, so the linear values are mildly conservative.)

33.3 What the ladder actually averages, and what the corrections remove

The formula wants $\bar{\delta}$ over the volume the measurement actually samples, and the ladder samples two very different zones: the calibrators live at 10–40 Mpc, where the Hole is deepest, but the rate itself is read from the supernova Hubble flow at roughly 100–600 Mpc. Because the average in Eq. (40) is volume-weighted ($r^2 dr$), a deep core contributes little once the flow volume extends past it: only readings in which the under-density persists across the flow volume produce percent-level lifts. Figure 17 makes this quantitative for three illustrative profiles spanning the published readings: at the flow scale the shallow profile yields $+1.7\%$, the canonical $+2.7$ – 2.9% , the KBC-deep $+5.2\%$.

The second dilution is methodological. The H0DN corrects its redshifts with the 2M++ peculiar-velocity model out to ~ 200 Mpc: whatever part of the Hole’s outflow the model captures is *already removed* from the published 73.50, and variant V19 — no corrections at all — shifts the consensus by only 0.6%, the network’s own measure of the modeled flow’s total effect. The lift that survives into the residual bookkeeping is therefore only the *unmodeled* monopole: the part of the under-density the velocity model does not trace, plus whatever extends beyond its volume. This is precisely why the framework grades the response law Derived and the depth as other people’s contested measurement, and why no band is adopted: the lift that survives into the bookkeeping is whatever the unmodeled monopole turns out to be, and the surveys of this decade will read it.

33.4 The scenario ledger and the budget

Table 8 assembles the readings into a decision ledger: each scenario’s mean contrast, the lift it implies, where it carries the background, and — the column that matters for the next section — the common-mode calibration residual each scenario still *requires* to reach the measured 73.50. Table 9 is the lift’s own systematic budget in the reproducibility format.

Table 8: The scenario ledger for a local-void lift, recorded and not carried. Each row: a published depth, the mean interior contrast over the supernova flow volume, the lift the master formula implies (Eq. (40), profile-weighted as in Figure 17), where it would carry the derived 66.8, and the common-mode calibration systematic still required to reach H0DN’s 73.50. No row is adopted: the response law is derived, the depth is an observational input on which the counts and flows literatures disagree, the counts branch is itself an $\approx 11\sigma$ Gaussian fluctuation on the σ_δ of §33.1 (charged in print), and the flows branch leaves at most a percent. The first row is the reading the parent GG-A10-Cosmos carries and this paper now carries with it. The KBC-deep row is retained for bookkeeping (field-level velocity modeling disfavors it); the shell test (F-Hubble-4, as re-registered) reads the profile directly, and a step consistent with zero confirms the first row.

Scenario	$\bar{\delta}$ (flow vol.)	Lift	H_0 lifted	Requirement
Not carried (GG-A10-Cosmos)	0	0	66.8	+10.0%
Pessimistic (corrections absorb most)	modeled away	+0.5–1%	67.3	+9.2%
Shallow counts	−0.10	+1.7%	67.9	+8.2%
Counts branch (recorded)	−0.16	+2.7–2.9%	68.7	+7.0%
KBC-deep (disfavored)	−0.30	+5.2%	70.3	+4.6%

Table 9: Systematic budget of the void channel itself (reproducibility format): each input, its range across the published readings, and its effect on the lift. The first row dominates the rest by a factor of three, and it is also the row the published counts disagree about most. The last two rows are included because a budget that lists only its large terms invites the reader to wonder what was left out. Every entry is a range rather than a value, which is the honest form for a quantity nobody has yet measured on the ladder’s own weight function.

Input	Range / justification	Effect on lift
Depth δ	−0.10 to −0.30 across count surveys (§33.1)	dominant; factor ~ 3
Extent vs. flow volume	core-only vs. 300 Mpc-wide readings; r^2 -weighted dilution	up to factor ~ 2 down
Flow-correction overlap	2M++ removes the modeled monopole (V19: total modeled effect 0.6%)	−50 to −100% of the modeled part
Nonlinearity	$ \delta \approx 0.3$ marginal for linear theory	$+\mathcal{O}(10\%)$, known sign
Growth rate f	inventory-fixed, $\Omega_m^{0.55} = 0.533$	$< 1\%$

What would settle the value is measurable this decade: a direct density–velocity cross-correlation over a $\gtrsim 1$ Gpc volume; full-sky near-infrared counts at Euclid/Roman depth, which will fix $\delta(r)$ outright; and the sharpest test of all, the rate measured in redshift shells — if the lift is real, $H_0(z_{\min})$ must decay toward the background as the low-redshift shells are cut away, at exactly the profile Figure 17 draws. And this signature is no longer merely prospective: an uncorrelated-bin




Figure 17: The lift a void would give: derived response, unobserved depth. *Left:* three illustrative under-density profiles spanning the published readings — shallow counts ($\delta \approx -0.10$), canonical (-0.16), and the contested KBC-deep (-0.30 ; Keenan et al. 2013) — against the ladder’s two sampling zones (calibrators, green; supernova Hubble flow, orange) and the edge of the 2M++ flow-correction volume (dotted). *Right:* the lift each profile implies for a rate measured within radius r , $\epsilon_{\text{Hole}} = -(f/3) \bar{\delta}(< r)$ with the inventory-fixed $f = 0.533$; no band is adopted, and the three published readings are drawn so the surveys can be read against them. The volume-weighted average is what the formula pays: only readings in which the under-density persists across the flow volume produce percent-level lifts, and the KBC-deep curve — which would carry the background to 70.3 — requires a contrast that standard structure formation disfavors at $\sim 20\sigma$ on that scale.

analysis of the DESI baryon-acoustic data finds the inferred H_0 decreasing with redshift at 6.4σ (Jia et al., 2025) — at low redshift, the shape a local void predicts, because the void’s outflow inflates only the shells that sample it, so the inferred rate must fall as those shells are cut away, with an amplitude that tracks the depth (Mazurenko et al., 2025). Three provenance cautions keep the entry honest — and the third cuts against us. First, the 6.4σ is a statement about the BAO-inferred rate in redshift bins, not the ladder-shell measurement the kill test above names — supporting evidence for the signature’s existence, not the test’s execution. Second, the deep-void amplitude the same modeling authors favor is precisely the reading the velocity data now disfavor (§33.1). Third, and against interest: a $\lesssim 300$ Mpc void lives at $z \lesssim 0.1$, while DESI’s bins reach $z \approx 0.3$ – 2.3 — a decline that *persists* at those redshifts is not a void signature at all but a background-evolution one, exactly the $w(z)$ reading this paper stakes its life against at F-Hubble-5. The trend would support a lift only insofar as it concentrates at low z ; its high- z tail, if real, is adverse. The paper therefore records the observed trend as consistent with some local outflow at low redshift; it adopts no lift on its strength, and a persistent high- z tail would be adverse at F-Hubble-5. None of these requires a sound horizon or a gravitational wave; whether any share of the ladder’s excess is environmental is the one question in this paper’s reading that ordinary galaxy surveys can close by themselves — and the framework’s prediction is that the share is zero.

34 The systematic against the error budget: where can ten percent hide?

The predicted systematic is this paper’s soft underbelly, and the honest response is not to look away but to quantify it from both sides. The H0DN quotes a total uncertainty of 1.09%; the reading of this paper requires a distance-calibration bias nine times larger. When the channel accounting of this section was first assembled, the candidate list was a table of small allowances — honest, and empty-handed: no published mechanism stood above the percent class. The 2025–2026 literature has changed the inventory. Two medium-size, physically motivated, common-mode channels have been proposed, sized, and contested in print, and the residual now has *named candidate anatomy*. This section states the requirement exactly, explains why the network’s error budget cannot exclude it, presents the two medium channels with their supporting *and* opposing evidence, walks the remaining channels one by one, and reports the verdict of the arithmetic — in both directions, and under both readings of the contested channel. One grading rule governs everything below, stated once: the external channels are *cited allowances*, *never GG-Theory results*. They enter no derivation, they move no grade in the honesty triage, and if every one of them dissolves under its kill test, nothing upstream of this section changes by a digit.

34.1 The requirement, stated exactly

Define the required common-mode systematic as what a void of depth $\bar{\delta}$, lifting the derived rate by ϵ_{Hole} , would leave for the calibration chain to supply:

$$R(\epsilon_{\text{Hole}}) \equiv \frac{73.50}{66.8(1 + \epsilon_{\text{Hole}})} - 1, \quad R(0) = +10.0\%. \quad (41)$$

With no lift carried the requirement is the whole excess, $R = +10.0\%$; a void of the depth the counts literature reports would reduce it to $+7.0\%$, and the surveys have not supplied that depth (§33). Throughout this paper the excess is quoted by the shorthand “ $\approx 10\%$ ” (the parent GG-A10-Cosmos prints “ $\sim 9\%$ ” against ≈ 67 ; the two are one number); the exact value on the derived rate is the $+10.0\%$ above, and every quantitative statement in this section works with it. An earlier version of this section worked with $+7.3\%$, the requirement after a $+2.8\%$ lift; the lift is withdrawn, and the arithmetic below is rerun at the whole excess.

34.2 What the 1.09% contains, and what it structurally cannot

The network’s quoted uncertainty is the quadrature of its statistical terms and every systematic its covariance matrix *models* — and the modeling is the H0DN’s great advance: shared anchors, shared photometric systems, and shared calibrator sets are tracked across paths, which is why the thirty variants scatter by only $0.40 \text{ km s}^{-1} \text{ Mpc}^{-1}$. But a covariance matrix can only price what is written into it. A *common-mode* term — one that moves every anchor, every rung, and every path together — leaves the variant catalog unmoved by construction: variants that share calibration DNA move together, so their stability tests internal consistency, not external truth. The requirement R is

a hypothesis of exactly this type. It cannot be excluded by the 1.09%; it can only be confronted channel by channel, from outside the network’s own bookkeeping.

34.3 The two medium channels: dust, and the statistics of the ladder itself

The two channels below are the ones that changed the arithmetic of this section. Both are common-mode in precisely the sense just defined: each moves every anchor, rung, and path together, which is why neither is priced inside the 1.09% and why neither is betrayed by the thirty variants. Both were identified by others, in the field’s own refereed literature; both are contested; and both carry kill tests with dates. They are presented with their opposing evidence in the same breath as their supporting evidence, because that is what Asserted-external means.

Channel 1 — calibrator-versus-flow dust, and the two-population standardization (allowance 2–4%). A Type-Ia supernova is standardizable, not standard: its distance is inferred only after corrections for light-curve shape and color, and the color correction encodes a dust law. The ladder’s calibrator supernovae do not live in average galaxies. A Cepheid is a young star, so a calibrator host is necessarily a star-forming, late-type spiral — and the supernovae found in such hosts are systematically young, high-stretch events seen through Milky-Way-like dust ($R_B \approx 4$). The Hubble-flow training set, by contrast, is dominated by the general galaxy population, and the mass-binned dust model trained on it applies a substantially lower effective $R_B \sim 3$ to exactly those calibrator-like events. The mismatch is therefore not scatter but a coherent offset with a definite sign: correct Milky-Way dust with a flow-trained law and every calibrator distance comes out short, every rung above inherits the error, and H_0 reads high. The intuition is a photographer’s: the calibrators were photographed through a redder filter than the one the calibration assumed, and every exposure downstream inherits the color cast — no amount of cross-checking between photographs taken through the *same* filter can reveal it. Quantitatively: imposing a consistent Milky-Way-like extinction model on the calibrator galaxies lowers H_0 by $\approx 2.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ — about half the gap (Wojtak and Hjorth, 2024) — and the full two-population hierarchical reanalysis of the Pantheon+ ladder returns $H_0 = 70.59 \pm 1.15$ (Wojtak and Hjorth, 2025) — with the same model’s most conservative, equal-prior variant reading 71.45 ± 1.03 , a spread the channel carries as part of its own honesty: the correction is real in every variant, and its size moves by almost a full unit with the prior on how different the two populations are allowed to be. (That reddening distributions differ between samples is itself demonstrable in simulation: imposing a universal extinction prior on intrinsically diverse hosts provably biases the recovered relative distances (Hallgren et al., 2025) — the statistical face of the same physics.) An independent cross-check sidesteps dust altogether: a ladder built only from intrinsically blue, low-reddening supernovae lands at 70.0 ± 2.1 (Gall et al., 2026). The channel hits Cepheid-, TRGB-, and JAGB-calibrated ladders identically (all share the supernova rungs), and it survives the two tests that killed earlier suspects: it is untouched by the JWST crowding verdict (it is not a photometry error), and it survives the network’s own V13 no-supernova variant — V13 carries the flow with SBF, fundamental-plane, and maser tracers instead of supernovae, so it tests whether supernovae are a *private* culprit, not whether a standardization term shared through the calibrator scale biases the flow intercept every path uses. The opposing

evidence is stated in the same breath: the Dark Energy Survey’s five-year systematics analysis argues, in the body of its error budget, that part of the host-mass step is intrinsic to the supernova populations rather than dust-driven (Vincenzi et al., 2024); no SH0ES reanalysis has reproduced the -2.5 ; and the blue-supernova calibrator sample is small. (The fitter-level residue beyond the two-population term — SALT2 versus SALT3 — is separately bounded below 0.5% and is not double-counted here.) The kill test is dated and specific: dust nearly vanishes in the near-infrared, so a NIR/JWST supernova ladder still converging at 73 kills this channel outright, and the large ZTF/Rubin siblings samples now accumulating will test the two-population split directly. The registry carries exactly this test as F-Hubble-8 (§56).

Channel 2 — the Bayesian distance-prior and selection statistics of the ladder itself (allowance 0–3%). The second channel is not in the photometry at all; it is in the inference. Every Bayesian regression of the ladder places a prior on the distance moduli, and the innocuous-looking uniform prior is not innocent: real objects are distributed uniformly in *volume*, not in distance modulus, so a uniform-in-modulus prior under-weights the far side of every error bar — distances bias low, H_0 biases high, rung after rung, in the direction that helps the tension. This is an Eddington-class bias, and it afflicts every Bayesian ladder, TRGB included. The intuition is a fisherman’s: cast a net of symmetric errors into a sea whose fish multiply with distance as $r^2 dr$, and the catch is biased toward the near side unless the census of the sea is built into the inference; the proposal’s charge is that the ladder’s regressions have been fishing with a flat sea. The pathology is formalized by Desmond et al. (2026) — whose own authors concede that the ladder’s simulation-based bias corrections likely already capture part of the effect, a concession printed here because Asserted-external means quoting the proposal’s caveats with its claims — and the recalibration with physically motivated priors has been executed: it shifts $73.0 \rightarrow 70.6 \pm 1.0$ and the quoted tension from $\sim 5\sigma$ to $\sim 2\sigma$ (Högås and Mörtzell, 2026). But this channel, alone on the list, already carries its rebuttal in print — from an author list overlapping the proposal’s own: a forward model of the Milky Way Cepheid selection function finds that the uniform-in-volume prior *without* a selection model manufactures a counter-bias of its own (~ 0.05 mag in the zero-point) and reinforces the published H_0 (Stiskalek et al., 2026). The two analyses disagree about which incomplete inference is less wrong — which is exactly the state “undecided” names. The status is therefore printed as it stands: *proposed, executed, and rebutted — undecided*; the allowance is carried as a range whose low edge is zero, 0–3%; and both readings are priced in the verdict below. The kill test is cheap, decisive, and — remarkably — not yet published: injection–recovery simulations of the full ladder pipeline on synthetic data with known truth. Whoever runs it settles the channel either way; the registry carries the test as F-Hubble-9 (§56).

The overlap caveat. The two channels are not perfectly independent: the prior recalibration re-treats some of the same supernova terms the dust channel corrects, so naive addition double-counts at the $\sim 1\%$ level, and no joint analysis of the two exists. The arithmetic below therefore combines channels in quadrature, never linearly, and quotes the aligned stack only as the worst-case envelope it has always been.

34.4 The remaining channels, one by one

The remaining channels are the percent-class terms the ladder literature has audited for years; the list and its sizes draw on the ladder-critical literature (Efstathiou, 2020; Freedman, 2021; Mörtzell et al., 2022) as much as on the teams’ own audits. Before walking them, the position of GG-Theory on this list should be stated once, because it governs every entry. The framework has already *closed* the channels it owns: a new baryon-coupled force is closed by the boundary-condition termination of GG-A3-B16 and, independently, excluded by eight orders of magnitude; the void channel is priced and not carried; and the coarse-graining term is computed and bounded below 10^{-3} (§35). What remain are instrumental and astrophysical terms — metrology — and about metrology a theory of the constants is deliberately silent: it neither derives a parallax zero-point nor is threatened by one. The framework’s contribution to this subsection is therefore *negative and disciplined*: having removed every physics escape route, it earns the right to say that if the residual exists, it lives below, in the chain of calibrations — and it accepts the referee that can say whether it does.

Primary-indicator family and calibrator subsampling (1–2%). The famous SH0ES–CCHP split has not converged; it has *localized*. On JWST data the CCHP program’s published tip-of-the-red-giant-branch value stands at 70.39 ± 1.94 , with its JAGB and Cepheid variants spanning 67.8–72.1 (Freedman et al., 2025); the SH0ES reanalysis of the same JWST fields returns 72.6 ± 2.0 , arguing that the early JWST subsample is statistically low — the full-sample prediction for those same galaxies is 72.8 — and that the difference is sample selection, not measurement (Riess et al., 2024b). Meanwhile the tip’s measured contrast varies field to field at the 0.05–0.1 mag level unless explicitly standardized — a 5σ effect, and a percent-class lever sitting inside a convention (Wu et al., 2023). Both sides’ arithmetic is internally consistent; what they disagree about is which subsample of calibrators, and which standardization convention, the ladder should stand on — a genuinely open dispute located in choices, not in anchors and not in crowding. Allowance 1–2%; the kill test is the complete JWST coverage of all ~ 37 SH0ES hosts, in progress now, plus an agreed tip/JAGB convention — the registry carries it as F-Hubble-10 (§56).

Gaia parallax zero-point (0.5–1%). The floor of the ladder: every Milky-Way Cepheid distance — and through the calibrators, every rung above — inherits Gaia’s astrometric zero-point, so an offset here is the purest possible common mode, multiplied silently down the whole chain. The EDR3 offset ($-17 \mu\text{as}$) was mapped and removed only by a dedicated recalibration (Lindgren et al., 2021), and the bright-star ($G < 8$) regime is genuinely unsettled: an independent orbital-parallax check finds a bright-end zero-point of $-38.9 \pm 10.3 \mu\text{as}$ (Ding et al., 2025) — with the formal errors themselves underestimated by a factor of ~ 2 — and no established sign for the H_0 shift such an offset would induce once propagated through the anchor fit. Two facts size the allowance at the low edge: the direction of the bias is itself unestablished (an allowance in “the helpful direction” cannot borrow more than the uncertainty grants), and the cluster-Cepheid route — a 0.9% Milky-Way calibration at faint G , where the zero-point is well behaved, agreeing with the ladder’s at 0.3σ (Cruz Reyes et al., 2022) — already bypasses the bright regime. The entry now carries a date: Gaia DR4, due December 2026, improves the parallaxes by the $\sqrt{t}$ factor ($\times 1.4$ on the 66-versus-34-month

baseline; it is the proper motions that improve as $t^{3/2}$, $\times 2.7$) and — the part that matters for the ladder — attacks the bright-star ($G \lesssim 8$) systematics floor with full epoch astrometry. It is the nearest-term decisive test of the entire anchor rung. In GG-Theory terms: purely instrumental; nothing in the ledger touches astrometry.

Cepheid metallicity term ($\sim 0.5\%$). The period–luminosity zero-point depends on the stars’ chemical composition, and the anchor Cepheids do not share the metallicity of the calibrator-host Cepheids; a mis-calibrated slope biases every host distance in the same direction. The slope has substantially converged (measured as significant and, decisively, wavelength-independent from 0.5 to 4.5 μm (Breuval et al., 2022) — a chromatic test a photometric artifact would fail), leaving roughly half a percent of leverage. For the framework this is stellar astrophysics — envelope opacities, not geometry; its one relevant statement is negative: the α_B bound guarantees no new force hides inside the pulsation physics.

Photometric crowding ($\lesssim 0.3\%$). Unresolved neighbors brighten Cepheid photometry in crowded fields, shortening inferred distances and pushing H_0 up — historically the favorite suspect, and the one channel with a completed kill test: JWST re-photometry of the same fields rejects crowding as the explanation of the tension at 8σ (Riess et al., 2024). The allowance honors that verdict. No GG-Theory content enters; the empirical test already did the work a theory would otherwise owe.

Cross-instrument photometry ($\sim 0.5\%$). The ladder’s photometry is stitched across HST, JWST, and ground-based systems, and each stitch carries a zero-point. The HST $\leftrightarrow$ JWST leg is now cross-checked directly (Riess et al., 2024); the allowance covers the remaining seams. Instrumental; no GG-Theory content.

Peculiar-flow modeling ($\sim 0.5\%$). The one channel adjacent to something the framework *does* derive. The network corrects its redshifts with the 2M++ flow model, and V19 (no corrections) bounds the total modeled effect at 0.6%. The bookkeeping guard matters here: a local-void lift, if any, would live in the *unmodeled* monopole — the part the flow model does not trace — and none is booked (§33), so this allowance covers errors in what 2M++ *did* model and nothing is counted twice. This is the seam where the framework’s environmental physics and the ladder’s metrology meet, and the ledger keeps them on opposite sides of the line.

34.5 The verdict of the arithmetic, in both directions

Read the tables together, and read them twice — once with the contested statistics channel at its proposal’s central value, once with the rebuttal taken seriously and the channel halved. Under the proposal-central reading the eight channels quadrature to 4.3% and stack to 9.3%; under the post-rebuttal conservative reading, 3.8% and 8.3%. Against these the requirement is $R = +10.0\%$, and the comparison is the section’s headline, printed against interest: in *both* scenarios the requirement sits *above* the fully aligned stack. The published channels, taken at their central values and stacked with every correlation assumed maximal, reach 9.3% (8.3% post-rebuttal) and do not close it. An

Table 10: Candidate common-mode calibration channels for the named residual: nature, published status, and the allowance adopted here (“allowance” = plausible undetected common-mode contribution in the helpful direction, not a claimed error; all channels Asserted-external — cited from the field’s literature, never GG-Theory results). The contested statistics channel is priced twice: at its proposal’s central value and at a post-rebuttal conservative value; the verdict (§34.5) carries both scenarios.

Channel	Nature and published status	Central value used	
		<i>proposal</i>	<i>post-rebuttal</i>
Calibrator-vs-flow dust / two-population	coherent R_B mismatch; 70.59 ± 1.15 reanalysis; blue-SN check 70.0 ± 2.1 ; DES counterweight; range 2–4%	3.0%	3.0%
Distance-prior / selection statistics	proposed and executed ($73.0 \rightarrow 70.6 \pm 1.0$); rebutted (MW-Cepheid forward model); undecided; range 0–3%	2.5%	1.5%
Indicator family / calibrator subsample	CCHP 70.39 ± 1.94 vs SH0ES 72.6 ± 2.0 on JWST data; tip-contrast convention; range 1–2%	1.5%	1.5%
Gaia parallax zero-point	bright-end zero-point $-38.9 \pm 10.3 \mu\text{as}$, sign of the H_0 shift unestablished; DR4 due Dec 2, 2026; range 0.5–1%	0.5%	0.5%
Cepheid metallicity term	slope significant, wavelength-independent $0.5\text{--}4.5 \mu\text{m}$; residual leverage	0.5%	0.5%
Photometric crowding	rejected as the explanation at 8σ	0.3%	0.3%
Cross-instrument photometry	HST $\leftrightarrow$ JWST $\leftrightarrow$ ground seams	0.5%	0.5%
Peculiar-flow modeling	V19 bounds the modeled effect; unmodeled monopole not booked (void channel not carried)	0.5%	0.5%
Quadrature (independent)		4.3%	3.8%
Aligned linear stack (worst case)		9.3%	8.3%

earlier version of this section, crediting a +2.8% void lift, found the requirement *between* quadrature and stack and called the geometry of a two-maxima conspiracy “gone”; with the lift withdrawn the geometry is back, in a milder form — not a deep void and a conspiracy, but a full alignment of every published channel plus roughly another percent that no published mechanism supplies. We retire the temptation to dress this in sigmas: the channel entries are one-sided plausibility allowances, not Gaussian half-widths, and dividing a requirement by their quadrature manufactures a statistic that does not exist. The honest statement is the comparison itself — the requirement sits above the range the published channels’ own arithmetic spans — read as a *consistency* statement, not a probability. The residual’s candidate anatomy is real, published, and dated, and it is not yet large enough. This paper prints that as plainly as it printed the 11σ void: the framework’s claim is that the excess lives in the ladder, not that the field has found all of it.

That last statement deserves its own arithmetic, and the arithmetic deserves two honesty clauses that this literature — ourselves included — has blurred. First, the two medium corrections are *alternative whole-ladder reanalyses of the same Pantheon+ supernovae*, each landing near 70.6 on its own; chaining both ($73.50 \times 0.970 \times 0.975 \approx 69.5$) assumes a joint analysis that does not exist, so the chained point is quoted only as the correlated floor of the range. Second, both corrections live in supernova standardization, so they move the consensus only through its SN-weighted share: the network’s own SN-free variant V13 stands at 73.43 ± 1.80 as the printed counter-constraint, and a consensus-level shift of the full -3% would require the non-SN tracers to move with it. With both clauses charged, the published recalibrations carry the ladder to $\approx 70.6\text{--}72$ (either reanalysis alone, scoped to the SN share, overlap unquantified), with $69.5\text{--}70.2$ as the unproven fully-chained floor. The corrected ladder then lands $2.7\text{--}3.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$ above the derived 66.8 — roughly three standard deviations of the published channels’ own $\pm 1.0\text{--}1.15$ — and 3.8 above at the single-reanalysis 70.6. No conspiracy is required for the corrections that exist; what is required, and not yet published, is roughly another four to five percent of common mode. The middle band into which the corrections drain is real; it is not yet the derived rate. And the band is not this paper’s invention: the CCHP JWST program reads 70.39 ± 1.94 on its own, and — as Part V’s landscape audit tabulates (Table 13; drawn in Figure 20) — the fast-radio-burst, dark-siren, chronometer, and ladder-free meta-average determinations independently cluster in the same 69–71 neighborhood. An earlier version of this paragraph called that band “the band the derived background plus its lift has predicted from the geometry side all along”; with the lift withdrawn the sentence is false and is deleted. The framework predicts the corrections to drain all the way to 66.8; the surveys will say whether they do.

Read in the other direction, the deflation arithmetic is equally mechanical: a common-mode term of the central quadrature size (4.3%) added to the ladder’s budget collapses the headline 7.1σ to $\approx 1.9\sigma$ (conservative: 3.8%, $\approx 2.1\sigma$; a common mode of the full 10% would take it to 0.8σ ; §16) — which is why the tension’s statistical grandeur and the residual hypothesis coexist without contradiction: significance computed inside a budget is conditional on the budget’s completeness, and the completeness is precisely what the two medium channels dispute.

Four honesty clauses complete the picture, one per exposure. *First, the status.* Both medium channels are young and contested: the dust channel awaits a quantitative SH0ES rebuttal, and the

statistics channel already carries one — from authors overlapping the proposal’s own team — which is why it is priced with a zero-edge range and the verdict is run in both scenarios. Nothing here is settled; everything here is dated. *Second, the overlap.* The two channels partially re-treat the same supernova terms; the quadrature convention absorbs the known $\sim 1\%$ double-count, and a joint analysis — which does not yet exist — is the clean fix. *Third, the counterweight.* The strongest evidence *against* any pure-ladder-bias reading is not in the ladder at all: one-step geometric probes — megamasers, Type-II supernovae, the surface-brightness route, time-delay lensing — read 72–75 without Type-Ia standardization and largely without Cepheids, and neither medium channel touches them. Each carries 3–5% errors and its own soft systematic, and none is individually decisive — but together they are the standing objection to this section’s hypothesis, and Part V confronts them head-on rather than in a footnote. And the counterweight is not only external: the ladder’s own robustness tests — the background-free Perfect-Host field at 73.49 ± 0.93 (Riess et al., 2025), the standardized-TRGB grid whose 108 analysis variants never fall below 71.6 (Scolnic et al., 2023), and the 0.994 ± 0.010 linearity constraint (Riess et al., 2024b) — argue that whatever moves the ladder is not a photometric error. The medium channels are consistent with exactly that verdict: neither is photometric — one lives in a dust law, the other in a prior — which is why both survive the tests that killed the photometric suspects, and why their kill tests had to be designed anew. *Fourth, symmetry.* Everything in this section cuts both ways, and more sharply than before: a bias story with named channels predicts that the ladder-free siren lands at 66.8, while a real-73 story predicts 73 — so the siren now discriminates between *named* alternatives, not between a theory and a shrug. If it converges at 73, the anatomy evaporates, the residual hypothesis dies — and so does the derived rate, and this paper with it (F-Hubble-2). The registry carries the three channels’ own kill tests as F-Hubble-8, -9, and -10 — the NIR ladder, the injection–recovery simulation, and the complete JWST sample — each two-sided by construction: the opposite outcome confirms the channel and collapses the tension toward the derived value (§56). Figure 18 draws the whole account.

35 The limit of backreaction as a remedy

The environmental channel of §33 is not carried; the honest question is whether any of the excess can come from anywhere *within* the framework rather than from the measurement chain. The most natural candidate is native to GG-Theory: if observation is a coarse-grained boundary projection (GDOS), then the expansion rate a local observer infers is a volume-averaged quantity, and volume-averaging an inhomogeneous universe need not reproduce the smooth background rate. This is the *backreaction* of Buchert’s spatial-averaging formalism (Buchert, 2000; reviewed in Buchert and Räsänen, 2012), in which the averaged expansion of a domain D carries a kinematic term

$$\mathcal{Q}_D = \frac{2}{3}(\langle\theta^2\rangle_D - \langle\theta\rangle_D^2) - 2\langle\sigma^2\rangle_D, \quad (42)$$

with θ the local expansion scalar and σ the shear. $\mathcal{Q}_D$ is precisely the object GDOS coarse-graining would source: it is the honest way a locally-measured $\langle H \rangle_D$ can exceed the background. The intuition: averaging does not commute with evolving — the average of many patches, each expanding at its own rate, is not the expansion of the average patch — and $\mathcal{Q}_D$ is the exact price of the swap: the




Figure 18: The bridge, drawn honestly. *Left:* the early side is closed — the derived background 66.8 plus the coherent π -offset is the CMB-inferred 67.4; nothing is added. *Center:* the late side — the hatched span from the derived 66.8 to H0DN’s 73.50 is the required common-mode systematic, +10.0%: predicted, bounded, and *not* derived. The violet bracket beside the hatched span is the systematic’s candidate anatomy (2026, external): dust 2–4%, distance priors 0–3% (proposed and rebutted, undecided), calibrator subsample 1–2% — contested allowances from the field’s own literature, never GG-Theory results, whose published central values carry the ladder to ≈ 69.5 –70.6 (dotted tick), still 2.7–3.8 above the derived rate; their kill tests are registered as F-Hubble-8/9/10 (§56). *Right:* the void channel (faded), priced and not carried: a counts-branch void would reduce the requirement to +7.0%, the disfavored KBC depth to +4.6% — recorded as scenarios, not claims. Hatched is the predicted systematic, violet is external, faded is priced and not carried: the figure is the paper’s honesty policy in one picture.

variance of the local expansion rates, minus the shear that collapse builds. In a universe like ours the shear term wins, which is why the honest computation comes out not only small but of the wrong sign.

We have evaluated it rigorously on the GG-Theory background (the derivation is Appendix H.4), and the result is a firm null. *First*, the systematic backreaction is bounded, $\langle Q_D \rangle / H^2 \lesssim 10^{-3}$, so the induced shift of the volume-averaged rate is $\lesssim 0.1\%$ and of the *wrong* sign (the shear term dominates) — consistent with the mainstream Λ CDM value $\sim 10^{-5}$, and unenhanced in GG-Theory because its structure formation is standard. *Second*, the cosmic variance of a locally measured H_0 — the *scatter* from sitting in a random over/under-density, $\sigma_H / H = (f/3) \sigma_\delta(R)$ — is only $\lesssim 0.5\%$ on the ladder’s ~ 100 –300 Mpc averaging scale (Marra et al., 2013). *Third*, to produce the full +10% excess from a pure under-density would require $\langle \delta \rangle \approx -0.56$ on that scale, a $> 12\sigma$ fluctuation; even the few percent the counts literature once credited would need a valley the same σ_δ prices at $\approx 11\sigma$ (§33.1), which is why neither is carried. (An earlier version of this sentence called that valley “ $\sim 2\sigma$ ”; it is not, on the paper’s own numbers, and the two statements are now the same statement.) These bounds are drawn in Figure 19, and they agree with the same statement made from the rival-model side in Part V.




Figure 19: Backreaction cannot supply the excess. Linear theory on the exact derived background ($\Omega_m = 1/\pi$; Eisenstein–Hu $P(k)$). *Left*: the cosmic variance of the locally measured H_0 , $\sigma_H/H = (f/3)\sigma_\delta(R)$, against averaging radius; on the ladder’s effective scale (shaded) it is below a percent, an order of magnitude under the +10.0% requirement (red). *Right*: the rarity $\nu = |\delta|/\sigma_\delta(R)$ of the under-density that would give the whole excess ($\langle\delta\rangle = -0.56$): 3.4 at 40 Mpc/h, 12σ at 100, 22σ at 150. Drawn at the parent’s $\sigma_8 = 0.80$ on the carried background ($h = 0.6676$, $n_s = 0.9634$); no void band is shaded because no lift is carried. The systematic Buchert term is $\langle Q_D \rangle/H^2 \lesssim 10^{-3}$ (Appendix H.4).

The one surviving loophole is honest to state and honest to bound. Only a fully *non-linear*, two-phase (void-versus-wall) reading of the kind Wiltshire’s timescape cosmology proposes (Wiltshire, 2007) claims percent-level differences between the statistically-averaged and locally-measured rate, through clock and lapse effects between bound structures and expanding voids. Timescape is philosophically the closest thing in the literature to a GDOS coarse-graining, and so it is the framework’s honest best hope for a larger effect — but it reaches its magnitude by *replacing* the standard FRW averaging with a specific bare/dressed computation, it is model-dependent and contested, and it does not cleanly turn a background 67 into a local 73. It is a door to a different cosmology, not a correction to the derived $H(t)$ of this paper. The conclusion therefore stands: the excess is *not* a cosmological-averaging effect. What the framework derives is the bound; the excess is a distance-ladder calibration question it predicts, and the ladder-free siren is its arbiter.

36 The why-now window: the coincidence as a factor-of-a-few timing

A reader may object that reading the present epoch at the $\Omega_m = 1/\pi$ crossing is itself a coincidence — why should “now” land where the matter and vacuum shares cross? The framework’s answer is quantitative. The one time-dependent ratio in the inventory is

$$\frac{\rho_\Lambda}{\rho_m(a)} = (\pi - 1) a^3 = (\pi - 1) e^{3(r-r_{\text{now}})}, \quad r = \ln \frac{E_{\text{GUT}}}{\mu} = r_{\text{now}} + \ln a, \quad (43)$$

with $r_{\text{now}} \simeq 66.6$ the present value of the radial/RG coordinate. The ratio climbs monotonically through unity, and the “order-unity” window in which it lies between 0.1 and 10 — the window in which the coincidence looks like a coincidence — has width

$$\Delta r = \frac{2 \ln 10}{3} = 1.535, \quad a \in [0.36, 1.67], \quad z \in [1.78, -0.40]. \quad (44)$$

Out of $r_{\text{now}} \simeq 66.6$ units of radial history (roughly 69 e-folds from the end of inflation), a window of width $\Delta r \simeq 1.5$ is about *two percent of the logarithmic history*. Living in it is therefore a factor-of-a-few timing coincidence, not a fine-tuning — the same order of “luck” as being alive during the few hundred million years a given star is on the main sequence. The present epoch is read at the $\Omega_m = 1/\pi$ crossing, 75 Myr before today at $H_0 = 66.8$ (a 0.5σ offset against the ± 160 Myr Planck timing blur), and that reading is the one declared input the geometry does not remove. The full form of the argument — the window drawn on the whole radial history, the crossing resolved from today, and the residual timing question stated as such — is the why-now chapter of the parent GG-A10-Cosmos; what this paper needs from it is only the price tag: one declared epoch, factor-of-a-few luck, no dial.

37 Why the derived reading is not an ad-hoc fit

A reader meeting the reading of Eq. (31) for the first time might ask whether its ingredients were chosen after the fact to reproduce the data. For the early side the answer is immediate and decisive: there is *no fitted number there at all*. The background rate is derived by the geometry alone (Route I, §24) and confirmed by dividing the CMB’s own model-independent densities (25) by the derived inventory (16); the CMB endpoint 67.4 is simply where that derived rate is read. Nothing on the early side is tuned, because nothing on the early side is free.

The late side carries no derived quantity and one honestly undetermined systematic, ϵ_{syst} , that the framework labels a distance-ladder calibration systematic rather than a parameter. Nothing on the late side is tuned either, because nothing on the late side is free: the one environmental channel the framework could have credited was computed — its sign unconditional, its coefficient fixed by the inventory — and found to need a depth the surveys do not show, so it was not carried rather than adjusted.

The 2026 anatomy of the residual strengthens this defense rather than straining it, for a reason worth stating plainly: every ingredient of the candidate anatomy was identified by *others*, in the field’s own refereed literature, by authors with no stake in this framework — one of the channels with the ladder’s own architects among the rebutting authors. The reading borrows no GG-Theory freedom whatsoever: the framework’s numbers — the derived 66.8 and the +10% requirement — were fixed by the geometry and printed before the channels were sized, and the channels were sized by people optimizing nothing of this paper’s. If the anatomy dissolves under its dated kill tests, the systematic reverts to a named blank and this section’s defense loses nothing; if it survives and grows to the requirement, the gap closes without a single GG-Theory dial having moved — because there is no dial.

37.1 The strong- CP contrast

The cleanest way to see that the reading is forced rather than fitted is to note what it does *not* need. One might try to reconcile a high background rate with the CMB by invoking a pre-recombination modification (extra radiation, an early-dark-energy lobe); the derived reading needs none of it: the microwave spectrum is fit at ≈ 67 by GG’s own fractions with the standard recombination history untouched. The only early-universe number the cohort predicts in this vicinity is a null: $\Delta N_{\text{eff}} \leq 0.03$ (GG-A10-Cosmos, codebook row D12), because below the lock scale the field content is the Standard Model’s and the Kaxion is not thermal — consistent with the Planck bound 0.04 ± 0.18 , and unrelated to the Hubble split. A future CMB bound tightening toward zero confirms it; a detection at $\Delta N_{\text{eff}} \gtrsim 0.2$ kills it. (An earlier version of this paper carried $\Delta N_{\text{eff}} = 0.04$ as a value inherited from the Standard-Model companion; that companion carries no such number, and the parent’s derived null replaces it.)

38 Cross-ledger coherence: why this is the hard test

One could always invent a theory that produces a single number matching observation. The interesting question for GG-Theory is whether the *same inherited ledger* that fixes $\sin^2 \theta_W$, $\hat{\alpha}^{-1}(M_Z)$, the 28 Standard-Model parameters, the DNA alphabet, the conscious-episode duration, the Kaxion line at 1.95 GHz, and the π -locked inventory $\Omega_m = 1/\pi$ *also* yields the right H_0 . It does, and by the tightest possible route: the inventory that fixes H_0 is the *same* inventory audited against the whole acoustic spectrum in GG-A10-Cosmos. If $\Omega_m = 1/\pi$ or $\Omega_b = 1/2\pi^2$ were tuned to fix H_0 , the CMB peak positions, the baryon-loading ratio, and the damping scale — all reproduced to sub-percent by CLASS — would break simultaneously. The agreement is therefore not a one-parameter victory but a consistency witness spanning the entire microwave sky.

Table 11: Cross-ledger coherence: the ingredients of the derived rate have sibling predictions in other sectors, so tuning any one to match H_0 would spoil the sibling. The table is the answer to the obvious objection, which is that a framework with enough structure can be tuned to any rate. Each ingredient of the derived value already carries a prediction somewhere else in the cohort, so moving it to improve the fit here breaks something there. The constraint is not that these numbers are hard to change; it is that none of them can be changed alone.

Ingredient	Role in H_0	Same source, other predictions
$\Omega_m = 1/\pi$, $\Omega_b = 1/2\pi^2$	divided into the CMB densities to give $H_0 = 66.6$, beside the geometric 66.8	the full acoustic spectrum (peaks, baryon loading, damping scale) fit at the same value; $\Omega_b/\Omega_m = 1/2\pi$ at every epoch
$M_{\text{lock}} \approx 3.85$ TeV	sets the seesaw cross-check $H_0 = 498.3\sqrt{\zeta_\Lambda}$	the KK-cluster collider scale; the meV vacuum floor; the whole scale chain of GG-A5-GDOS T8.7
Ladder systematic ϵ_{syst}	+10% on the local ladder, predicted	tested by every ladder-free probe (BAO, sirens, time delays), not by H_0 itself; no share of it is environmental (§33)

39 The S_8 tension: shared, not resolved

The same logic does *not* settle the companion S_8 tension, and the paper says so. Because GG makes no pre-recombination modification, the growth history is standard growth on the derived inventory — $\Omega_m = 1/\pi$, the derived tilt $n_s = 0.9634$, $w = -1$, three Standard-Model neutrinos at $\Sigma m_\nu \simeq 59$ meV — and the amplitude follows from the launch sector’s one declared input, $A_s = 2.10 \times 10^{-9}$, with nothing left to tune (GG-A10-Cosmos, its released `sigma8_gg.py`):

$$\sigma_8 = 0.80 \pm 0.01, \quad S_8^{\text{GG}} = \sigma_8 \sqrt{\Omega_m/0.3} = 0.825 \pm 0.010 \quad (\text{App. I.5}). \quad (45)$$

One number, not two branches. It sits 0.4σ below the microwave background’s 0.832 ± 0.013 and 2.5σ above the cosmic-shear $0.759^{+0.024}_{-0.021}$: the framework shares the S_8 tension with Λ CDM rather than resolving it. If the tension survives the next generation of lensing surveys, it is between the CMB-normalized amplitude and the shear-normalized one, and not between this framework and either. (An earlier version of this paper printed $S_8 = 0.759$ from a $\sigma_8 = 0.737$ attributed to the parent; the parent never computed that number, and its released chain returns 0.80.) The sign structure that remains on one background is summarized in Table 12.

Table 12: Sign structure of the growth reading on one background, as it now stands. One row is live: the inventory sets both the rate and the amplitude, and neither can be moved without the other. The local-void row an earlier version carried — raising the local rate and lowering local growth, so that the two tensions were read as one environmental effect with two signs — is withdrawn with the lift and kept for the record. The S_8 tension is shared, not resolved: the framework carries one amplitude and inherits the comparison.

Effect	on H_0	on S_8
Derived background (inventory, A_s)	sets 66.8	sets 0.825 ± 0.010
Local void (withdrawn: not carried)	—	—
Coherent π -offset (+1% on Ω_m)	−0.9% on h	+0.5%, inside the ± 0.010

The broad-brush growth relation used in the comparison Part below,

$$\frac{\Delta S_8}{S_8} \approx \frac{1}{2} \frac{\Delta \Omega_m}{\Omega_m}, \quad (46)$$

carries no early-universe correction and, on one background, no late-time environmental term either.

39.1 The sign structure, as it now stands

The reason S_8 belongs in the same discussion as H_0 is that they draw on the same inherited ledger: the amplitude and the rate are two shadows of one inventory, so the growth sector is a cross-consistency audit, not a free knob — σ_8 , f , Ω_m , and H_0 move together or not at all. An earlier version of this subsection quantified a sign structure in which a derived 0.759 landed on the

lensing value and called the agreement a success of the inventory; that number had no generator, and the reading is withdrawn with it. The honesty here is essential and we state it once, plainly: the computed $S_8 = 0.825 \pm 0.010$ does not land on the shear value, the early-inference-versus-late-measurement gap is not explained by this framework, and the paper carries the tension as shared rather than reading it, as it once did, the same way as the local H_0 excess. The same discipline governs both tensions: one ledger, one honest account of what it has and has not reached.

40 Revisiting Λ CDM in the light of a derived $H(t)$

There is a larger consequence to draw, and we state it plainly, as an invitation. Λ CDM 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-Theory proposes that they are not free: that the densities, and therefore the entire expansion history $H(t)$ of Figure 13, are derived from one six-dimensional geometry with no continuous dial. If that proposal holds, the parameters of Λ CDM 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 Λ CDM — GG-Theory *runs on* its Friedmann/FRW machinery, and its empirical fit is not in question. We are asking the community to revisit it 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 $H_0^{\text{GW}} \rightarrow 67$, the exact $w = -1$, the tensor ratio $r = 0.0038$, and the coherent π -offset across the acoustic peaks. The strongest way to settle the tension is not to manufacture the residual but to *over-determine* the derived curve: direct $H(z)$ from cosmic chronometers and baryon acoustic oscillations, redshift drift, and standard sirens all measure $H(t)$ without the sound horizon or the distance ladder, and if they track the 66.8-anchored curve the local 73 is confirmed as a local excess on an otherwise-correct history. The honest banner is not “ Λ CDM is wrong” but “ Λ CDM 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.

41 Summary: the reconciliation, and the road to Part V

Part IV set out to do the hard half of the tension, and this closing section gathers what it established, ending where Part V must begin.

The tension changed character at the outset (§30): with the background rate derived, it is no longer a contradiction between two measurements of an unknown number but one known number, and the question of which instruments read it. The early dictionary needs nothing: the derived 66.8

plus the inherited sound-horizon calibration *is* the CMB value, the ruler lands on the measured sound horizon, and no pre-recombination physics is invoked because none is missing. The switch at $z_* = e^7$, fixed by the integer $J = 7$ and landing 0.6% from the measured decoupling, marks where the dictionaries change and does not license two answers: every ladder-free late probe is predicted to read the same rate, and only the calibrated ladder is displaced (§31).

The late side then carries one piece (§32), and the chain is one step: 66.8 (derived), plus a predicted common-mode systematic of +10.0% $\rightarrow$ 73.5. The one step carries a name and an open verdict.

The void channel is priced, not asserted (§33): the under-density is observed and its depth is contested, the outflow follows from continuity, the master formula $\epsilon = -(f/3)\bar{\delta}$ has an unconditional sign and an inventory-fixed coefficient, and the two dilutions — the ladder’s volume-weighted averaging and the flow corrections already applied — are worked explicitly. The depth that would let it carry even +2–3% is an $\approx 11\sigma$ fluctuation on the framework’s own power spectrum that the velocity data disfavor, which is why the paper carries no lift and records the scenarios as what the surveys can measure.

The residual is priced from both directions, and the pricing has sharpened (§34). The 2025–26 literature has given it candidate anatomy: two published medium-size common-mode channels — calibrator dust (2–4%) and the ladder’s own Bayesian distance priors (0–3%; proposed, executed, and rebutted — undecided) — joined to the localized CCHP subsample dispute (1–2%) and the ordinary small terms. Under the proposal-central reading the channels quadrature to 4.3% and stack to 9.3%; under the post-rebuttal conservative reading, 3.8% and 8.3%; in both, the requirement +10.0% sits *above* the aligned stack — the published anatomy is real and not yet sufficient, and the paper says so — and the published central values carry 73.5 to ≈ 70.6 (each reanalysis on its own; 69.5–70.2 only if fully chained), still 2.7–3.8 above the derived rate. Read in the other direction, a common mode of the quadrature size deflates the headline 7.1σ to ≈ 1.9 – 2.1σ , which is why the tension’s statistical grandeur and the residual hypothesis coexist without contradiction. The standing counterweight is carried openly: the one-step geometric probes stay high without supernova standardization, and Part V confronts them head-on. And in the light of GG-Theory the channel list retains its clean structure: the framework closes the channels it owns — the $U(1)_B$ vector by the boundary-condition termination of GG-A3-B16 and, independently, by eight orders of magnitude; the void channel, priced and not carried; the backreaction term below 10^{-3} with the wrong sign (§35) — and is deliberately silent on the metrology it does not, having earned, by that closing, the right to locate the residual in the measurement chain; the chain’s own literature has now begun to name the suspects, and every suspect carries a dated kill test reserved in the registry (§56).

The wider audit closes the Part. The why-now question is priced, not hidden: one declared epoch, a two-percent logarithmic window, factor-of-a-few luck (§36). The reading is shown not to be a fit: the early side contains no free number at all, and the late side carries no derived ingredient and no dial (§37). The same ledger is on the hook everywhere at once — tune any ingredient to help H_0 and the acoustic spectrum, the collider stake, or the meV vacuum scale breaks instead (§38, Table 11) — and the S_8 companion tension is not resolved: the computed 0.825 shares it, 0.4σ

below the microwave background and 2.5σ above cosmic shear, and the paper says so (§39). The Part ends with an invitation rather than a verdict: revisit Λ CDM not as wrong but as unexplained, and over-determine the derived curve with every ladder-free, horizon-free probe available (§40).

What remains is evidence, competition, and verdict. Part V first lays the entire 2026 measurement landscape on the table — all forty-five audited determinations, drawn and tabulated — and then lines up the rival explanations — the named-systematics reading, now a hypothesis with its own anatomy rather than a shrug, early dark energy, extra radiation, modified gravity, local voids — against the same derived rate and the same kill switches, and asks which reading survives its own falsifiers. Part VI then hands the decision to the channel this Part has invoked at every turn: the gravitational-wave standard siren, which needs neither a sound horizon nor a ladder, and which this paper commits to in advance — ≈ 67 and the reading stands; ≈ 73 and it dies (F-Hubble-2). The reconciliation is done; the adversarial test begins.

42 Part IV in plain English: why the local ladder reads high, and what the framework does and does not claim about it

The derived curve lands on the early number, so all that remains is to say why the local measurement reads about ten percent above it. This Part answers with a prediction rather than a sum, and the honest part of the answer is that the prediction is about *where* the ten percent lives, not how large it is.

There is one term, and it is not derived. On one background every instrument that reads the expansion without a calibrated rung — the sound horizon, the acoustic scale in galaxy surveys, the sirens, the lensing time delays — must read the same rate, so the one instrument that disagrees is predicted to be the one with rungs. The framework says the ladder reads about ten percent high through a shared calibration, and it declines to invent physics for any of that. The one environmental effect it could have credited — we live in a region emptier than average, and such a region expands slightly faster — is computed from linear growth theory and found to need a depth the surveys do not show, so it is priced and not carried. An earlier version of this Part credited two to three percent to it; the parent paper withdrew that, and this Part now agrees.

The ladder’s error budget is then walked channel by channel — the dust law, the calibration priors, the stability of the subsamples, the statistics of the ladder itself — and each channel is asked how much of ten percent it could hide. Several can hide a little. None can hide ten percent, and together, at their published sizes, they do not yet reach it: that is a finding against this paper, and it is printed as one. The arithmetic is then run in both directions, so that a reader who thinks a systematic of that size impossible can see exactly what would have to be true, and so can a reader who thinks it easy.

That is the state of the reconciliation and the paper states it the same way everywhere: nothing derived, everything named and open. The exposure here is total, because a systematic of that size in the distance ladder is precisely the sort of thing a better ladder analysis will either find or exclude within a few years, and this paper has already agreed to accept the verdict.

Part V

The 2026 Measurement Landscape and the Rival Explanations

Every reported H_0 of the audited reference base, drawn and tabulated; the three bands; the six rival classes read against them; and the one measurement that bisects the field

43 The measurement landscape in 2026

Before the rival explanations can be judged, the evidence they are rivals *about* must be laid on the table — all of it, not a curated pair. The clean “67 versus 73” picture in which the Hubble tension was born has not survived its own decade: the measurement landscape of 2026 is a *spectrum* of forty-five published determinations from nine method families, and its structure — three bands, not two camps — is itself a datum that any explanation must now explain. This section assembles that landscape from the primary literature, audited value by value; the sections that follow put each rival class in front of it.

43.1 The three bands, and how to read them

Figure 20 draws, and Table 13 tabulates, every reported H_0 in the audited reference base of this paper: forty-five determinations, each read from its primary arXiv or journal record on July 12, 2026, each carried with its published uncertainty, its identifier, and a one-line status remark. The audit discipline matters as much as the values: a claim enters the chart only at the number its source actually prints (two body-text values are flagged as such), contested entries carry their contest in the remark, and the entire set is reproduced by a released script in which every plotted value is asserted against the audit record. Nothing below is a poll of abstracts remembered; it is the literature re-read.

How to read the chart, in one paragraph. Rows are grouped into the nine method families and colored by them; the horizontal axis is the present expansion rate; three shaded bands mark where the determinations themselves cluster — band 1 at 67.2–68.5, band 2 at 69–71, band 3 at 72.5–77. Two green guides carry this paper’s numbers: the derived background 66.8 ± 0.2 (Route I, geometry alone) with the two-road reading 66.6 ± 0.2 beside it. The dotted line is the paper’s early reference anchor $H_0^{\text{ref}} = 67.24 \pm 0.35$, the three-instrument CMB combination. One bar (GWTC-4.0) runs off the right edge and is clipped with an arrow. The remarks column of the table is part of the data: it is where each value’s soft spot is named.

Band 1 (67.2–68.5): the absolute-density band. The five tightest determinations in the field live here — the three-instrument CMB combination (67.24 ± 0.35), P-ACT with lensing and DESI (68.22 ± 0.36), Planck alone (67.4 ± 0.5), DESI BAO with BBN (68.5 ± 0.6), and, newest and sharpest

as a single instrument, SPT-3G alone at 66.66 ± 0.60 . The physics of their agreement is the physics of Part III: the young plasma responds to *absolute* densities — grams per cubic meter, not fractions of a rate — so three telescopes with different sites, detectors, and calibration chains, and one route (BAO+BBN) that uses no CMB anisotropies at all, read the same universe because they are all reading the same thermodynamics. There is no shared hardware here to conspire. And the band’s floor is not empty: SPT-3G alone lands directly on the two-road 66.6 ± 0.2 and within its error on the geometry’s own 66.8, printed before that instrument’s value existed. Balance requires the whole band, however, and we print it against ourselves: band 1’s own members span 66.66–68.51 — an internal spread of $\approx 2.5\sigma$ the eye should not flatten — and its DESI-weighted combinations pull the *upper* edge (68.2–68.5), $3\text{--}4\sigma$ above the derived rate. The derived 66.8 lives at the band’s floor, not its centroid; whether the early anchor drifts toward the floor or the ceiling is itself a registered exposure (the commitments ledger, §58.2).

Band 2 (69–71): the corrected middle. The middle band is the newest structure in the landscape, and the most instructive, because almost nothing in it is anyone’s *home* value. It is where the distance ladder goes when any one of its published medium-size corrections is applied — the consistent-dust ladder (70.5 ± 1.0), the two-population remodeling (70.59 ± 1.15 , conservatively 71.45 ± 1.03), the dust-free blue-supernova ladder (70.0 ± 2.1), the physically-motivated-priors recalibration (70.6 ± 1.0 ; carried with its rebuttal) — and where the CCHP JWST program already stands (70.39 ± 1.94 , with its JWST-only TRGB and JAGB variants at 68.8 and 67.8). It is also where the probes that owe nothing to either camp natively sit: the localized fast radio bursts ($69.40^{+2.14}_{-1.97}$; Gao et al., 2024), the sharpest dark-siren value ($69.9^{+4.1}_{-4.0}$; Alfradique et al., 2026), the AGN-flare siren candidates (70.5, associations unconfirmed; Kumar et al., 2026), the ladder-free meta-average of fifty-eight measurements (69.37 ± 0.34 ; Pantos and Perivolaropoulos, 2026), the model-agnostic jackknife posterior (≈ 68 ; Hughes et al., 2025), and the upper half of the sound-horizon-free readings (70.0 ± 1.7 , $70.2^{+2.2}_{-2.4}$). The intuition is worth one sentence: band 2 is not a camp, it is a *destination* — the place measurements arrive when a systematic is removed or was never present, which is precisely what the residual anatomy of Part IV predicts the ladder would do under its named corrections ($73.5 \rightarrow \approx 70.6$ per reanalysis; 69.5–70.2 fully chained).

Band 3 (72.5–77): the ladder complex and the geometric high probes. The high band contains two different populations, and the distinction organizes the whole comparison problem. The first is the Cepheid-ladder complex — the H0DN consensus (73.50 ± 0.81), the full SH0ES set (73.2 ± 0.9), the covariance combination (73.30 ± 0.92), the background-free Perfect-Host field (73.49 ± 0.93), the standardized-TRGB grid (73.22 ± 2.06), the Mira rungs (72.4–73.1; Bhardwaj et al., 2025; Huang et al., 2023) — a family of formidable internal consistency whose members share calibration DNA, which is why its robustness tests certify photometry rather than immunity to common modes (Part IV, §34). The second population is the one-step geometric probes — megamasers (73.9 ± 3.0), Type-II supernovae (74.9 ± 1.9 stat), TRGB-SBF (73.8 ± 2.4), time-delay lensing ($71.6^{+3.9}_{-3.3}$), the Coma fundamental plane (76.5 ± 2.2), Tully–Fisher (73.3 ± 4.1) — which share no supernova standardization and mostly no Cepheids, and which are the standing counterweight to any pure-ladder-bias reading (§43.4). Across the whole band a single pattern holds and deserves its own line: *precision and directness are inversely ordered*. The sub-percent claims sit on the longest

calibration chains; the one-step probes that could arbitrate carry 3–5% errors and one named soft systematic each. The gamma-ray attenuation outlier ($61.9^{+2.9}_{-2.4}$, EBL-model dependent; Domínguez et al., 2023) is retained at the bottom of the chart as the honest reminder that not every anomaly points the same way.

Where this paper sits is then a matter of reading coordinates, not argument. The derived 66.8 ± 0.2 stands at the floor of band 1, with the newest single-instrument CMB value on top of it, and the framework predicts that every ladder-free probe drains to it; the corrected-ladder cluster at the floor of band 2 sits 2.7–3.8 above it, roughly three of its own error bars. The region into which the field’s center of gravity has been drifting from both directions — the early side down-anchored all along, the late side descending correction by correction — and both were printed by the geometry before most of the rows in the chart existed. That is the landscape form of the claim Part III staked: one number, exposed many ways.

Two things the landscape does *not* show, stated before any rival is judged. First, it does not show independent scatter: the 216-measurement survey of the field finds the literature’s spread exceeds its quoted errors (Wang et al., 2024), so the eye should weight bands, not count rows. Second, it does not show a culprit: the model-agnostic jackknife finds “no bias anywhere” strongly disfavored but declines to single out the ladder — many bias assignments fit comparably. The landscape constrains explanations; it does not convict for them. What it does do, decisively, is change the question every rival must answer: not “67 or 73?” but “*why three bands — and why is the middle one filling?*”

Figure 20: The 2026 H_0 landscape. Every reported H_0 in this paper’s audited reference base — forty-five determinations from nine method families, each read from its primary record on July 12, 2026 — with the three landscape bands shaded and this paper’s numbers shown twice, as the top category (*GG predictions*) and as vertical guides: the derived background 66.8 ± 0.2 (dashed green; the two-road 66.6 ± 0.2 solid beside it). The dotted line is the early reference anchor $H_0^{\text{ref}} = 67.24 \pm 0.35$ (this combination’s abstract variant reads 67.19 ± 0.38 ; the paper uses the H_0^{ref} convention throughout). The GWTC-4.0 and GW170817 upper error bars are clipped at the right edge (arrows). Row identifiers, posted dates, and per-value remarks are carried in Table 13; sources are cited in §43.1–§43.4 and Part B of the references. Generated by `c2_showcase_landscape.py`, released with the paper; every plotted value is asserted in-script against the audit record.

The 2026 H_0 Landscape: Every Reported Value in the Audited Reference Base (GG predictions on top; the CMB combination as guide)




Table 13: The 2026 H_0 landscape, tabulated. The forty-five determinations of Figure 20 in the figure’s order and categories, with this paper’s two reference rows on top. Posted = first arXiv posting (year-month from the identifier; day where verified). “In-paper” flags a value quoted from the source’s body rather than its abstract. H_0 in $\text{km s}^{-1} \text{Mpc}^{-1}$. All sources are carried in Part B of the references and cited at their first textual use in this Part.

Entry	arXiv	Posted	H_0	Brief remarks
GG Route I (carried)	—	—	66.8 ± 0.2	geometry alone, zero sky input; the rate this paper carries; Part III
GG two-road	—	—	66.6 ± 0.2	Route II (densities $\div$ derived fractions) read against Route I; the same background; Part III
Early / sound-horizon route				
SPT-3G D1 alone	2506.20707	2025-06	66.66 ± 0.60	tightest single-instrument CMB value; lands on the GG derived rate
CMB combined SPT+ACT+Planck	2506.20707	2025-06	67.24 ± 0.35	the paper’s H_0^{ref} (CMB-SPA convention; abstract variant 67.19 ± 0.38); 7.1σ below H0DN
Planck TT,TE,EE+lowE+lens.	2018 1807.06209	2018-07	67.4 ± 0.5	the legacy early anchor; A_L anomaly since resolved by ACT lensing
P-ACT + lensing + DESI DR1	2503.14452	2025-03	68.22 ± 0.36	68.43 ± 0.27 with DESI DR2; agrees with DESI at 1.6σ
ACT DR6 lensing + BAO	2304.05203	2023-04	68.3 ± 1.1	lensing route, no primary-anisotropy fit; Λ CDM-consistent at $\sim 1\%$
DESI DR2 BAO + BBN	2503.14738	2025-03	68.5 ± 0.6	in-paper; CMB-anisotropy-free ruler (deuterium + LUNA ω_b)
Sound-horizon-free				
DESI r_s -free, DESY5 SNe	2411.16677	2024-11	$66.7^{+1.7}_{-1.9}$	equality-scale route; closes “shorten the ruler” as a road to 73
DESI r_s -free, Pantheon+	2411.16677	2024-11	$67.9^{+1.9}_{-2.1}$	same analysis, Pantheon+ (Union3: $67.8^{+2.0}_{-2.2}$); $2.2\text{--}3\sigma$ below SH0ES
r_s -agnostic BAO+lens.+ $3\times 2\text{pt}$	2509.16202	2025-09	70.0 ± 1.7	higher r_d -free reading; prior-sensitive; widens the honest span
EFTofLSS $P+B$ r_s -free	2602.21733	2026-02	$70.2^{+2.2}_{-2.4}$	first $P+B$ sound-horizon-free H_0 ; 68.6 ± 1.8 with Pantheon+
Distance ladder (Cepheid / TRGB / JAGB / Mira)				
H0DN / LDN consensus	2510.23823	2025-10	73.50 ± 0.81	covariance-weighted community baseline; excludes the CCHP group
SH0ES full HST set	2408.11770	2024-08-21	73.2 ± 0.9	42 SNe, 4 anchors; Cepheid linearity 0.994 ± 0.010 ($\sim 7\sigma$ vs drift)
Ladder covariance combo	2606.26831	2026-06	73.30 ± 0.92	2026 review’s seven-route combination; 5.6σ vs CMB
Perfect-Host JWST Cepheids	2509.01667	2025-09	73.49 ± 0.93	background-free field test; 19/24 JWST hosts unbiased
SH0ES joint JWST 3-method	2408.11770	2024-08-21	72.6 ± 2.0	Cepheid+TRGB+JAGB on the joint JWST subsample; full-sample prediction 72.8
CATS TRGB+SNIa	standardized 2304.06693	2023-04	73.22 ± 2.06	tip-contrast-standardized; 108 variants, none below 71.6

continued on the next page

Table 13, continued.

Entry	arXiv	Posted	H_0	Brief remarks
SH0ES N4258-only anchor	1908.05625	2019-08	72.0 ± 1.9	single (maser-galaxy) anchor variant
Miras of M101	2312.08423	2023-12	72.37 ± 2.97	211 Miras; independent stellar-variable rung, sides high
Cluster-Mira anchor	2507.10658	2025-07	73.06 ± 2.67	41 globular-cluster Miras; no metallicity dependence over 1.6 dex
CCHP TRGB, published	2408.06153	2024-08	70.39 ± 1.94	published CCHP value (ApJ 985, 203 + erratum); low side of the split; the program's own Cepheid variant reads ≈ 72.0 — CCHP's variants span 67.8–72.1
CCHP JWST-only TRGB	2408.06153	2024-08	68.81 ± 2.23	JWST-only subsample (stat $\oplus$ sys); SH0ES: sample selection
CCHP JWST JAGB	2408.03474	2024-08	67.80 ± 2.72	blind carbon-star route, N4258 zero point; lowest CCHP value
Ladder re-analyses (dust / statistics)				
Consistent calibrator dust	2403.10388	2024-03	70.5 ± 1.0	MW-like $R_B \approx 4.3$ on calibrators fits better ($\Delta\text{BIC} = -11$); $5.2 \rightarrow 2.8\sigma$; same Pantheon+ calibrator SNe as the re-analysis rows below
Two-population SNe, baseline	2506.22150	2025-06-27	70.59 ± 1.15	hierarchical two-population Pantheon+ remodeling (stretch bimodality)
Two-population, conservative	2506.22150	2025-06-27	71.45 ± 1.03	same model, equal-prior variant; the channel's own internal spread
Blue (dust-free) SNe Ia	2411.05642	2024-11	70.0 ± 2.1	unextinguished blue SNe only, no dust law; ~ 6 calibrators
Physically-motivated priors	2601.22215	2026-01	70.6 ± 1.0	distance-prior recalibration, $5\sigma \rightarrow 2\sigma$; rebutted (2603.09880) — undecided; alternative to (not additive with) the dust rows
One-step geometric probes				
Megamasers, MCP	2001.09213	2020-01	73.9 ± 3.0	six maser galaxies, geometric; PV treatment spans 71.8–76.9
SNe II EPM, no rungs	2411.04968	2024-11	74.9 ± 1.9	tailored EPM, 10 SNe II (stat only); spectral-emulator systematics
TRGB-SBF via N4258	2502.15935	2025-02	73.8 ± 2.4	JWST TRGB $\rightarrow$ SBF zero point, Cepheid/SNIa-free; single anchor
TDCOSMO 2025 lensing	2506.03023	2025-06	$71.6^{+3.9}_{-3.3}$	8 time-delay lenses + JWST/Keck kinematics, blinded; MSD-conservative
DESI FP + Coma	2409.14546	2024-09	76.5 ± 2.2	fundamental plane on Coma (D from 12 SNe Ia — not SNIa-independent)
Tully–Fisher CF4	2408.03660	2024-08	73.3 ± 4.1	CF4 i /W1 TF (stat $\oplus$ sys); discordant zero-point calibrators
Gamma-ray attenuation	2306.09878	2023-06	$61.9^{+2.9}_{-2.4}$	EBL route at $\Omega_m = 0.32$ (free fit $65.6^{+5.6}_{-5.0}$); anomalously low
Gravitational-wave sirens				
GW170817 bright siren	1710.05835	2017-10	70^{+12}_{-8}	the one bright siren in nine years; ladder- and ruler-free

continued on the next page

Table 13, continued.

Entry	arXiv	Posted	H_0	Brief remarks
GWTC-4.0 + GW170817	2509.04348	2025-09	$76.6^{+13.0}_{-9.5}$	142 dark sirens + GLADE+; mass-spectrum-model dependent
GWTC-5.0 + GW170817	2605.27227	2026-05	$71.0^{+9.0}_{-7.1}$	236 sources + DES-Y6; 25.7% tighter than GWTC-4.0
O4a dark sirens + GW170817	2603.20195	2026-03	$69.9^{+4.1}_{-4.0}$	17 dark sirens, DESI Legacy photo- z ; sharpest siren value
AGN-flare BBH sirens	2607.03672	2026-07	$70.5^{+3.4}_{-2.9}$	13 claimed AGN counterparts (± 1.56 cal.); associations unconfirmed
Fast radio bursts				
108 localized FRBs	2410.03994	2024-10	$69.40^{+2.14}_{-1.97}$	IllustrisTNG DM priors; squarely between the camps
527 CHIME FRBs, nonlocalized	2410.03994	2024-10	68.81 ± 0.68	statistical sample; formally tight, DM-model dependent
Meta-analyses / statistics				
30 ladder measurements, mean	2601.00650	2026-01	72.73 ± 0.39	mean of the distance-ladder family
58 ladder-free, mean	2601.00650	2026-01	69.37 ± 0.34	ladder-independent mean; internal split up to 6.5σ
Bayesian jackknife, 16 probes	2511.19341	2025-11	$68.0^{+4.7}_{-1.3}$	model-agnostic posterior (95%); no-bias strongly disfavored
Model-dependent early fixes				
EDE profile likelihood, no SH0ES	2505.08051	2025-05	71.0 ± 1.1	$f_{\text{EDE}} = 0.09 \pm 0.03$ (ACT+DESI); vs the ACT team’s own null

43.2 The consensus that is not a consensus

The H_0 DN consensus 73.50 ± 0.81 is the subject measurement of this paper, and Part II took its internal architecture seriously — thirty structured variants, orthogonal paths, leave-one-out stability. One sociological fact must now be set beside that architecture, because the landscape makes it quantitative: the consensus’s ~ 37 authors do not include the CCHP group — the team whose JWST measurements (70.39 ± 1.94 published; 68.8 and 67.8 in its JWST-only variants) disagree with the baseline most (Freedman et al., 2025; Lee et al., 2024). The Freedman–Riess split has not converged; it has *localized*: the 2026 ladder review (Chen et al., 2026) and the SH0ES same-data reanalysis (Riess et al., 2024b) place the residual disagreement in supernova-calibrator subsampling and tip/JAGB standardization conventions — not in anchors, not in crowding, not in photometry. The practical consequence for every comparison in this Part: the ladder’s headline number and its principal internal dissent must both be carried, and this paper carries both — 73.50 ± 0.81 as the subject measurement, 70.39 ± 1.94 as the standing demonstration that calibrator choices worth $\sim 3 \text{ km s}^{-1} \text{ Mpc}^{-1}$ are live inside the ladder community itself. An internal-consistency architecture, however excellent, certifies what a collaboration *shares*; it is silent on what its members do not share — which is precisely where the subsample channel of Part IV lives (F-Hubble-10).

43.3 The early side is closed

The landscape’s band 1 is not merely populated; it is *overdetermined*, and the 2024–26 data have closed, one by one, every route by which the early side could have been blamed for the tension. This subsection collects the closures, because Part V’s class verdicts lean on them repeatedly.

Three instruments, one thermodynamics. Planck, ACT DR6, and SPT-3G — different sites, detectors, scan strategies, and calibration chains — agree at a fraction of the offset in play (Camphuis et al., 2025; Louis et al., 2025); their combination, $H_0^{\text{ref}} = 67.24 \pm 0.35$, carries an empirical inter-experiment spread that bounds any CMB-side systematic at a quarter of what the tension needs. The old A_L lensing anomaly, once the early side’s one internal blemish, is resolved: ACT DR6 lensing agrees with Λ CDM at the percent level (Madhavacheril et al., 2023), and the lensing+BAO route — no primary anisotropies at all — reads 68.3 ± 1.1 .

The ruler is audited twice. The sound horizon’s raw material is the baryon density, and deuterium fixes it without any microwave sky: the LUNA-pinned BBN value $\omega_b = 0.02218 \pm 0.00055$ (Schöneberg, 2024) agrees with the acoustic reading, and DESI’s BAO+BBN route lands at 68.5 ± 0.6 (Abdul Karim et al., 2025). The intuition: recombination is the ruler’s factory, but the factory’s books are audited by an accountant (BBN) who has never seen the factory — and the books balance.

Even the ruler can be discarded. The sound-horizon-free determinations — DESI’s equality-scale route at $66.7\text{--}67.9$ (Zaborowski et al., 2024), the r_d -agnostic large-scale-structure reading at 70.0 ± 1.7 (García Escudero et al., 2025), and the EFTofLSS power-spectrum-plus- bispectrum route at $70.2^{+2.2}_{-2.4}$ (Lu et al., 2026) — span $66.7\text{--}70.2$: every one of them well below 73. The honest breadth of that span is printed deliberately: the r_s -free band overlaps bands 1 and 2, so it does not adjudicate between the derived background and the corrected middle — but it closes the “shrink the ruler” road to the ladder from every side, because if sound-horizon inference were the culprit, the ruler-free readings would sit high, and none does.

What survives on the early side, capped. The one genuinely live early-side lever is non-standard recombination, and it is now squeezed from four directions; its own cap — $H_0 < 70.7$ at 95% confidence with CMB+BAO, from the non-parametric ionization reconstruction (Lynch et al., 2024a) — reaches band 2, not the ladder. The full accounting is Class 3 (§47). For the derived reading the consequence of this subsection is structural: the early anchor now stands on three instruments, a CMB-free ruler audit, and ruler-free cross-checks, with the newest single-instrument value sitting on the derived rate itself — so every rival that begins “the early inference is wrong” begins three closed doors behind the evidence.

43.4 The high geometric counterweight

Honesty requires the strongest opposing evidence to be stated at full strength, in its own subsection, before any class is judged. The counterweight to every measurement-side reading — this paper’s included — is band 3’s second population: probes that reach the Hubble flow in one step, without Type-Ia standardization and mostly without Cepheids, and read high anyway. The megamaser galaxies are pure geometry — water-maser orbits give angular-diameter distances directly — and read 73.9 ± 3.0 , though the value slides across $71.8\text{--}76.9$ with the peculiar-velocity treatment of just six objects (Pesce et al., 2020). The Type-II expanding-photosphere method needs no ladder

at all — the photosphere’s physical size comes from spectra — and reads 74.9 ± 1.9 (statistical), with its systematics budget resting on a spectral-emulator calibration still in its proof-of-principle phase (Vogl et al., 2025). The TRGB-SBF route ties surface-brightness fluctuations to a JWST tip calibration — no Cepheids, no supernovae — and reads 73.8 ± 2.4 , on a single anchor galaxy (Jensen et al., 2025). Time-delay lensing, blinded and mass-sheet-conservative, reads $71.6^{+3.9}_{-3.3}$, straddling bands 2 and 3 (Birrer et al., 2025). The Coma fundamental plane reads highest of all (76.5 ± 2.2) but calibrates Coma’s distance with twelve Type-Ia supernovae, so it is not independent of the standardization question (Scolnic et al., 2024); Tully–Fisher reads 73.3 with a $3.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ systematic term driven by discordant zero-point calibrators (Boubel et al., 2024).

Read as a population, these probes say something no advocate of a pure-bias reading may wave away: *whatever is high in the local universe is not confined to the Cepheid–supernova chain*. Read individually, none is decisive — each carries 3–5% errors and one named soft systematic (six-object peculiar velocities; an emulator; a single anchor; the mass-sheet degeneracy; a supernova-calibrated cluster distance; discordant zero points) — and their weighted center of gravity, together with everything else that is not the Cepheid ladder, remains ~ 69 –71. The derived reading’s obligations to this subsection are recorded in the registry: if the maser sample triples and holds 74, if the Type-II emulator survives its systematics audit at 75, the measurement-side story is in real trouble well before any siren flies — and the converse collapses band 3 to one calibration family. The counterweight is not an inconvenience to be footnoted; it is the second-sharpest open question the landscape poses, after the tension itself.

44 Statement of the comparison problem

The landscape sharpens the problem every proposed resolution must solve. It is no longer enough to move one number from 67 to 73: a viable explanation must now account for a *three-band structure* — an overdetermined absolute-density band at 67.2–68.5; a filling middle band at 69–71 into which every corrected ladder and every camp-free probe drains; and a high band whose two populations (one calibration family, one set of soft-systematic geometric probes) it must explain *differently* or explain why not. It must survive the meta-split — thirty ladder measurements averaging 72.73 ± 0.39 against fifty-eight ladder-free averaging 69.37 ± 0.34 , internally discrepant at up to 6.5σ (Pantos and Perivolaropoulos, 2026) — and it must survive the sound-horizon-free readings, which cap what any ruler-side fix can deliver. The literature has produced many candidates (Di Valentino et al., 2021, 2025; Poulin et al., 2023); we examine six representative classes, each against two instruments: the *cross-ledger coherence* of Part IV, and the *siren discriminator* of §51, which now carries three outcomes rather than two.

The six classes, in the taxonomy the 2026 landscape itself suggests:

1. Λ CDM plus measurement systematics — now a *named* hypothesis, not a shrug;
2. phenomenological early dark energy;
3. modified recombination (split out from EDE: same lever, different physics, different caps);
4. interacting dark sectors and extra radiation (merged: both re-time or re-load the early plasma

- through added dark content);
5. modified gravity and evolving dark energy;
 6. local voids and inhomogeneity.

45 Class 1 — Λ CDM plus measurement systematics: a named hypothesis

The hypothesis: the concordance model is correct, the early inference is unbiased, and the ladder’s excess is measurement systematics. For most of the tension’s decade this class was a shrug — “some unknown error, somewhere” — and it was treated accordingly.

What the network’s architecture retires — and what it cannot. The H0DN architecture genuinely retires the class’s naive form. The orthogonal-path test shows two statistically independent local routes agreeing at $< 0.2\sigma$ (Fig. 12 of Casertano et al., 2026); leave-one-out shows no single method drives the result; HST and JWST agree (Riess et al., 2022, 2024). Any *private* systematic — one instrument’s bug, one method’s bias — is dead. But the same architecture, by construction, cannot price a *common-mode* term: variants that share calibration DNA move together, so their mutual stability certifies internal consistency, not external truth (§34). The intuition is an auditor’s: thirty cross-checked copies of one ledger prove the copying is faithful, not that the ledger is right. What the class needed, to be more than a shrug, was *named* common-mode mechanisms with published sizes — and that is exactly what 2024–26 supplied.

The named anatomy. The class now runs on the three channels priced in Part IV: the calibrator-versus-flow dust and two-population standardization mismatch (allowance 2–4%; published central values 70.5 ± 1.0 , 70.59 ± 1.15 , with the dust-free cross-check at 70.0 ± 2.1); the Bayesian distance-prior statistics of the ladder itself (allowance 0–3%; proposed, executed at 70.6 ± 1.0 , and rebutted — undecided); and the calibrator-subsample selection localized by the CCHP split (allowance 1–2%). Under the proposal-central reading the channels quadrature to 4.3% and stack to 9.3% against the +10.0% requirement; the published central values alone carry $73.5 \rightarrow \approx 70.6$ –70.2-if-chained — into band 2, still 2.7–3.8 above the derived 66.8.

What the class explains, and where it strains. Explains: why the middle band is filling (corrected ladders drain into it); why the CCHP split exists (subsampling is one of its channels); why the $H_0(z)$ trend descends (§33.4); why the meta-split separates ladder from ladder-free. Strains: at the geometric counterweight (§43.4) — masers, Type-II supernovae, and SBF are touched by *neither* the dust nor the priors channel, and a complete Class-1 story must either find their own soft systematics doing coherent work or concede band 3’s second population to coincidence. The class is honest only when it carries that strain visibly.

What GG-Theory offers that this class does not. A *positive* decomposition rather than a corrected suspicion: a background *derived* — twice, once with zero sky input (Route I, its coefficient

pre-registered as F-Hubble-11) — where Class 1 assumes one; a void channel computed and found short where Class 1 would fit one; a systematic priced in both directions against the ladder’s own budget; and a stake. And one thing should be said plainly, because a careful reader will notice it anyway: for the ten percent, the derived reading is itself a disciplined member of this class. The differences are the discipline — and one number: Class 1 has no opinion about *where* the corrected ladder should land, while the geometry printed 66.8 ± 0.2 before the corrections existed. The field’s corrections have so far carried the ladder about a third of the way there; whether they carry it the rest of the way is either the gap closing as predicted, or the point at which this reading fails — and Class 1, which predicts no endpoint, cannot fail that way.

Siren verdict required. ≈ 67 –68: if the excess is measurement, the true rate is low, and the ladder-free channel must say so. The class lives or dies with the derived reading on the same catalog — the kinship above, made operational.

46 Class 2 — phenomenological early dark energy

The EDE program (Kamionkowski and Riess, 2023; Karwal and Kamionkowski, 2016; Poulin et al., 2019, 2023) introduces a scalar field contributing a narrow lobe of energy density around $z \sim 3500$ that dilutes away. It works because shortening r_s by $\sim 5\%$ near that epoch is exactly what the lever of Eq. (35) converts into a higher inferred rate: hold the measured angle θ_* fixed, shave the stick, and the inference must stand closer.

Status after the 2026 landscape: two-sided, and capped. The class’s empirical position is now genuinely two-sided, and the paper prints both sides. On one side, the ACT collaboration’s own extended-model search finds *no evidence* for EDE (Calabrese et al., 2025), and the joint CMB+BAO pressure confines the lobe to the low end (Camphuis et al., 2025; Kreisch et al., 2024). On the other, a profile-likelihood analysis of the same ACT DR6 + DESI DR2 data — deliberately prior-free where the Bayesian search is prior-limited — reports $f_{\text{EDE}} = 0.09 \pm 0.03$ and $H_0 = 71.0 \pm 1.1$ *without* any SH0ES input (Poulin et al., 2025). An honest scorer calls that a live dispute about statistical method, not a settled null. What is settled is the ceiling: the sound-horizon-free determinations (§43.3) span 66.7–70.2, and a mechanism whose entire output is a shorter ruler cannot move determinations that never use the ruler. EDE has become a *band-2 machine*: it can reach 70–71; it cannot reach the ladder, and the DESI-era Ω_m coherence squeezes it further (Poulin et al., 2023). The deuterium audit is the same wall from a second side: early fixes that raise H_0 generically raise the inferred ω_b against the BBN value (Giovanetti et al., 2026).

What GG-Theory offers that EDE phenomenology does not. The decisive difference: the derived reading needs no early modification at all. EDE buys a higher inferred H_0 by shortening the sound horizon; the derived background already sits on the CMB value, so there is no early-side gap to close. Four consequences:

1. **No new field, no new potential.** EDE adds a scalar with a $V(\phi) \propto (1 - \cos(\phi/f))^n$ potential

- and free n , f , and mass scale. The derived reading adds nothing: the acoustic spectrum is fit by the π -locked fractions with the standard recombination history untouched (§26).
2. **No S_8 cost.** The early lobe that shortens r_s also suppresses growth, and in full likelihood refits the net S_8 shift typically comes out worse, not better (Hill et al., 2020; Poulin et al., 2023). A reading with no early lobe pays no such cost (§39).
 3. **Opposite exposure.** EDE is confirmed by detecting a lobe; the derived reading is *refuted* by one (F-Hubble-3, §56). The profile-likelihood hint above is therefore not an embarrassment to be argued away but this paper’s own kill test in motion: if the $f_{\text{EDE}} = 0.09 \pm 0.03$ reading survives the likelihood dispute and sharpens past 3σ , F-Hubble-3 fires. The two frameworks make opposite pre-registered bets on the same forthcoming analyses.
 4. **Cross-domain reach.** The same integer ledger that fixes the inventory also fixes the Standard Model’s 28 parameters and the Kaxion spectrum. EDE is a cosmology-only fix.

Siren verdict required. ≈ 73 to solve the tension it was built for — but the landscape has narrowed its own demand: since its caps now reach only band 2, a siren at ≈ 70 is the best confirmation EDE can still claim, and a siren at 67 retires it outright. One catalog, opposite bets: early dark energy and the derived reading cannot both survive it.

47 Class 3 — modified recombination

Split out from the EDE class because it pulls the same lever with different physics — and meets different walls. Instead of adding energy density, this class changes *when and how the fog lifted*: shift the ionization history $x_e(z)$, move recombination earlier, and the sound horizon frozen at the switch is shorter, so the fixed acoustic angle again converts into a higher inferred rate. Concrete mechanisms include spline-level freedom in $x_e(z)$ (Lynch et al., 2024a,b), and primordial magnetic fields that clump baryons and hasten recombination (Jedamzik et al., 2025).

Status: the factory audit. The intuition of §43.3 does the work: recombination is the ruler’s factory, and the factory’s books are audited by accountants who never enter it. Four audits now cap the class. *First*, the non-parametric reconstruction itself: given full spline freedom in the ionization history, CMB+BAO data allow at most $H_0 < 70.7$ at 95% (Lynch et al., 2024a) — the class’s own ceiling, printed by its own machinery, reaching band 2 and no further. *Second*, deuterium: raising H_0 this way raises the inferred ω_b , and the LUNA-pinned BBN value stands in the way generically (Giovannetti et al., 2026; Schöneberg, 2024). *Third*, DESI: the spline histories that help H_0 drag Ω_m against the DR2 measurement (Lynch et al., 2024b); the electron-mass variant of the same idea, which resolves the CMB-only fit, *fails* once DESI DR2 is included (Lee et al., 2026). *Fourth*, the magnetic-field mechanism’s evidence stands at 1.8σ — a hint, dataset dependent, honestly below claim level (Jedamzik et al., 2025). Within its caps the class remains the most respectable early-side survivor: spline reconstructions genuinely prefer nonstandard histories at the 2–3 σ level in some data combinations (Mirpoorian et al., 2024), and a band-2 landing is within its reach.

What GG-Theory offers. No need for the lever, and no freedom to pull it. The derived chart lands on the *measured* sound horizon ($r_s = 144.7$ against 144.43 ± 0.26) with the standard ionization history untouched, and the corridor switch is not adjustable: $z_* = e^7$ is the exponential of a ledger integer, a topological quantity with no continuous dial (§31). A framework whose recombination redshift is an integer’s exponential cannot join this class even if it wanted to — there is nothing to tune — and its agreement with the measured decoupling at 0.6% is therefore a checkable coincidence-or-theorem, not a fit.

Siren verdict required. ≈ 73 to be the tension’s resolution — which its own 70.7 ceiling already forbids; ≈ 70 to survive as a band-2 contributor. A siren at 67 retires the class; at 70 it lives, but as physics of the middle band, not as the explanation of the ladder.

48 Class 4 — interacting dark sectors and extra radiation

Merged from two classes of the earlier taxonomy, because both work by re-loading the early plasma with dark content: couple the dark fluids so energy migrates between them, or add relativistic species (ΔN_{eff}) so the pre-recombination expansion runs faster. Both re-time equality or shorten the horizon; both are early-side fixes wearing different hardware.

Status. The interacting-sector limb writes coupled continuity equations — $\dot{\rho}_{\text{dm}} + 3H\rho_{\text{dm}} = +Q$, $\dot{\rho}_{\text{de}} + 3(1+w)H\rho_{\text{de}} = -Q$ — whose exchange rate Q is a new function with new couplings; case by case such models can fit CMB+ H_0 + S_8 combinations, at the cost of parameters with no independent motivation. The radiation limb is cleaner and more cornered: closing the gap to the ladder needs $\Delta N_{\text{eff}} \gtrsim 0.4$, which BBN’s deuterium and helium refuse (Schöneberg, 2024), and which the current CMB bound $\Delta N_{\text{eff}} = 0.04 \pm 0.18$ leaves no room to hide. Both limbs also answer to the ruler-free wall: whatever they do to r_s , the sound-horizon-free band at 66.7–70.2 does not move with them.

What GG-Theory offers. No interaction, and no room for one — and that is the point. The dark sectors the framework contains are locked: the $U(1)_{Q_i} \leftrightarrow U(1)_B$ channel is fixed by the BI-6 coefficients $(1, -1, -2/5, -2/5)$ with no free strength anywhere, and the only radiation term the cohort carries is a derived null, $\Delta N_{\text{eff}} \leq 0.03$ (GG-A10-Cosmos) — consistent with the bound, unrelated to the Hubble split, and two-sided: a bound tightening toward zero supports the ledger; a detection at $\Delta N_{\text{eff}} \gtrsim 0.2$ is new physics the framework neither needs nor survives unchanged. A framework whose dark-sector couplings are pinned by the same integers that fix $\sin^2 \theta_W = 25/108$ is not in the same complexity class as a model with three couplings tuned to a fit.

Siren verdict required. ≈ 73 for the class’s viable early-side members; late-coupling variants that bend $H(z)$ at low redshift face the $w(z)$ and BAO constraints directly and inherit whichever verdict their bending implies. A siren at 67 retires the class alongside every other early-side fix.

49 Class 5 — modified gravity and evolving dark energy

This class alters the gravitational sector itself: a running $G_{\text{eff}}(z)$, a scalar–tensor extension, or — its newly prominent limb — an evolving dark-energy equation of state $w(z)$ of the kind the DESI $w_0 w_a$ analyses have made headline news.

Status, with the wrong-sign surprise. The modified-gravity limb can shift H_0 , but typically at the cost of scale-dependent growth modifications that redshift-space distortions and lensing rule out, or equivalence-principle violations that solar-system tests forbid (Di Valentino et al., 2025). The evolving-dark-energy limb is where the data actually moved: DESI DR2 with supernovae prefers $w_0 > -1, w_a < 0$ at $2.8\text{--}4.2\sigma$ depending on the sample (Abdul Karim et al., 2025) — a genuine, contested, parametrization- sensitive hint. But its sign is the striking fact: evolving dark energy of the preferred kind pushes the inferred H_0 *down*, not up. The one live new-physics hint of the DESI era is a rope thrown in the wrong direction — it entangles the early anchor’s error budget, it cannot rescue 73, and if it solidifies it wounds the ladder’s side of the tension more than the CMB’s. (For this paper it is a two-sided exposure of its own: $w = -1$ exactly is kill switch F-Hubble-5, and the DESI hint is that switch under live test.)

What GG-Theory offers. Gravity is not modified — AX-2 locks the geometry–gauge structure, solar-system tests are preserved structurally rather than by tuning — and the vacuum does not run: $w = -1$ is the ground-attractor theorem, staked two-sidedly. The derived reading needs neither limb: the rate is what it is because the inventory is what it is.

Siren verdict required. Either — and that ambivalence is the class’s own signature. A running effective Planck mass generically makes the gravitational-wave luminosity distance differ from the electromagnetic one, so the siren tests this class twice: once through the rate, once through the GW–EM distance duality checked event by event on the same detections. A duality violation would be its smoking gun; a null, one more constraint tightening around it.

50 Class 6 — local voids and inhomogeneity

The hypothesis: the local universe sits in an underdensity that lifts every locally inferred rate above the global mean. After Part IV this class is largely *native* machinery to this paper, so the update is brief and quantitative.

Status after the 2025–26 data. Three depths must be distinguished, because the data now treat them differently. The *modest* void — $\bar{\delta} \sim -0.1$ to -0.2 over the flow volume, worth $+2\text{--}3\%$ — is observed in sign, contested in depth, and not carried: the framework computes the class’s response and credits none of it, because the depth that would make even the modest member viable is an $\approx 11\sigma$ fluctuation on its own power spectrum that the velocity data disfavor (§33). The *deep* void — the KBC reading at -0.3 over 300 Mpc, worth $+5\%$ — is disfavored from within the void camp itself: field-level forward modeling of the CosmicFlows-4 velocities prefers voids smaller

than ~ 70 Mpc, and the profiles that help H_0 most contradict the measured bulk flow (Stiskalek et al., 2025); supernova smoothness independently excludes $|\delta| > 0.2$ voids at $4\text{--}5\sigma$ (Kenworthy et al., 2019). The *full-tension* void — $+10\%$ from underdensity alone — requires $\langle\delta\rangle \approx -0.56$ on flow scales, a $> 12\sigma$ fluctuation (§35): excluded in standard structure formation, and invocable only through the fully non-linear timescape loophole, which remains this class’s honest extreme (Wiltshire, 2007). The class’s supporting evidence has meanwhile sharpened in exactly the modest direction: the $H_0(z)$ trend now stands at 6.4σ in DESI’s uncorrelated bins (Jia et al., 2025) — at low redshift the generic signature of *some* local outflow with depth-tracking amplitude, though a tail persisting at $z \gtrsim 0.3$ would instead be background evolution, adverse at F-Hubble-5; and the CF4 bulk-flow excess is carried as contested companion evidence, its significance estimator-dependent (428 ± 108 under the audit of Whitford et al., 2023).

What GG-Theory offers. The weakest viable member of the class, computed rather than invoked: sign unconditional, size bounded by the same linear theory that excludes its deep cousins, profile-diluted, flow-correction-guarded — and not carried, because the depth that would make even the weakest member viable is an $\approx 11\sigma$ fluctuation the velocity data disfavor (§33). The framework’s one exotic local channel, the $U(1)_B$ vector, is closed by the boundary-condition termination of GG-A3-B16 and, independently, excluded by eight orders of magnitude.

Siren verdict required. ≈ 67 : a void biases only the observers inside it, and a catalog reaching well beyond it reads the true background. The class’s quantitative apparatus — profiles, dilutions, rarity bounds — is Part IV’s own, and it caps the pure-void explanation far below the requirement.

51 The siren discriminator: one measurement, three worlds

The class-by-class verdicts assemble into the sharpest table in this Part — now with three columns of fate, because the landscape has three bands. The classes disagree, at bottom, about a single question: *where does the true background rate sit?* Every early-side fix needs it high — the tension is real physics and the ladder was right. Every measurement-side reading needs it low — the ladder is the displaced instrument. And the landscape’s new middle band forces a third possibility onto the table: a true rate near 70, which no class predicted outright and several could partially absorb. The gravitational-wave standard siren reads the rate with neither a sound horizon nor a ladder (§64), and its catalogs are already walking: $76.6^{+13.0}_{-9.5}$ at GWTC-4.0, $71.0^{+9.0}_{-7.1}$ at GWTC-5.0, $69.9^{+4.1}_{-4.0}$ with the O4a dark sirens (Alfradique et al., 2026; LVK Collaboration, 2025, 2026) — so one converged catalog still bisects — now trisects — the entire solution space (Table 14).

The three worlds, spelled out. *At ≈ 67 :* the derived reading and its Class-1 kin are confirmed; every early-side fix retires at once; the residual anatomy of Part IV graduates from hypothesis to history. *At ≈ 73 :* the derivation dies in public (F-Hubble-2), the anatomy evaporates, and Classes 2–4 revive with the ladder vindicated. *At ≈ 70 :* the subtlest world, and the one the middle band now makes live — the H0DN centroid falls, the corrected-ladder cluster is vindicated, EDE and modified recombination survive as band-2 physics but lose the tension they were built to solve; for this paper the derived 66.8 ± 0.2 is strained at $1.4\text{--}2.4\sigma$ (the exact arithmetic, at the two verdict

thresholds, is §70) — pressure this paper would report as exactly that, not as confirmation. The stake does not soften: the pre-registered rule of Part VII scores the catalog against 66.8, and the honesty of the three-world table is that the middle world is uncomfortable for everyone, this reading included.

Table 14: The siren discriminator, three ways. Each class’s required location of the true background rate and its fate under a converged siren catalog at ≈ 67 , ≈ 70 , or ≈ 73 . A 67 retires every early-side fix at once; a 73 retires every measurement-side reading, this paper included; a 70 vindicates the corrected middle, strains the pure derived rate at $1.5\text{--}2.4\sigma$, and demotes the early-side fixes to band-2 physics. The derived reading is the only entry whose low bet is a derivation rather than a preference.

Class	True ground	back-	Needs	If 67	If 70	If 73
Λ CDM + systemat- ics (named)	low		$\approx 67\text{--}68$	confirmed	partly confirmed	dies
Pheno. EDE	high		≈ 73 (caps: ~ 71)	dies	survives as band-2	vindicated
Modified recombini- nation	high		≈ 73 (cap: 70.7)	dies	survives as band-2	strained by cap
Interacting DS / ΔN_{eff}	high		≈ 73	dies	squeezed	vindicated
Mod. gravity / evolving DE	either		either + duality	constrained	constrained	constrained
Local void (mod- est)	low		$\approx 67\text{--}68$	confirmed (capped)	partly confirmed	dies
GG-Theory (derived)	(de- low — 66.8	<i>derived</i>	≈ 67 Hubble-2)	(F- confirmed	strained: $1.4\text{--}2.4\sigma$, ad- judication to the inven- tory and the shell test	dies

Three remarks complete the point. First, the discriminator remains total: whichever of the three worlds the catalog lands in, most of the standing literature retires or is demoted in one measurement — a rarity in a field where proposals usually fail slowly. Second, modified gravity is still the one class the rate alone cannot bisect, and it compensates by exposing itself twice through the GW–EM duality. Third, the derived reading’s row differs from its allies in kind, not in verdict: Classes 1 and 6 *prefer* a low rate; the geometry *computed* one, in writing, before the catalog that will check it exists. Sharing a bet is not sharing a claim.

52 Comparison summary

Table 15 scores the six classes and the derived reading against the standing constraints; Table 16 adds the column the 2026 landscape itself supplies — each class’s status after the audits of §43.3, in one line each.

The key distinction in the last three rows of Table 15 is structural. The six classes either carry parameters fitted after the fact, or are single-purpose cosmology extensions with no cross-domain reach, or have no unique falsifier that retires the entire framework. GG-Theory is the only candidate

Table 15: Six classes of Hubble-tension fix versus GG-Theory. $\checkmark$ = achieved; $\times$ = not achieved; $\sim$ = partial; $*$ = with a GW–EM distance-duality signature (§51). No column here is clean, including the last one. The rows are ordered so that the ones a reader is most likely to care about — the acoustic fit, the deuterium abundance, the r_s -free readings — come before the rows that favor this framework, and the final row is the one that will settle the question without any of the arguments above it.

Feature	Λ CDM +sys- tem.	Pheno. EDE	Mod. re- comb.	Int. DS / ΔN_{eff}	Mod. grav. /DDE	Loc. void	GG- Theory
Accounts for the local $\sim 73.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$	$\sim$	$\sim$	$\times$	$\checkmark$	$\sim$	$\sim$	$\sim$
Preserves CMB acous- tic fit	$\checkmark$	$\sim$	$\sim$	$\sim$	$\sim$	$\checkmark$	$\checkmark$
Does not worsen S_8	$\checkmark$	$\times$	$\sim$	$\sim$	$\times$	$\checkmark$	$\sim$
Consistent with BBN, D/H	$\checkmark$	$\sim$	$\times$	$\times$	$\checkmark$	$\checkmark$	$\checkmark$
Consistent with BAO (incl. DESI DR2)	$\checkmark$	$\sim$	$\times$	$\sim$	$\sim$	$\checkmark$	$\checkmark$
Survives the r_s -free readings	$\checkmark$	$\times$	$\times$	$\times$	$\sim$	$\checkmark$	$\checkmark$
Consistent with solar- system tests	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\sim$	$\checkmark$	$\checkmark$
Integer-level inheri- tance / no free fit	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\sim$
Formulated indepen- dently of the H0DN analysis	$\sim$	$\sim$	$\sim$	$\sim$	$\sim$	$\sim$	$\checkmark$
Cross-ledger reach (particles, biology, mind)	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\checkmark$
Single falsifier retires the whole theory	—	$\times$	$\times$	$\times$	$\times$	$\times$	$\checkmark$ (pro- ton decay)
Standard-siren verdict required (§51)	≈ 67	≈ 73	≈ 73	≈ 73	either*	≈ 67	≈ 67

Table 16: Status of each class after the 2026 landscape: the one-line verdict the audits of §43 impose, independent of any preference. The verdicts are the 2026 record’s and not this paper’s. Each is a one-line compression of an audit given in full in the section named above, and each was written to be readable by someone who prefers the class it describes. The framework’s own row is included on the same terms and is not the strongest entry in the table.

Class	Status after the 2026 landscape
Λ CDM + systematics	Promoted from shrug to named hypothesis: three published channels, quadrature 4.3%, central values reach band 2; strained by the geometric counterweight
Phenomenological EDE	Two-sided (ACT null vs. profile-likelihood $f_{\text{EDE}} = 0.09 \pm 0.03$); capped below the ladder by the r_s -free span; a band-2 machine
Modified recombination	Own ceiling $H_0 < 70.7$ (95%); deuterium and DESI Ω_m audits; m_e variant fails DR2; PMF hint at 1.8σ
Interacting DS / ΔN_{eff}	Radiation limb refused by BBN ($\Delta N_{\text{eff}} \gtrsim 0.4$ needed, 0.04 ± 0.18 measured); coupled limb unmotivated parameters; r_s -free wall
Mod. gravity / evolving DE	DESI $w_0 w_a$ hint ($2.8\text{--}4.2\sigma$, contested) points <i>down</i> — wrong sign for the tension; duality test pending
Local void	Modest (+2–3%) observed in sign, contested in depth, not carried; deep disfavored by the void camp’s own velocities; full-tension void at $> 12\sigma$ excluded
GG-Theory (derived)	Derived 66.8 sits at band 1’s floor under the newest single-instrument value; the corrected-ladder cluster at band 2’s floor sits 2.7–3.8 above it, and the framework predicts it drains further

listed in which the numbers driving the Hubble sector are the same integers that drive the SU(5) charge algebra and the 4-letter DNA alphabet — with the honest qualification, encoded by the $\sim$ entries in its column, that the derived reading does not account for the local 73.5 at all: it names the whole $\approx 10\%$ as a calibration systematic it predicts, with published, kill-testable candidate anatomy that does not yet reach it, rather than producing it (§34).

53 Why “uniqueness” means more than “agreement”

The H0DN paper itself notes (Section 6 of Casertano et al. 2026) that the local direct measurement of H_0 is “the most accurate direct H_0 measurement to date, with a relative uncertainty of 1.09%, including systematic uncertainties.” Any theory that merely *agrees* with some row of the landscape is already on a crowded shelf — the chart offers forty-five numbers to agree with. The uniqueness of the GG-Theory reading lies elsewhere:

1. **Derived anchor.** The background rate is not fitted to any row of the landscape: the π -locked fractions, fixed by the integer ledger (1, 3, 5; 7, 35) independently of any cosmological measurement, deliver the rate twice over — 66.8 from geometry alone, 66.6 from the measured densities (Section 24). No competing framework derives its background rate; all fit it to some band.
2. **The ladder as the predicted outlier.** On one background every probe without a calibrated rung reads 66.8, so the local excess is predicted to live in the ladder’s rungs; the one environmental

channel the framework could have credited is computed and not carried (Appendix G.5), and the band the corrected measurements are now filling is where the corrections drain, not yet where the framework says they end.

3. **Forced cross-domain tests.** The same integers forced a long list of other predictions (SM parameters, the Kaxion window, $\Omega_m = 1/\pi$, $\bar{\theta} \lesssim 10^{-18}$, DNA 4-letter alphabet, conscious episode ~ 400 ms), none of which could have been tuned after the fact without breaking the Hubble sector’s match.
4. **Sharp single falsifier.** Proton decay retires the entire program, including the $\epsilon_{\text{local}}^{(B)}$ piece. No comparable unique falsifier exists for the other six classes.

The landscape is therefore a very specific kind of test for GG-Theory: not whether it can fit a new data point — it fits none of them, in the fitting sense — but whether the *already committed* integer ledger stays coherent as the measurement world reorganizes itself around it. Three bands later, it does.

54 Summary: the landscape, the rivals, and the road to Part VI

Part V set out to put the derived reading in front of all the evidence and all its rivals at once, and this closing section gathers what it found.

The evidence first (§43). The 2026 measurement landscape is not two camps but three bands, and the paper now carries all of it: forty-five audited determinations, drawn (Figure 20) and tabulated (Table 13) with their identifiers, dates, and soft spots. Band 1 is overdetermined thermodynamics — three instruments and a CMB-free ruler audit agreeing because absolute densities leave nothing to conspire about — with the newest single-instrument value, SPT-3G’s 66.66 ± 0.60 , sitting on the derived rate. Band 2 is a destination, not a camp: corrected ladders, camp-free probes, and the ladder-free mean drain into 69–71 from both sides. Band 3 is two populations — one calibration family of formidable internal consistency, and the one-step geometric probes that are the standing counterweight to every bias reading, this paper’s included, and are carried here at full strength (§43.4).

The rivals second (§45–§50). Read against the landscape, the six classes sort cleanly. The systematics class has been promoted from shrug to named hypothesis — the paper’s own Part IV supplied the pricing — and strains only at the geometric counterweight. The early-side fixes are capped by walls their own machinery erected: EDE by the r_s -free span, modified recombination by its own 70.7 ceiling and the deuterium audit, extra radiation by BBN. Evolving dark energy points the wrong way. The deep void is disfavored by its own camp’s velocity data, while the modest void this paper adopts is alive and now carries a 6.4σ signature trend. None of the six reaches 73.5 any more without contradicting an audit; most reach exactly the band the corrected measurements are filling — which is why the middle world had to be added to the siren table.

The verdict machinery third (§51). One catalog now trisects the field: 67 confirms the derived reading and retires the early fixes; 73 kills this paper as designed; 70 — the world the middle band makes live — vindicates the corrected ladder, demotes the early fixes to band-2 physics, and leaves

this reading honestly strained at its one value, with the adjudication handed to the inventory and the shell test by its own rule. Every class's requirement is printed; nobody gets to renegotiate after the flyby.

What remains is the verdict's schedule and machinery: the kill switches, the dated decision calendar, and the siren forecast — the business of Parts VI and VII. The landscape is on the table; the clocks are running.

55 Part V in plain English: all forty-five measurements on one chart, and the six ways people have tried to close the gap

An argument about two numbers can be won by choosing which two. This Part refuses that move by putting every audited determination on a single chart — forty-five of them, from nine families of method, each read from its own primary record rather than from a review.

What the chart shows is that the tidy picture of two camps is gone. There are three bands and not two, and the middle band is populated by recent reanalyses that neither camp asked for. That is itself a fact any explanation now has to accommodate, and several of the standard explanations were designed for a world with only two camps in it.

The six classes of fix are then lined up against the same evidence: systematics in the ladder, early dark energy, modified recombination, extra radiation and interacting dark sectors, modified gravity with evolving dark energy, and a local void. Each is judged against audits that were not run for this paper's benefit — the deuterium abundance, the readings that need no sound horizon, the void camp's own velocity data. Each class runs into a wall its own machinery built. The systematics reading comes out of this Part considerably stronger than it went in, and it should be said plainly that it is a rival to this paper rather than an ally of it.

Nothing is settled here, and nothing could be. What emerges is that one future measurement now separates the survivors from one another cleanly, which is the business of the two Parts that follow.

Part VI

Kill Switches, the Decision Calendar, and Wider Implications

The eleven named falsifiers with their pre-declared rules; the dated decision calendar 2026–2035 and what the paper commits to at each date; and what the H0DN measurement means for the broader GG-Theory program

56 The falsifier registry: F-Hubble-1 through F-Hubble-11

A responsible interpretation of a new measurement must commit, in advance, to what would refute it. The GG-Theory reading of the H0DN result carries eleven named falsifiers, in four tiers — three at the program level, which retire the whole framework; four at the parameter-block level, which retire this paper’s reading of the Hubble tension alone; three at the channel level, which settle the individual residual-anatomy channels of §34 (the dust law, the calibration priors, the subsample stability) without touching the framework either way; and one at the coefficient level, which cashes the codebook selection of §24.2 against a computation that does not yet exist. Read operationally, and counted the way a careful reader would: five entries are direct kills of the derived reading (F-Hubble-2, -3, -4, -5, -6); three adjudicate *other groups’* corrections either way without touching the derivation (F-8/9/10); one refutes the dark-matter count doctrine without touching the rate (F-7); one is a theory window with no observation date (F-11); and one is the cohort’s standing single-event kill (F-1). Eleven registered rules; five gallows. (An earlier version counted four, with F-4 a pressure gauge on a local lift; the lift is withdrawn and F-4 is re-registered below as a kill.) We print the distinction so the registry’s breadth is never mistaken for its depth. Each is stated in the cohort’s F-Hubble- N convention with a *pre-declared numerical decision rule*: a threshold fixed here, before the data, so that the verdict cannot be negotiated after the fact. The tiers are not equal in what they stake: the channel-level trio adjudicates *external* allowances, quoted from the field’s own literature and never promoted to framework claims, so firing one of them reshapes the residual’s anatomy rather than the derivation — and the registry says so explicitly, entry by entry. We list all eleven, and draw the dated ones on the decision calendar of §58.

56.1 Program-level kill tests

F-Hubble-1. Proton decay. Observation of a single unambiguous proton-decay event at any rate retires GG-Theory in its entirety; the proton’s own account within the framework is GG-P6-Proton (Arisaka, 2026, P6-Proton). The exact local $U(1)_B$ of the GG-6 stage is what makes proton stability a theorem; a decay would retire the ledger and, with it, every number in this paper. Current Super-Kamiokande, Hyper-K, JUNO, and DUNE searches are relevant.

F-Hubble-2. A standard-siren H_0 near 73. The sharp, model-independent prediction of the derived reading is that the gravitational-wave standard siren — which uses neither a sound horizon nor a distance ladder — converges near 67 (Section 30). A siren determination that settled instead near 73, with adequate statistics, would mean the true background rate is high, the CMB densities cannot be divided by GG’s fractions to give it, and the derivation of Section 24 is wrong. *Pre-declared rule — one decision function, defending one hypothesis.* The hypothesis this rule defends is the *background* 66.8 ± 0.2 , the object the geometry derives and the one number this paper carries; a deep siren catalog and a shallow subset are both predicted to read it (OQ-Hubble-7). (An earlier version assigned a second, lifted reading to the shallow bright subsets; the lift is withdrawn and there is nothing for a shallow subset to test but the same background.) For a converged catalog of central value μ and precision $\sigma_{\text{cat}} \leq 2.2 \text{ km s}^{-1} \text{ Mpc}^{-1}$: *confirmed* if μ lies within $2\sigma_{\text{cat}}$ of 66.8; *refuted* if $\mu \geq 71.5$ with $(\mu - 66.8)/\sigma_{\text{cat}} \geq 3$; and for μ in the pre-declared *middle world* (68–71.5), the background is strained at the printed $(\mu - 66.8)/\sigma_{\text{cat}}$ and the adjudication passes, by this rule, to the sub-percent inventory determinations and the shell test (F-Hubble-4) — no renegotiation, the handoff is part of the rule. And one rung so that precision cannot hide inside the middle world: a catalog converging at *any* μ with $(\mu - 66.8)/\sigma_{\text{cat}} \geq 5$ *and* a host-shell profile flat in depth (no gradient for a bias to live in) refutes the background directly — a flat 5σ miss leaves no channel, and the rule says so now. One honest clause the confirmation arm must carry: a catalog at ≈ 67 confirms every low-rate reading at once — ΛCDM -plus-ladder-systematics included; distinguishing the derived 66.8 from the CMB combination’s 67.24 needs $\sigma_{\text{cat}} \approx 0.2$, beyond every scenario priced in Part VII. The kill is uniquely this reading’s exposure; the confirmation is shared, and we say so. The coming decade of catalogs will execute the rule; the full forecast — principle, pipelines, systematic budget, and timetable — is Part VII.

F-Hubble-3. A detected early-dark-energy lobe. The derived reading requires *no* early modification: the acoustic spectrum is fit by the π -locked fractions with the standard recombination history untouched. A joint Planck+SPT+ACT+DESI *detection* of a genuine early-dark-energy lobe — $f_{\text{EDE}} \gtrsim 0.02$ at $\geq 3\sigma$ near $z_c \sim 3500$ — would establish exactly the early-side new physics this reading says is absent, and refute the claim that the CMB endpoint is the unbiased derived background. The same holds for a radiation excess: the early side carries none at all — $\Delta N_{\text{eff}} \leq 0.03$ is a derived null (GG-A10-Cosmos, row D12) — and a detection at $\Delta N_{\text{eff}} \gtrsim 0.2$ kills the module as the lobe would. *Pre-declared rule:* a converged lobe detection at that level, or a radiation excess $\Delta N_{\text{eff}} \gtrsim 0.2$, refutes; a tightening null on either confirms.

56.2 Parameter-block-level kill tests

These tests retire this paper’s reading of the Hubble tension without automatically retiring the full program.

F-Hubble-4. A ladder-free late probe at the ladder’s value. The framework predicts that every probe without a calibrated rung — baryon acoustic oscillations on the derived r_d , the time-delay cosmography program, and the sirens of F-Hubble-2 — reads the background 66.8.

Pre-declared rule: a converged lensing-time-delay determination at 1% precision reading ≈ 73 refutes the one-background reading, exactly as a siren would (the present blinded TDCOSMO value, $71.6^{+3.9}_{-3.3}$, sits 1.5σ above 66.8 and adjudicates nothing yet); convergence on ≈ 67 confirms it. A controlled distance ladder that converges on the derived rate is the *expected end of the tension*, not a kill. And a ladder-shell profile flat in depth — a step in H_0 across $z_{\min} = 0.05$ consistent with zero at $\pm 0.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$ — confirms, as GG-A10-Cosmos reads it, that no environmental share of the excess exists; a measured step of $1\text{--}2 \text{ km s}^{-1} \text{ Mpc}^{-1}$ would re-open the question the parent closed. The entry carries one further clause on the early chain, because it is the same kind of test: a future independent determination of the decoupling redshift at the ± 5 level (post-2030 spectral-distortion missions, complemented by 21-cm cosmology) refutes the reading if the measured z_{dec} sits off GG-A10-Cosmos’s atomic-chain value 1091.5 by more than 1.4%, three standard deviations of the proposed measurement (present agreement +0.15%). (An earlier version of this entry made a ladder converging on ≈ 67 the pressure gauge on a local lift; under the one-background reading that convergence is the prediction, and the entry is inverted accordingly.)

F-Hubble-5. Detection of $w \neq -1$. The GG-A10-Cosmos prediction is $w = -1$ exactly, forced by the GSSB cap geometry. A detection of w different from -1 at the 3σ level by DESI Y5, Euclid, or Roman would retire the interpretation that ρ_Λ is the throat residue, and with it the exact- w leg of the derived expansion history.

F-Hubble-6. A primordial tensor amplitude inconsistent with $r = 0.0038$. The GG-A10-Cosmos prediction is the exact plateau value $r = 0.0038$ at $N^* = 52.6$ on the radion (Starobinsky/ ξ) plateau ($12/N^{*2} = 0.0043$ at leading order, the Starobinsky relation), with a spread $0.0037\text{--}0.0040$ from the declared half e-fold. *Pre-declared rule:* a Simons Observatory + LiteBIRD measurement outside $[10^{-3}, 10^{-2}]$ in either direction — a detection above 0.01, or a null pushing r below 10^{-3} — retires the interpretation of the early sector that anchors the derived history’s inflationary end. The tilt that comes with it, $n_s = 0.9634$, is not a registry entry of this paper but is printed beside it because it is the parent’s one currently failing test: the 2025 CMB-only value sits 1.5σ above it and the DESI-combined value 3.2σ , with no dial in the chain that moves it.

F-Hubble-7. A second Kaxion line in the window. The dark-matter companion predicts one laboratory line, standing alone, in the window $\nu_K \in [1.59, 2.13] \text{ GHz}$ on the DFSZ coupling line — the discrete Ξ_{cap} ladder, $m_K = 6.6\text{--}8.8 \text{ } \mu\text{eV}$, with the benchmark rung at $8.1 \text{ } \mu\text{eV}$ (1.95 GHz) (Arisaka, 2026, C1-Kaxion). The $1/35$ partner an earlier reading of this entry predicted is, under the boundary-condition termination of GG-A3-BI6, the longitudinal Z'_B at $\sim 10^{13} \text{ GeV}$, far above any receiver; its *absence* in the window is the prediction. *Pre-declared rule:* a second light QCD-coupled line in the window, established at the first line’s grade and at any spacing, refutes the count; if that second line sits specifically at $\Delta\nu/\nu = 1/35$ it is additionally evidence against the termination, which is a separate claim about a separate object; a window swept at DFSZ depth and silent kills the calibrated placement and its ladder (GG-C1-Kaxion’s F-Kaxion-1c); one line alone, in the window, confirms. Neither outcome bears on the derived rate, and neither sources ΔN_{eff} , which is a derived null.

56.3 Channel-level kill tests: the residual anatomy

These three tests settle the named external channels of §34 one at a time. They are deliberately graded below the parameter-block tier: every one of them adjudicates a *cited allowance* — a correction proposed inside the distance-ladder literature itself, never a GG-Theory result — so each verdict moves the anatomy of the $\sim 10\%$ systematic, not the derived rate. Their honesty value lies in the symmetry: each is two-sided by construction, and the confirming outcome is printed with the same prominence as the killing one.

F-Hubble-8. The near-infrared ladder and the dust channel. Dust nearly vanishes in the near-infrared, so a NIR supernova ladder strips away the dust-law freedom that the channel of §34 prices at 2–4%; and the accumulating ZTF/Rubin sibling samples test the two-population split directly, with statistics the calibrator set cannot match. *Pre-declared rule:* a published NIR/JWST supernova ladder, with the dust-law freedom removed, still converging at 73 kills the dust channel outright; the same ladder landing at 70–71 — where the published re-analyses put it (70.5 ± 1.0 , 70.59 ± 1.15) — confirms the channel and removes that share of the systematic. A ladder converging in the gap, 71.5–72.5, adjudicates nothing: the channel stays priced and undecided, and we say so now rather than negotiate it later. Channel-level only: either outcome leaves the derived rate untouched.

F-Hubble-9. Injection–recovery and the statistics channel. The distance-prior channel is priced at 0–3% and carried as *proposed, executed, and rebutted* — *undecided*. The test that settles it is cheap, decisive, and not yet published: injection–recovery simulation of the full ladder pipeline on synthetic data with known truth. *Pre-declared rule:* a published injection–recovery campaign validating the ladder’s priors at $< 1\%$ bias kills the statistics channel and zeroes its allowance; a demonstrated prior-induced bias at the 2–3% level confirms it; a demonstrated bias of 1–2% is a partial confirmation — the channel’s allowance shrinks to the demonstrated size, no more. Channel-level only; whoever runs the simulation settles the channel either way.

F-Hubble-10. The complete JWST sample and the subsample channel. The subsample channel (1–2%) lives in which calibrator hosts enter the fit: the same JWST-era data read 72.6 ± 2.0 under the SH0ES joint three-method analysis and 70.39 ± 1.94 in the CCHP program. *Pre-declared rule:* the complete JWST coverage of all ~ 37 SH0ES calibrator hosts, reduced under a tip/JAGB standardization convention agreed by both groups in advance, confirming 73.5 kills the subsample channel; a complete-sample value near 70.4 confirms the channel and locates $\sim 2 \text{ km s}^{-1} \text{ Mpc}^{-1}$ of the tension in calibrator selection. And the single most probable outcome deserves its sentence in 2026 rather than a negotiation in 2028: a complete sample near the SH0ES joint reading 72.6 ± 2.0 adjudicates *nothing* — the two conventions would still disagree inside their errors, and the channel would stay open, undecided, at its printed price. Channel-level only; the data are already being taken.

56.4 The coefficient-level test

F-Hubble-11. The vacuum coefficient off its pre-registered target. The codebook selection of §24.2 commits the framework to a number before the computation exists: the Wilson-line (Hosotani) evaluation of the vacuum residue on S_θ^1 (GG-A5-GDOS, Open Problem O6), when completed and checked, must return $\zeta_\Lambda = \pi/(LK) = \pi/175 = 0.017952$ to within $\pm 0.6\%$ (the tolerance the rate’s own ± 0.2 implies; equivalently H_0 within ± 0.2 of 66.8). *Decision rule:* a completed computation returning any other cluster member — or any value outside $[0.01784, 0.01806]$, a window that 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 the geometric route from a point back to the admitted stratum cluster of §24.2, whose nearest members to the observed decade are 66.8 and 74.8; a value inside confirms the one remaining conditional link in the derived rate. (An earlier version registered $\pm 1\%$, $[0.0177, 0.0182]$; that window contained $2/(K\pi)$, and the parent has since narrowed it.) The falsifier is two-sided by construction: the target is printed here, before the lemma, and cannot be renegotiated after it.

56.5 Kill-criteria summary

Table 17 compiles these tests.

The existence of F-Hubble-1 as a unique single-event kill criterion is the strongest distinguishing feature of GG-Theory as a cosmological framework. No other proposed resolution of the Hubble tension can be falsified by a single particle-physics observation.

57 The decisive observation: how the tension will be settled

Eleven falsifiers are listed above, but they are not of equal weight. One of them settles the central question — is the true background rate 67 or 73? — by itself, and it does so without a sound horizon and without a distance ladder. That is the gravitational-wave standard siren (F-Hubble-2). A binary-neutron-star or dark-siren catalog reads the luminosity distance straight off the waveform and the redshift off the host, so it measures H_0 through a channel that shares *no* systematic with either of the two discordant methods. The derived reading stakes everything on its verdict: $H_0^{\text{GW}} \rightarrow 67$, and a converged value at 73 refutes the whole construction (§56).

Around that decisive test sits a near-term program that will sharpen the picture well before the sirens converge. On the *early* side, a precision CMB determination (the Simons Observatory, then its enhanced large-aperture surveys) that drives the radiation term toward zero — the cohort carries a derived null, $\Delta N_{\text{eff}} \leq 0.03$ — *confirms* the derived reading rather than threatening it, and a detection at $\gtrsim 0.2$ refutes it, because the early side carries no excess at all; and a joint CMB+BAO search for an early-dark-energy lobe (F-Hubble-3) tests the no-early-modification claim directly. On the *late* side, the distance-ladder program itself — continued H0DN updates, and the JWST cross-checks of the Cepheid and TRGB calibrations (CCHP) — will show whether the local excess is stable or drifts toward the background as calibration effects are controlled; on this reading the drift is the expected end of the tension, and the kill sits with the ladder-free probes (F-Hubble-4, as re-registered). And the dark-energy equation of state (F-Hubble-5, DESI Y5, Euclid, Roman) and

Table 17: Kill criteria for the GG-Theory reading of the H0DN measurement, in four tiers. Program-level tests (F-Hubble-1–3) retire the entire framework; parameter-block tests (F-Hubble-4–7) retire this paper’s reading of the Hubble tension while leaving other sectors open; channel-level tests (F-Hubble-8–10) settle the residual anatomy’s individual external channels either way, reshaping the residual rather than the derivation; and the coefficient test (F-Hubble-11) cashes the codebook selection against a computation that does not yet exist.

Tag	Pre-declared decision rule	Facility or dataset	Structural target
F-Hubble-1	Observation of proton decay at any rate	Hyper-K, JUNO, DUNE	Exact local $U(1)_B$; retires whole program
F-Hubble-2	Converged standard-siren H_0 near 73 ($> 3\sigma$ from 66.8)	LVK O5, dark-siren catalogs	Derived background rate; the whole reading
F-Hubble-3	Detected EDE lobe $f_{\text{EDE}} \gtrsim 0.02$ at $\geq 3\sigma$, or $\Delta N_{\text{eff}} \gtrsim 0.2$	Planck+SPT+ACT, DESI Y5	The no-early-modification claim
F-Hubble-4	Converged ladder-free time delay at ≈ 73 (1%); or z_{dec} off the atomic chain by $> 1.4\%$	TDCOSMO/LSST lenses; post-2030 distortion missions	One background; a converging ladder confirms
F-Hubble-5	Detection of $w \neq -1$ at 3σ	DESI Y5, Euclid, Roman	Throat-cap ρ_Λ interpretation
F-Hubble-6	Primordial tensors outside $[10^{-3}, 10^{-2}]$ (stake 0.0038)	Simons Obs., LiteBIRD	Radion-plateau inflation
F-Hubble-7	A second Kaxion line in the window at any spacing (at 1/35, also evidence against the termination), or a swept silent window	ADMX, HAYSTAC, DMRadio	The count (one line alone) and the placement
F-Hubble-8	NIR/JWST SN ladder (dust freedom removed) still at 73	JWST NIR programs; ZTF/Rubin siblings	Dust/two-population channel (2–4%); channel-level
F-Hubble-9	Injection–recovery validating ladder priors at $< 1\%$ bias	full-pipeline simulation on synthetic truth	Distance-prior/statistics channel (0–3%); channel-level
F-Hubble-10	Complete ~ 37 -host JWST sample (agreed convention) confirming 73.5	JWST (SH0ES + CCHP hosts)	Calibrator-subsample channel (1–2%); channel-level
F-Hubble-11	Hosotani computation returns ζ_Λ outside $[0.01784, 0.01806]$	theory (GG-A5-GDOS O6), pre-registered target $\pi/175$	The codebook selection; Route I as a point rather than a bracket

the primordial tensor ratio (F-Hubble-6, Simons Obs., LiteBIRD) test the two ends of the same inventory the rate is read from.

These tests do not each carry a date so much as a data era, and the next section lays them — together with the channel-level trio, which *does* carry near dates — on one forward axis, entry by entry, with what this paper has committed to at each (§58). The wager is dated, and the dates are near. Nor is the wager this paper’s alone: as §51 tabulates, the same converged catalog bisects the entire solution space of Part V — every class has printed, or has been forced here to print, which verdict it needs.

58 The decision calendar, 2026–2035

A registry states rules; a calendar states dates. This section walks the falsifiers in the order the data will actually arrive, because the discipline of pre-registration is only as good as the specificity of the appointment: a rule with no date can be deferred forever, and a date with no printed number can be renegotiated when it comes. Neither escape is available here. Figure 21 draws the whole schedule in the cohort’s kill-test-timeline convention — three eras, each test card naming its rule and venue, the violet siren band carrying the decisive F-Hubble-2 — and the two subsections below spell out the entries and the commitments, date by date.

58.1 The dated entries

December 2026 — Gaia DR4. The nearest and sharpest date on the calendar, and the floor of the ladder: DR4 improves the parallaxes by $\sqrt{t}$ ($\times 1.4$) and — decisively for the anchors — maps the bright-star zero-point, with full epoch astrometry, that every Milky-Way calibrator inherits. The anatomy of §34 carries the Gaia entry at 0.5–1% with the sign of the induced H_0 shift explicitly unestablished ($-38.9 \pm 10.3 \mu\text{as}$ at bright G , formal errors understated); DR4 reads that entry either way. Physically this is the purest appointment on the calendar: astrometry knows nothing of cosmology, so whatever DR4 returns is a statement about the measurement chain alone — which is exactly why it can move the anatomy and cannot move the derivation.

2026–2028 — the channel-level trio. The injection–recovery simulation (F-Hubble-9) has no scheduled date because it needs no telescope: it is a software campaign any group could publish in a season, and the calendar carries it as *open now*. The complete JWST coverage of the ~ 37 SH0ES hosts (F-Hubble-10) is in progress, with the agreed standardization convention the real schedule risk — the data will exist before the convention does. The NIR ladder and the ZTF/Rubin sibling samples (F-Hubble-8) accumulate through the late 2020s. These three appointments are the cheapest on the calendar relative to what they settle: together they adjudicate the full quadrature-to-stack span of Part IV’s two scenarios (≈ 4 –9%).

2027–2029 — the Local-Hole census and the shell-level $H_0(z)$. Euclid and DESI count galaxies straight through the Hole, fixing its depth profile from the density side, while the peculiar-velocity programs read the outflow a void of that depth would drive — between them they measure




Figure 21: The decision calendar as a kill-test timeline. The dated falsifiers of §56 drawn on a forward time axis (2026–2035+, broad data-era bands rather than precise dates), in the cohort’s kill-test-timeline convention. The decisive test is F-Hubble-2 (violet, bordered): a converged standard-siren H_0 landing at 73 rather than 67 refutes the derived reading outright; the violet band is the siren accumulation, with the O5 segment marked. Alongside the program- and parameter-block-level tests, the calendar carries the channel-level trio F-Hubble-8/9/10 (blue cards, 2026–2028 era), which adjudicates the residual anatomy of Part IV either way, and the Gaia DR4 appointment (December 2026), which reads the ladder’s parallax floor. The stake card (top right) collects the printed numbers that cannot move. F-Hubble-11, the coefficient test, is a theory milestone with no observation date, tracked in Table 17; F-Hubble-1 (proton decay, gray) is the single-event kill live at every moment.

whether any share of the ladder’s excess is environmental. The framework carries no lift and predicts none: a shell profile flat in depth confirms the parent’s withdrawal, a measured lift at 2–3% would re-open the question (§33.1). And one standing entry is already on the board: the descending $H_0(z)$ trend reported at 6.4σ in DESI’s uncorrelated baryon-acoustic bins (Jia et al., 2025), carried with the three provenance cautions of §33.1 (the third: a high- z tail would be adverse at F-Hubble-5), deepens or dissolves with each DESI release. A trend that sharpens as the low shells are cut away is what a local lift would look like from inside; a constant-bias ladder, which is what this paper predicts, imprints no such shape.

2028–2030 — the early-side precision era. The full SPT-3G survey, the Simons Observatory, and DESI DR3 tighten every early-side number this paper stakes: the early-dark-energy search (F-Hubble-3) reaches the $f_{\text{EDE}} \sim 0.02$ discovery floor, the radiation term is driven toward the cohort’s null, $\Delta N_{\text{eff}} \leq 0.03$, and the reference combination $H_0^{\text{ref}} = 67.24 \pm 0.35$ sharpens above the derived 66.8 ± 0.2 — which must continue to sit one coherent π -offset below it, or the reading fails at its own anchor. Alongside, DESI’s $w(z)$ program (F-Hubble-5) approaches the ± 0.02 regime where $w = -1$ exactly is either locked in or broken.

2028–2030 — O5 and the siren accumulation. The decisive channel (F-Hubble-2). The O5 observing run is the first with a realistic path toward a converged percent-class siren value, and the dark-siren catalogs sharpen meanwhile — the current best, $69.9^{+4.1}_{-4.0}$, already sits 3.1 above the derived rate, inside its own error. The timeline’s violet band carries this accumulation; Part VII prices it event by event, pipeline by pipeline.

The 2030s — the tensor verdict. LiteBIRD (launch ~ 2032 ; the full-survey $\sigma(r) \approx 10^{-3}$ arrives toward the decade’s end) and the Simons Observatory’s enhanced large-aperture surveys read r at the 10^{-3} floor. The ground-based partner originally written into this test, CMB-S4, is no longer available: DOE and NSF ended their support in July 2025, and the prediction now rests on the two programs above. The prediction is two-sided (F-Hubble-6): $r = 0.0038$ exact on the plateau (0.0043 at leading order), and data outside $[10^{-3}, 10^{-2}]$ in either direction retire the early-sector anchor of the derived history. The honest date is later than the tension literature’s habitual “early 2030s,” and this calendar says so rather than inheriting the habit.

Standing entries, undated. Proton decay (F-Hubble-1) is live at every moment in Hyper-K, JUNO, and DUNE — the single-event kill. The Kaxion window (F-Hubble-7) is open to any haloscope that reaches 1.59–2.13 GHz. And the one theory milestone, F-Hubble-11, waits on the Hosotani computation: no observation can advance it and no date binds it, but its target — $\zeta_\Lambda \in [0.01784, 0.01806]$ — has been on the record since §24.2.

58.2 What the paper commits to at each date

A standing adverse condition, registered. One exposure carries no single date and belongs at the head of this ledger: the early anchor itself. The derived rate requires H_0^{ref} to continue sitting one coherent π -offset *above* 66.8; if the DESI-weighted early-side combinations continue their present drift and the consensus anchor converges above ≈ 68 , the π -offset reading is strained at its own anchor — $\geq 2\sigma$ at the paper’s structural error, with no siren required. We register the condition here, against interest, because the trend is in today’s data and a reader should not have to discover it.

Pre-registration is a ledger of hostages. For each appointment above, this paper’s exposure is a specific printed number, and it is worth collecting them in one place — because after the datum arrives, none of them can move.

At Gaia DR4 (December 2026): the 0.5–1% floor allowance and its *unestablished* sign. A mapped zero-point near zero deletes the anatomy’s floor entry, and the stack shrinks accordingly; a confirmed offset in the helpful direction cashes it. Neither outcome touches the derived rate; both readings are printed here in advance.

At the channel-level verdicts (2026–2028): the three allowances 2–4%, 0–3%, and 1–2%, and with them the two scenario totals of §34.5 — quadrature 4.3%/3.8%, stack 9.3%/8.3%. Every channel killed moves the corrected-ladder landing point back toward 73.50; every channel confirmed moves it toward the published chain’s ≈ 70.6 (69.5–70.2 fully chained). The middle band of Part V is the scoreboard on which this plays out in public.

At the early-side era (2028–2030): $\Delta N_{\text{eff}} \leq 0.03$ (no detection at $\gtrsim 0.2$); no early-dark-energy lobe above $f_{\text{EDE}} \approx 0.02$; $w = -1$ exactly; and the derived 66.8 ± 0.2 one coherent π -offset below whatever H_0^{ref} sharpens to. Any of these failing is a parameter-block or program-level kill, not a reshuffle.

At the siren convergence (O5 and beyond): 67, not 73, under the 3σ rule of F-Hubble-2 — the single cleanest exposure in the register, and the one this paper has staked its central claim on since Part III.

And at every moment in between: no proton decays; and one Kaxion line, alone, if any haloscope reaches the window first. The calendar closes where the registry began: eleven rules, all printed before their data — the posture a derived number can afford, and a fitted one cannot.

59 Over-determining the curve: chronometers, calibrated rulers, and the redshift drift

The decisive test adjudicates one point — the present value. But the derived reading’s exposure is far larger than one point, because what Part III derived is the whole function $H(z)$, and three channels can read that function directly, with neither a sound horizon nor a distance ladder anywhere in the chain. The strongest strategy against a derived curve is not to argue about its endpoint but to measure it everywhere — as the community call of §40 anticipated — and every value these channels will confront is already fixed.

Cosmic chronometers. A passively evolving galaxy population is a clock: its stars age at a known rate, so the differential age of two populations separated by a small redshift interval reads the expansion directly. From $1 + z = a_0/a(t)$, differentiating along the line of sight gives, in two lines,

$$\frac{dz}{dt} = -(1+z) \frac{\dot{a}}{a} \implies H(z) = -\frac{1}{1+z} \frac{dz}{dt}, \quad (47)$$

no cosmology assumed, no ruler, no ladder — galaxies as clocks (Jimenez and Loeb, 2002; Moresco et al., 2022). Current compilations reach five-to-fifteen-percent precision per point and are consistent with the derived curve; a percent-class chronometer program would arbitrate the whole reading by itself.

Calibrated rulers. Baryon acoustic oscillations measure the combination $H(z) r_d$. Ordinarily r_d is fitted; but the derived chart *computes* it — the drag-epoch ruler $r_d = 147.2$ Mpc on the derived background (the CMB landmark itself is $r_* = 144.7$, landmark 4; App. E.2) — so every BAO measurement becomes an absolute reading of $H(z)$ with no freedom left on either side of the product. Table 18 lists what the curve requires at the redshifts these surveys actually probe.

The redshift drift. The purest channel of all watches a single redshift change in real time. For a comoving source, $1 + z = a(t_0)/a(t_e)$; observe again after Δt_0 , and both epochs advance, giving (to first order)

$$\dot{z} = (1+z) H_0 - H(z), \quad \Delta v = \frac{c \Delta z}{1+z} = c \Delta t_0 \left[H_0 - \frac{H(z)}{1+z} \right] \quad (48)$$

(Loeb, 1998; Sandage, 1962). This is cosmology with no proxy at all: no standard candle, no ruler,

no clock but the observer’s own — the expansion history watched, not inferred. On the derived curve the numbers are fixed: $\Delta v = +2.1 \text{ cm s}^{-1}$ per decade at $z = 1$, falling through zero at $z \approx 1.9$, and -5.6 cm s^{-1} per decade at $z = 4$ (Table 18). Centimeters per second per decade is brutal territory, but it is exactly the design goal of the ELT’s high-resolution spectrograph (Liske et al., 2008) and of SKA-era absorption monitoring; a drift map inconsistent with the 66.8-anchored curve would kill the derived reading independently of every other channel in this paper — including the siren.

Table 18: The derived expansion history at the redshifts the curve-reading channels probe. Every entry is a zero-parameter output of the π -locked inventory (Eq. (85)); the drift column is Eq. (48) over one decade of observer time. Consistency of future chronometer, calibrated-BAO, and drift data with this column set is an independent kill test of the whole reading.

z	$H(z)$ derived [$\text{km s}^{-1} \text{ Mpc}^{-1}$]	Drift Δv per decade	Who reads it
0.5	88.6	$+2.4 \text{ cm s}^{-1}$	chronometers; DESI/Euclid BAO
1.0	120.1	$+2.1 \text{ cm s}^{-1}$	chronometers; BAO; ELT drift
2.0	203.5	-0.3 cm s^{-1} (sign change at $z \approx 1.9$)	Lyman- α BAO; ELT drift
4.0	425.3	-5.6 cm s^{-1}	ELT/SKA drift

Emerging channels. Two young channels are approaching the same curve from unexpected directions. Localized fast radio bursts — whose dispersion measures integrate the electron column along the line of sight — already return $H_0 = 69.40^{+2.14}_{-1.97}$ from 108 events (Gao et al., 2024), squarely between the camps and tightening with every localization; and γ -ray attenuation on the extragalactic background light reads $61.9^{+2.9}_{-2.4}$ (Domínguez et al., 2023) — anomalously low and strongly EBL-model-dependent, but a channel with no ladder and no sound horizon anywhere in it. Neither is yet decisive; both mature on the calendar’s timescale, and the derived curve is already committed to what they must find. Meanwhile the first curve-level *trend* datum exists: the descending $H_0(z)$ at 6.4σ in DESI’s uncorrelated bins (Jia et al., 2025), read with the provenance cautions of §33.1 — the generic signature of some local outflow, and a shape no constant-bias ladder can imprint.

The strategic point deserves one plain sentence. A fitted Λ CDM can absorb a surprising $H(z)$ point by re-fitting; the derived curve cannot absorb anything — every row of Table 18 is a place where the reading can die, and that multiplicity of exposure, not any single agreement, is what a derivation buys.

60 Observational priorities opened by the H0DN result

The H0DN result transforms the observational agenda because it locks in the local band (6) at 1%. The next pressure points, in order of potential impact on the GG-Theory reading — a ranking updated by the anatomy of Part IV, which promotes the two channel-level tests that can settle named allowances outright above everything except the joint-likelihood confrontation:

60.1 Priority 1: Joint CMB+BAO+ H_0 analysis within the inherited ledger

The most urgent task is a fully joint CMB+BAO+ H_0 likelihood analysis with the background *fixed* at the π -locked inventory — zero background parameters, only the seed amplitude A_s and the optical depth τ inherited (the tilt is derived, $n_s = 0.9634$) — run head-to-head against six-parameter Λ CDM and against EDE. This is the over-determined test of §26 promoted to full-likelihood level: if the derived chart’s χ^2 per *free* parameter beats the fitted alternatives, the case stops being philosophical.

60.2 Priority 2: The injection–recovery campaign (F-Hubble-9)

Promoted to second place on its cost-to-consequence ratio: a software-only campaign, publishable in a season, that settles a 0–3% allowance currently frozen at “undecided.” No other entry on this list resolves an open channel so cheaply, and the verdict is symmetric by construction (§56.3).

60.3 Priority 3: The near-infrared ladder and the sibling samples (F-Hubble-8)

Promoted alongside: dust nearly vanishes in the near-infrared, so the NIR/JWST supernova ladder is the one photometric program that can kill or confirm the largest single allowance (2–4%) in the residual’s anatomy, and the ZTF/Rubin sibling samples test the two-population split with statistics the calibrator set cannot match.

60.4 Priority 4: Increasing the number of geometric anchors

The H0DN paper explicitly suggests (Section 6 of Casertano et al., 2026) that “establishing additional anchors via geometric methods (e.g., using population-II pulsating stars calibrated by parallax)” is a key direction. From the GG-Theory perspective, more geometric anchors sharpen the confrontation: they narrow the measured local band around its stable value, and every halving of its width makes the predicted $\approx 10\%$ systematic harder to dissolve quietly — either the controlled ladder drifts toward the background ≈ 67 , the expected end of the tension on this reading, or the systematic stands and the siren adjudicates (F-Hubble-2).

60.5 Priority 5: Megamaser distances beyond NGC 4258

Five galaxies with megamaser systems (UGC 3789, NGC 6264, NGC 6323, NGC 5765b, CGCG 074-064) currently anchor the megamaser Hubble-flow tracer class. Their individual uncertainties are $\gtrsim 10\%$. Additional Megamaser Cosmology Project work (or SKA-era AGN surveys) that brings another five to ten systems into the network would both tighten the central H0DN value and, crucially, further independently test the no-systematic hypothesis. The derived reading expects them to agree with the network consensus: the stability of the local value is a measurement-chain fact, whatever its ultimate origin.

60.6 Priority 6: JWST extension of the primary ladders

JWST observations of Cepheids, TRGB, and JAGB in SN Ia host galaxies provide the first fully independent check of HST-era Cepheids in these hosts (Freedman et al., 2025; Riess et al., 2024). The H0DN paper finds (Section 5.2.2 of Casertano et al. 2026) that JWST-including variants favor H_0 slightly higher than HST-only variants, but within 1σ . From the GG-Theory viewpoint, these JWST data provide the instrument-swap hard test that confirms $\epsilon_{\text{local}}^{(B)} > 0$ is not an HST artifact. The completion of this program across all ~ 37 hosts, under a convention agreed in advance, is the registry’s F-Hubble-10.

60.7 Priority 7: DESI $w(z)$ constraints

If DESI DR3 or DR4 reports a statistically significant detection of $w(z) \neq -1$ at 3σ , this would activate F-Hubble-5. If DESI confirms $w = -1$ to within ± 0.02 , this would firmly lock in the throat-cap interpretation of ρ_Λ — the exact- w leg of the derived expansion history.

60.8 Priority 8: Kaxion haloscope search in the window

ADMX, HAYSTAC, and DMRadio should search for dark matter in the unique laboratory window predicted by GG-C1-Kaxion (Arisaka, 2026, C1-Kaxion), $\nu_K \in [1.59, 2.13]$ GHz ($m_K = 6.6\text{--}8.8$ μeV ; benchmark 8.1 μeV , 1.95 GHz). One line standing alone is the confirmation; a second line at $\Delta\nu/\nu = 1/35$ would refute the count and count against the termination of the gauged baryon symmetry; a window swept at DFSZ depth and silent kills the placement (F-Hubble-7).

61 Wider implications: S_8 , early galaxies, and the H0DN program

Three wider readings follow from the same ledger. None of them is load-bearing for the Hubble claim, and each is carried at exactly the weight the framework can afford — one derived number, one consistency statement, and one caveat.

61.1 The S_8 tension

We have so far focused on H_0 . The H0DN paper itself does not address S_8 . But the same ledger that fixes the background rate must, by the cross-consistency argument of Section 5, also fix a specific S_8 .

In the derived reading there is no early-time suppression to invoke (Section 39): the growth history is standard growth on the derived inventory, and the S_8 prediction is one number, $S_8^{\text{GG}} = 0.825 \pm 0.010$ (Eq. (45), App. I.5), computed in the parent GG-A10-Cosmos from the declared amplitude. It sits 2.5σ above the benchmark cosmic-shear determination $S_8 = 0.759^{+0.024}_{-0.021}$ (KiDS-1000; Asgari et al. 2021) and 0.4σ below the CMB- ΛCDM inference: unlike the local H_0 , the S_8 tension is not read as an environmental question — the framework shares it with ΛCDM , and a combined shear-plus-clustering determination that excludes 0.825 at three standard deviations, with the CMB-normalized value moving away from it too, would falsify the amplitude chain and, since that chain has no dial, the background (GG-A10-Cosmos).

61.2 JWST-era early galaxies

The JWST early-massive-galaxy puzzle (unexpectedly massive galaxies at $z \sim 10\text{--}14$) is a high- ν tail problem, and the derived reading treats it the way it treats the growth sector generally: standard structure formation on the derived inventory, with no Hubble-specific mechanism invoked. The equality epoch that seeds the small-scale tail is fixed by the same ledger ($z_{\text{eq}} \simeq 3400$ from $\Omega_m = 1/\pi$ and the measured radiation density; App. K.5), so the growth clock is not adjustable. The quantitative confrontation — halo-mass functions and baryon-feedback modeling at $z \sim 10$ — is carried by the parent GG-A10-Cosmos on the same inventory; this paper imports no credit from it, and stakes no claim on it beyond consistency.

61.3 The meaning of the agreement for the H0DN program itself

Even setting aside GG-Theory, the H0DN collaboration itself suggests at the end of (Section 6 of Casertano et al., 2026) that “rather than serving solely to constrain dark energy models, as envisioned a decade ago, the improved accuracy of H_0 now exposes a broader inconsistency within the standard cosmological framework and strengthens the case for new physics or a deeper reassessment of early-Universe inferences.”

GG-Theory is a concrete candidate for what that “new physics” looks like. In particular:

- The new physics is not a late-time dark-energy species.
- The new physics is not a single scalar field.
- The new physics is not an anomalous local void.
- The new physics is the *same* integer ledger that already fixes the Standard Model parameters — applied to the cosmic inventory, from which the whole expansion history follows with nothing further added.

If GG-Theory is right, the Hubble tension is not a failure of our cosmology but a mis-stated question: the background rate was never a free number to be measured twice, but a derived function read at one point — and the residual disagreement is a statement about one measurement chain, decidable by a channel that owes nothing to either side.

One caveat travels with the word “consensus” wherever this paper uses it: the H0DN combination excludes the CCHP program, whose JWST-era values (70.39 ± 1.94 , and lower on the JWST-only subsamples) define the low edge of the ladder’s own internal spread (§43.2). What H0DN certifies is the covariance-weighted agreement of the teams inside it; the subsample channel of Part IV lives precisely at the boundary of that “inside.”

62 Cosmological benchmark sheet

For ease of reference, we collect the complete set of GG-Theory-inherited cosmological predictions in Table 19, with current observational status.

Table 19: GG-Theory-inherited cosmological benchmarks and their observational status as of 2026. Status labels follow the conditional triage of the companion paper **GG-A10-Cosmos** (Derived / Derived-conditional / Asserted); entries quoted from it carry that conditional status, and the local excess is carried whole as a calibration systematic the framework predicts and does not derive; the void channel is priced and not carried (Appendix G.5).

Observable	GG-Theory prediction	Current measurement	Status
Local ladder H_0	66.8, read $\approx 10\%$ high by a predicted calibration systematic	73.50 $\pm$ 0.81 km s $^{-1}$ Mpc $^{-1}$ (H0DN)	Systematic predicted, not derived; its anatomy priced and not yet sufficient (Part IV)
Background $H_0^{(0)}$ (derived)	66.8 $\pm$ 0.2 km s $^{-1}$ Mpc $^{-1}$ (two-road 66.6 $\pm$ 0.2)	67.24 $\pm$ 0.35 (CMB); 498.3 $\sqrt{\zeta_\Lambda}$ (seesaw)	Derived (cond.) / D*; -0.65% (1.1σ) from the CMB combination, -0.9% from Planck's 67.4
ΔN_{eff}	≤ 0.03 (derived null, GG-A10-Cosmos row D12)	0.04 $\pm$ 0.18 (Planck)	Consistent; a detection $\gtrsim 0.2$ kills
Standard-siren H_0^{GW}	≈ 67 (predicted)	(forthcoming)	The decisive model-independent test
2026 landscape structure	corrections move band 3 (72.5–77) into band 2 (69–71), not to 67	re-analyses land at 70–71.5; chained corrections 69.5–70.2 (Part V)	Consistent; the middle band is filling
Corrected-ladder scenario H_0	69.5–70.2 (published central corrections chained)	CCHP 70.39 $\pm$ 1.94; two-population 70.59 $\pm$ 1.15	2.7–3.8 above the derived 66.8; predicted to drain further
Local-void channel ϵ_{Hole}	0 carried (response law computed)	depth contested; velocity data disfavor the counts branch	Priced, not carried (GG-A10-Cosmos)
w (dark energy eq. of state)	-1 exactly	-1.03 ± 0.03 (DESI Y3)	Compatible
Ω_m	$1/\pi \approx 0.3183$	0.315 $\pm$ 0.007 (Planck)	0.47 σ
Ω_Λ	$1 - 1/\pi \approx 0.6817$	0.685 $\pm$ 0.007 (Planck)	0.47 σ (approx.; neglects Ω_r , massive ν)
$\rho_\Lambda^{1/4}$ (dark energy)	$\sim$ few meV	(2.24 $\pm$ 0.02) meV	Consistent
r (primordial tensors)	0.0038 (exact plateau; 0.0043 leading order)	< 0.032 (BI-CEP/Keck)	Consistent; a factor eight below the bound; kill window $[10^{-3}, 10^{-2}]$
$\sin^2 \theta_W(M_Z)$	25/108 $\approx$ 0.23148	0.23149 $\pm$ 0.00013	-0.07σ
$\hat{\alpha}_{\text{EM}}^{-1}(M_Z)$	$2^J = 128$	127.93 $\pm$ 0.008	$+0.05\%$ (fractional-agreement convention, as for z_* in GG-A10-Cosmos; $\gg 1\sigma$ at the quoted precision)
Proton lifetime τ_p	∞ (absolute stability)	$> 2.4 \times 10^{34}$ yr	Consistent
$\bar{\theta}$	$\lesssim 4.7 \times 10^{-18}$	$< 10^{-10}$ bound	Consistent
Ω_m/Ω_Λ crossover $z_{\Lambda m}$	≈ 0.289	0.30 $\pm$ 0.02	Consistent (implied by the Ω_m row; not independent)
S_8	0.825 $\pm$ 0.010 (computed from A_s ; Derived-cond. on Lemma V', GG-A10-Cosmos)	0.759 $^{+0.024}_{-0.021}$ (KiDS-1000), 0.832 $\pm$ 0.013 (CMB- Λ CDM)	Shares the tension: $+2.5\sigma$ vs shear, -0.4σ vs CMB
JWST massive-galaxy excess	standard growth on the derived inventory	observed 147	Consistent (quantitative case in GG-A10-Cosmos)
Redshift drift at $z = 1$	$+2.1$ cm s $^{-1}$ per decade	(ELT/SKA era)	Prediction of the derived curve
Kaxion line	one line, alone, in	not yet probed	Pending; a second line refutes the

The H0DN agreement sits at the top of Table 19 as the newest entry, but it is only one of twenty-one benchmark rows. The distinctive feature of the GG-Theory reading is this consistency across all entries, not any one of them individually.

Coda. Part VI has turned the reading into a schedule. The falsifier registry now runs to eleven named rules — program, parameter-block, channel, and coefficient — and the decisive one is still a converged standard siren, which reads the clock a third way, owing nothing to either endpoint, and lands near 67 or near 73. New on the board is the calendar itself: Gaia DR4 on December 2026; the channel-level trio through 2028; the early-side precision era to 2030; the sirens beyond — each date with its printed number, and each number a hostage. That is the posture a derived quantity can afford and a fitted one cannot: to name, in advance and in public, the observation that would end it. And beyond the one decisive point, the whole curve is on the schedule: chronometers, rulers calibrated by the derived sound horizon, the young fast-radio-burst and γ -ray channels — and, most beautifully, the redshift drift, the expansion watched in real time at centimeters per second per decade, each a place where the derived history has already committed and cannot move.

63 Part VI in plain English: the eleven ways this reading can be killed, and the dates they are scheduled for

A derived number is worth less than the list of things that would refute it, so this Part is that list, with dates attached.

Eleven named falsifiers, in four tiers. Three would retire the whole framework rather than merely this paper. Four would retire this paper's reading of the Hubble tension and leave the framework standing. Three settle individual channels of the residual without deciding anything about the framework in either direction. One cashes the choice of a single coefficient against a computation that nobody has yet performed.

Then the calendar, which is what makes a registry more than a gesture. Gaia's fourth data release in December 2026. The channel-level trio through 2028. The precision era on the early side out to 2030. The sirens after that. Each date carries a number this paper has already committed to, and the commitment is public before the measurement, which is the only order in which a commitment counts for anything.

The Part then opens the work outward: eight observational priorities, running from a joint likelihood that any group could run today with already released data, to a haloscope search in a frequency window the dark-matter companion fixes. It closes on a benchmark sheet of twenty-one rows, because the honest way to present a framework is on all of its entries at once and not on the one it does best.

Part VII

Forecasting the Verdict: The Standard-Siren Decade

The physics of the siren; the referee’s report card 2017–2026; the events-to-verdict forecast on empirical rates; the referee’s own systematic budget; and the three worlds the catalog can choose among

64 The principle: a distance from the waveform alone

Everything in this paper funnels to one measurement, so the measurement deserves a Part of its own: what a standard siren is, why it owes nothing to either discordant route, how many events the verdict needs, and what could bias the referee itself. The idea is forty years old (Schutz, 1986) and almost indecently clean. A compact binary in its final inspiral is a broadcasting machine whose absolute output is computed, not calibrated: general relativity fixes the entire waveform in terms of the source masses. The observed frequency f and its acceleration $\dot{f}$ give the *chirp mass* through

$$\dot{f} \propto f^{11/3} \mathcal{M}_c^{5/3} \quad \implies \quad \mathcal{M}_c \text{ read from the phasing alone,} \quad (49)$$

and the strain amplitude then contains only one unknown:

$$h \propto \frac{\mathcal{M}_c^{5/3} f^{2/3}}{d_L} \quad \implies \quad d_L \text{ read from the amplitude,} \quad (50)$$

an *absolute* luminosity distance, from first principles (Holz and Hughes, 2005). No parallax, no period–luminosity relation, no supernova standardization: the siren is a one-rung ladder whose single rung is the Einstein field equations. The redshift the waveform cannot supply (a redshifted binary looks exactly like a heavier one) comes from electromagnetic astronomy — the host galaxy of the counterpart — and at low redshift the rate follows from the two directly: $cz \simeq H_0 d_L$.

This has already been done once. GW170817, the first binary-neutron-star inspiral, came with a kilonova, a host galaxy (NGC 4993), and a Hubble-flow velocity of $3017 \pm 166 \text{ km s}^{-1}$; against the waveform’s distance of roughly forty megaparsecs the joint posterior gave $H_0 = 70_{-8}^{+12} \text{ km s}^{-1} \text{ Mpc}^{-1}$ (Abbott et al., 2017) — one event, fifteen-percent precision, straddling both camps with perfect impartiality. The entire question of this Part is how fast that fifteen percent divides by $\sqrt{N}$.

65 Three pipelines, disjoint by construction

Bright sirens are the GW170817 channel: an inspiral with an electromagnetic counterpart, hence an exact host redshift. They are rare — binary neutron stars with observable kilonovae — and each

carries a $\sim 10\text{--}20\%$ distance error dominated by a single geometric degeneracy: a binary seen face-on is louder, so inclination and distance trade against each other in the amplitude. The degeneracy is broken partly by the two gravitational-wave polarizations and, for the lucky events, decisively by electromagnetic viewing-angle information from the jet afterglow.

Dark sirens are the volume channel: binary black holes with no counterpart, redshifted statistically against galaxy catalogs over the localization volume. Each event is weak, but LVK-class runs collect them by the hundreds; the founding forecasts (Chen et al., 2018) promised percent-class H_0 within years of design sensitivity, and this Part dates the channel from its *observed* record instead (§66). Its systematics — catalog completeness, population priors — are entirely different from the bright channel’s.

Spectral sirens use features of the source mass distribution itself as a redshift reference; they are the youngest channel and serve here as a third cross-check. The structural point is the disjointness: none of the three shares a systematic with the sound horizon, with the Cepheid zero-point, with supernova standardization, or — crucially — with each other’s dominant error. Three referees, mutually independent, all reading the same clock.

66 The referee’s report card, 2017–2026

A Part built on a forecast owes the reader the record before the extrapolation. As of this writing the record is nine observing years long, and it reads: *one bright siren*.

GW170817 remains, in 2026, the only gravitational-wave event with a confirmed electromagnetic counterpart. The subsequent observing runs detected further binary-neutron-star candidates, but none yielded a kilonova identification: the events sat several times farther away, their sky localizations spanned tens to hundreds of square degrees, and a kilonova fades through the follow-up window in days. This is the O4 null-counterpart record, and its physics lesson is worth stating plainly: the bright channel’s bottleneck was never the interferometers. It is the compound accident of a nearby, well-oriented merger caught while wide-field telescopes are looking. GW170817, at forty megaparsecs, was a gift; the volume in which a kilonova is *cheap* to catch grows far more slowly than the volume in which the waveform is *loud* enough to detect, so detector improvements alone do not manufacture counterparts.

The dark-siren ladder, meanwhile, has been climbing exactly as the forecast’s structure — though not its dates — said it would. Three rungs are on the record: $76.6^{+13.0}_{-9.5}$ from the GWTC-4.0 catalog (142 dark sirens against the GLADE+ galaxy catalog, with GW170817; mass-spectrum-model dependent) (LVK Collaboration, 2025); $71.0^{+9.0}_{-7.1}$ from GWTC-5.0 (236 sources against DES-Y6 — a 25.7% tightening) (LVK Collaboration, 2026); and $69.9^{+4.1}_{-4.0}$ from seventeen well-localized O4a dark sirens against DESI Legacy photometric redshifts, again with GW170817 (Alfradique et al., 2026) — the sharpest siren value in existence. The precision is marching down as catalogs deepen, and the central value has walked from the ladder camp to within an error bar of the derived rate. One honesty clause belongs beside that observation: the three values share events and methods, so the walk is a single maturing measurement seen at three stages, not three independent confirmations — suggestive, and no more. (A fourth entry exists and is carried as exactly one line: thirteen binary

black holes with *claimed* AGN-flare counterparts give $70.5^{+3.4}_{-2.9}$, but the associations are unconfirmed, and this paper takes no support from it (Kumar et al., 2026).)

Then the recalibration, printed rather than papered over. The rate forecasts of the GW170817 era imagined tens of counterpart events within a few years of design sensitivity; nine years produced one. The dates in this Part are therefore drawn from the empirical record, not from those forecasts, and the bright-alone verdict slips accordingly — into the O5-to-next-generation era. What does *not* slip is the arithmetic: the twenty-two- and sixty-one-event thresholds of Eq. (51) depend only on the hypothesis separation and the per-event error, and are hostage to no rate assumption at all. A forecast’s dates are hostage to Poisson luck; a forecast’s thresholds are hostage to nothing. The rest of this Part keeps the two strictly apart.

67 The forecast: events to verdict

The arithmetic of the decision is elementary, and stating it exactly is the point of this Part. The two hypotheses sit $\Delta = 73.5 - 66.8 = 6.7 \text{ km s}^{-1} \text{ Mpc}^{-1}$ apart — ten percent. A catalog of N unbiased bright sirens with per-event fractional error σ_{ev} reaches $\sigma(N) = \sigma_{\text{ev}} \times 70/\sqrt{N}$, and separates the hypotheses at k standard deviations when $\sigma(N) = \Delta/k$:

$$N_{k\sigma} = \left(\frac{70 \sigma_{\text{ev}} k}{\Delta} \right)^2, \quad \sigma_{3\sigma} = 2.23, \quad \sigma_{5\sigma} = 1.34 \text{ km s}^{-1} \text{ Mpc}^{-1}. \quad (51)$$

Table 20 evaluates it for three per-event scenarios, and Figure 22 draws it: at the GW170817-like central scenario the 3σ verdict needs *twenty-two* bright events and the 5σ verdict *sixty-one*. Even the first five events of the O5 era deliver a 1.4σ lean — the community will watch this arrive in public, event by event. One scope clause, kept because an earlier version of this Part needed it and the record should show why it no longer does: the 22/61 is the separation of the background 66.8 from the ladder 73.5 ($\Delta = 6.7$), and it is the only separation the paper now prices. An earlier version carried a second, shallower one — a bright-siren subset at $z \lesssim 0.05$, sampling inside the Local Hole, where a conditional lift would have been a coherent common mode that does not average down, with its own pre-registered rule and a branch test assigned to the ladder shells. With the lift withdrawn there is no branch to decide: a shallow subset is predicted to read 66.8 like the deep catalog, and a shallow subset that reads high would be the one instrument that could re-open the environmental question. The ladder-shell test (F-Hubble-4, as re-registered) carries the confirmation: a step in H_0 across $z_{\text{min}} = 0.05$ consistent with zero at $\pm 0.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$ confirms, as the parent GG-A10-Cosmos reads it, that no environmental share exists; a measured step of $1\text{--}2 \text{ km s}^{-1} \text{ Mpc}^{-1}$ would re-open the question the parent closed.

Rates convert events to dates, and here the report card of §66 governs: the dates below come from the empirical record, not from the GW170817-era rate papers. Nine observing years have produced one counterpart event, so the honest bright-siren prior is of order one event per year in the O5 era *at best* — the deeper range multiplies the detection volume severalfold, but the counterpart bottleneck (kilonova depth, localization area, follow-up coverage) has proven the harder constraint. Under that prior the bright channel alone reaches its 22-event 3σ crossing only in the mid-to-late

Table 20: Events to verdict (Eq. (51); open-source script `c2_siren_forecast.py`). Bright-siren counts; dark-siren catalogs contribute in parallel and, on the empirical catalog-to-catalog improvement rate of §66, reach the same precision on a similar clock.

Scenario	Per-event error (justification)	$N_{3\sigma}$	$N_{5\sigma}$
Optimistic	10%: jet viewing-angle constraints on most events	10	27
Central	15%: GW170817-like, polarization-only inclination	22	61
Conservative	20%: poor localization, no viewing-angle help	39	109

2030s, at the O5-to-next-generation boundary. What has not slipped is the clock’s other hand: the dark-siren channel, accumulating hundreds of black-hole events per run and already at the ± 4 level (§66), reaches the two-percent class after three to four further catalog generations at the observed improvement rate — with the honesty clause attached: only one generation step is cleanly measured (GWTC-4.0 $\rightarrow$ 5.0, 25.7%; the further step to ± 4.05 mixes in a method change), so compounding it ($\pm 4.0 \times 0.74^{3-4} \approx \pm 1.2\text{--}1.6$) is a scaling from the record but still a two-point extrapolation, and we call it that. The early-to-mid 2030s on the current run cadence, dated from what has happened rather than from the 2018-era forecasts this Part retires for the bright channel. The timeline bands of Figure 21 therefore stand, with the near-term weight carried by the dark channel: first verdicts near 2030, the decisive era by ~ 2035 .

One outcome deserves its own sentence: a catalog that converges *between* the camps. This is not a separate verdict but one cell of F-Hubble-2’s single decision function — at $\mu = 70.0$, $\sigma_{\text{cat}} = 1.0$, both point hypotheses fail their 3σ tests simultaneously (3.2σ and 3.5σ), and the rule’s middle-world clause takes over: the strain on the pure background is printed, and the adjudication passes to the inventory and the shell test. That outcome would tell the field something more interesting than either camp expects — that the early-side inference and the late-side calibration are *both* displaced — and the most informative outcome remains the one nobody is betting on.

68 The referee’s own budget: can the siren itself be wrong by ten percent?

A paper that has spent a Part interrogating the ladder’s error budget owes the same scrutiny to its chosen referee. Table 21 walks the siren channel’s systematics with the one question that matters: could any of them manufacture a coherent ten-percent bias — the size needed to fake a verdict?

The table’s verdict is short. Every astrophysical term either averages down with the catalog or is printed above as a named coherent term with its mitigation and its size — and the named terms top out at a few percent, an order below the ten-percent question the referee is asked to settle. The one irreducible common mode is the amplitude calibration of the detectors themselves — currently at the one-to-two percent level, measured with photon calibrators in the laboratory, and required to reach sub-percent for the 5 σ era. It is an engineering quantity, with no conceivable coupling to the sound horizon, the Gaia zero-point, or the Cepheid period–luminosity relation. There is, simply, no candidate channel by which the siren could be coherently wrong by ten percent — and if one is

Table 21: The systematic budget of the referee. “Averages down” indicates a per-event term that shrinks as $1/\sqrt{N}$. Three terms do not: the detector amplitude calibration (the channel’s true floor, an instrument property with no conceivable coupling to either discordant route); the source mass-spectrum model, coherent across the dark/spectral channel; and, for shallow bright subsets only, a local-void monopole, which the framework carries at zero.

Source	Nature	Per event	Floor
Inclination–distance	geometric degeneracy; broken by polarizations and jet viewing angle	10–30%	averages down
Host peculiar velocity	$\sim 300 \text{ km s}^{-1}$ against the flow: 10% at 40 Mpc, 1% at 450 Mpc	1–10%	averages down
Detector amplitude calibration	photon-calibrator laboratory measurement; <i>common mode</i>	1–2%	$\sim 1\%$, must be engineered down
Selection (Malmquist)	louder events over-represented; modeled with the population	$\sim 1\%$	sub-% with population model
Dark-siren catalog completeness	missing hosts skew the statistical redshift	—	$\sim 1\%$, the dark channel’s floor
Weak lensing	line-of-sight magnification; grows with z	$< 1\%$ at low z	averages down
Source mass-spectrum model	the $\sim 35 M_{\odot}$ feature’s location and redshift evolution set the dark/spectral distance scale; mis-modeling biases coherently — the $76.6 \rightarrow 71.0 \rightarrow 69.9$ walk of §66 is in part this term relaxing; mitigated by population marginalization and the bright-siren cross-check	few %	coherent
Local-void monopole	not carried (§33); a shallow-host subset at $z \lesssim 0.05$ would read it if it existed, and 153 is predicted to read the background; absent for deep	$\leq$ few % if present	coherent (shallow)




Figure 22: The verdict, forecast. *Left:* catalog precision $\sigma(N) = \sigma_{\text{ev}}/\sqrt{N}$ for three per-event scenarios, against the pre-declared decision thresholds: $\sigma = 2.23$ separates the derived 66.8 from the ladder 73.5 at 3σ , $\sigma = 1.34$ at 5σ , and the central-scenario crossings are $N = 22$ and $N = 61$ (dots; Eq. (51)). The shaded band is the O5-era bright catalog under the *empirical* rate record of §66 (one counterpart in nine years) rather than the GW170817-era forecasts — the one assumption of the released script revised by that record; the thresholds and per-event scenarios are rate-independent. Dark-siren catalogs carry the near-term clock in parallel. *Right:* the same forecast as the community will experience it — the toy posterior after 22 (dashed) and 61 (solid) events under each hypothesis. Whichever pair materializes, the other reading is finished; a catalog converging *between* them lands in the pre-declared middle world of F-Hubble-2’s decision function — the background strained at the printed $(\mu - 66.8)/\sigma_{\text{cat}}$, adjudication handed to the inventory and shell tests. Script: `c2_siren_forecast.py`, released with the paper.

nervous anyway, the bright, dark, and spectral pipelines carry *different* floors and must agree. The referee is as clean as instruments get.

And the budget is no longer only a table — it is being demonstrated live. The dark-siren walk of §66, $76.6 \rightarrow 71.0 \rightarrow 69.9$, is Table 21’s middle rows seen in the wild: the several-unit swings between consecutive catalogs are the mass-spectrum-model and catalog-completeness terms at their current size, and their trajectory — error bars shrinking from $^{+13.0}_{-9.5}$ to $^{+4.1}_{-4.0}$ as the catalogs deepen and the host samples improve — is what “averages down” looks like empirically. A referee whose systematics are visibly shrinking on schedule, in public, is precisely the referee one wants; the spread between today’s dark-siren values is the size of the remaining budget, and it is already smaller than the gap it is being asked to adjudicate.

69 The verdict, simulated — and what failure looks like

The right panel of Figure 22 runs the decade forward as a toy simulation (seeded, elementary, and released): the posterior the community will actually watch, after 22 and after 61 central-scenario events under each hypothesis. The curves are deliberately unglamorous — two Gaussians marching apart as $\sqrt{N}$ — because that is what a clean decision looks like: no reanalysis, no tension indices, no committee. One pair of curves materializes, and the argument of a decade ends.

Failure is drawn on the same axes, and this paper reads it without flinching. If the red pair materializes — a catalog converged at 73.5 — then the true background rate is high; the CMB densities cannot be divided by the π -locked fractions to give it; Theorem T-Hubble-1’s inventory, whatever its geometric elegance, does not describe this universe’s budget; and the derived reading of this paper is finished, cleanly, in public, by the exact instrument it spent a Part praising. The personal note of the Conclusions was written with that outcome squarely in view. What the framework may not do — and commits here not to do — is discover, after such a verdict, that some neglected effect had shifted its prediction to 73 all along. The number is 66.8 ± 0.2 . It was printed before the catalog existed. There is nowhere to hide, which is precisely the point.

70 The sharpened stake: three outcomes, three worlds

Part V tabulated the three worlds for the six rival classes (Table 14); this section prices them for this paper alone, with the catalog precisions made explicit. Everywhere below, “a converged catalog” means one standing between the two verdict thresholds of Eq. (51), $\sigma_{\text{cat}} = 2.23 \rightarrow 1.34 \text{ km s}^{-1} \text{ Mpc}^{-1}$; every verdict in this section is a cell of the single decision function pre-registered at F-Hubble-2, which defends the background 66.8 — the one number this paper carries; a deep catalog and a shallow subset are both predicted to read it (OQ-Hubble-7).

At ≈ 67 — confirmation, and the tension ends as metrology. F-Hubble-2 confirms the derived background; the residual anatomy of Part IV graduates from priced hypothesis to history. The channel-level registry entries F-Hubble-8/9/10 stay interesting even then, because they would describe *how* a superb measurement chain came to read ten percent high — and the tension would end the way discrepant constants have most often ended and are least often predicted to: as metrology, not new physics. The oldest number in cosmology becomes what Part III said it was — a theorem’s output, read three independent ways.

At ≈ 70 — the middle world, priced exactly. The H0DN centroid falls and the corrected-ladder cluster of band 2 is vindicated. Here the rule’s honesty matters most, so we state it against the temptation a middle value invites: a converged catalog reads the *background*, so its verdict falls on the derived 66.8 ± 0.2 — strained at $1.4\text{--}2.4\sigma$, the exact figure depending on where between the thresholds the catalog stands — and there is no second, lifted reading for the catalog to vindicate instead; an earlier version of this section offered one, and it is withdrawn with the lift. At ≈ 70 , then: the background is pressured but not killed at the pre-registered level, the question of whether any share of the excess is environmental is re-opened by the data rather than by this paper, and the shell-level $H_0(z)$ program of §58.1 adjudicates it independently of any siren. The middle world is uncomfortable for everyone, this reading included; the discomfort is printed here in advance, with its numbers.

At ≈ 73 — refutation, as designed. F-Hubble-2 fires; the anatomy evaporates; the derivation dies in public, by the instrument this Part spent its pages praising (§69). Nothing in that sentence is softened by its being unlikely on the framework’s own reading: the rule was built so that this world, if it is the real one, wins quickly and unarguably.

Three worlds, one instrument, and a rule fixed before the catalog exists. No fitted resolution of

the tension can offer the same sentence.

71 Reproducibility and the pre-registered rule

Everything quantitative in this Part reproduces from one elementary, fully documented script (`c2_siren_forecast.py`, released with the paper): the thresholds of Eq. (51), the event counts of Table 20, and both panels of Figure 22. Its inputs are three declared numbers — the two hypotheses and the per-event error scenario — and nothing else; an independent group can regenerate the forecast, or sharpen it with a full waveform-level analysis, in an afternoon. The rule it serves is the pre-registered F-Hubble-2, restated once, verbatim: a converged siren value $H_0^{\text{GW}} = 73 \pm 1$ — a $> 3\sigma$ miss of the derived 67 — refutes; a value 67 ± 1 confirms; a value between the camps kills both readings and is the most informative outcome of all. The rule was fixed before any converged catalog existed, and this Part has now fixed the timetable on which it will be applied.

Coda. Since the first neutron stars locked into their death spirals, the universe has been broadcasting absolutely calibrated distances — a distance ladder with no rungs, built into the field equations themselves, transmitting for billions of years at frequencies no one could hear. For a century we built ladders of starlight because we could not receive that signal; now we can. Somewhere in the coming decade the catalog — bright, or dark and statistical — will pass its twenty-two-event equivalent, and then its sixty-one, and the posterior will settle onto one of the two curves this Part has already drawn, or between them, into the third world of §70. The oldest number in cosmology will then have been read a third way, by an instrument that owes nothing to any ladder, any horizon, or any theory — including the one whose fate it decides.

72 Part VII in plain English: the one instrument that owes nothing to either side, and when it will have spoken

Both routes to the Hubble constant are calibrated — one against stars, one against a sound horizon — and both calibrations are under suspicion. A third route is not calibrated at all.

When two neutron stars spiral into each other, general relativity fixes the entire waveform from the masses alone. The frequency and the rate at which it climbs give the masses; the masses give the intrinsic loudness; the observed loudness then gives the distance, with no ladder and no ruler anywhere in the chain. Pair that distance with a redshift and the result is the Hubble constant, measured by an instrument that has never heard of a Cepheid or an acoustic scale.

This Part does the accounting for that instrument. What the catalog has actually achieved since 2017, which is less than the enthusiasm suggests. How many events a verdict needs, computed from empirical event rates rather than from optimism. Whether the siren itself could be wrong by ten percent, which is asked seriously and answered with a budget rather than a reassurance.

And then the stake, fixed in advance and in public. A converged siren near 67 confirms this reading. A converged siren near 73 ends it. A value between the two camps kills both readings at once and is, for that reason, the most informative of the three outcomes. The rule was written before any converged catalog existed, and it is not available for renegotiation once one does.

Part VIII

Discussion: Significance, Uniqueness, and Innovation

What a derived expansion history means as a scientific claim; the problems it resolves and the moves it registers; the honest boundary of what it computes; and its open questions, predictions, and the road ahead

73 What this Part does

Parts III through VII did the work: derived the history, priced the residual, drew the landscape, dated the falsifiers, and forecast the verdict. This Part steps back and weighs the result as a scientific claim, in five movements, in the order a skeptical reader should meet them. *Grounding* (§74–§75): what kind of object this paper is — one derivation chain, read on one contested observable — and what the H0DN comparison actually establishes, graded strong, moderate, and untested. *Significance* (§76–§78): the case stated in ten points, read against the leading responses to the tension, and placed against the historical precedents — including the cautionary ones. *Accomplishments* (§79–§80): the standing puzzles resolved or dissolved, and the structural moves this paper makes that no prior treatment of the tension has made, each in before-and-after form. *The honest boundary* (§81–§85): exactly what is proved, what rides on a named condition, what the framework computes and what it hands to other instruments, and what the whole posture costs. *The road ahead* (§86–§89): the open questions stated as a numbered program, the forward predictions and recommendations addressed to the five communities that can act on them, and the conduct the paper binds itself to while it waits. A reader who wants the paper’s case in one Part should read this one; every claim below carries a pointer to the Part that earns it.

74 The essence: one chain, one observable

Before the particulars, it is worth naming what kind of object this paper is, because its shape is inherited, not improvised. The GG-Theory engine is a single chain: OGGE (axiom AX-3 of the GG-Constitution; [Arisaka, 2026, A1-Constitution](#)) declares every observable a boundary projection of a closed six-dimensional geometry; GDOS ([Arisaka, 2026, A5-GDOS](#)) is the law by which the open boundary evolves; the GRIP cascade $DGI \rightarrow GSSB \rightarrow \text{attractor}$ ([Arisaka, 2026, A6-GRIP](#)) selects the unique ground geometric attractor; and the observable *is* that attractor’s boundary reading. Run on the particle stratum, the chain returns the Standard Model’s twenty-eight parameters (GG-A9-SM). Run on the cosmological stratum, it returns the π -locked inventory and the whole background history (GG-A10-Cosmos). This paper is the chain run all the way down to a *single*

contested number: the same engine, specialized to the one observable the field has fought over for a decade, with every intermediate step exposed.

That siting explains the paper’s two most distinctive features at once. First, why the rate could be derived at all: $H(t)$ is not an extra postulate but the Friedmann reading of an inventory the chain had already fixed — the derivation cost nothing new, which is exactly what a genuine consequence should cost. Second, why the paper can afford its honesty: because the chain fixes the fractions *jointly*, the derived 66.8 ± 0.2 cannot be nudged toward 73.5 without breaking the acoustic fit, the age, the equation of state, and the sister papers simultaneously (§26, Table 11). A fitted model meets an unwelcome measurement by adjusting; this construction can only be right or wrong about it. Everything else in this Part — the significance claims, the innovations, the open questions, even the conduct clauses — is downstream of that single structural fact.

The chain, read on the two-fiber anatomy. The cohort’s recent campaign (Arisaka, 2026, W1-Symmetry; Arisaka, 2026, P5-CPT; Arisaka, 2026, A10-Cosmos) compresses the bulk’s non-spatial pair into one exact statement: *holonomy requires closed loops, and an interval has none*. The compact circle S^1_θ carries all of the bulk’s holonomy — quantized, protected, exact at every scale; the scale interval I_r carries none — only one-way transport, a semigroup. Read on that anatomy, this paper’s siting is no accident. The Hubble rate is the *interval’s* observable: expansion is the descent read in boundary time — cooling *is* descent, boundary time grows as $d \ln t / dr = 2$ in the radiation era, and $H = 1/(2t)$ is that slope inverted — so the question this paper answers is, structurally, *how fast does the fiber that forgets forget*. And everything that *pins* the derived curve is the *circle’s*: the π -locked fractions $\Omega_m = 1/\pi$ and $\Omega_b = 1/(2\pi^2)$ are geometric measures of the compact fiber, and the corridor’s switch sits at $z_* = e^J$ — a ledger integer in the exponent, exact and unable to run, planted in the middle of the chronology. That is the anatomy behind the two structural facts named above: the rate could be *derived* because rates are interval-side — they belong to the fiber whose whole content is its one-way history, which a derivation can own end to end — and it could be derived *without a dial* because everything pinning it is circle-side, where nothing is continuous enough to tune. A decade-long tension between two readings of one number is then no scandal of the instruments: it is what the ledger looks like while one rate is read through two dictionaries, one of which carries a calibrated rung; one background, read by every rung-free instrument, is what reconciliation looks like when the anatomy is read whole.

75 Is H0DN a good support for GG-Theory?

We have argued, in detail, that it is. Let us state the case as cleanly as possible in one place, separating the strong, weak, and merely plausible components.

75.1 What is strongly supported

1. The derived background $H_0^{(0)} = 66.8 \pm 0.2$ lands on the independent CMB-inferred 67.24 ± 0.35 within the coherent π -offset, and the same fractions reproduce the full acoustic spectrum at $\chi^2 = 187$ with zero background freedom (§26) — fixed before the H0DN comparison was ever

made, and never adjusted. The h -free baryon-loading ratio $1/2\pi$ agrees at 1.7σ with no freedom whatsoever.

2. The H0DN variant catalog shows a stable local value that no single method, anchor, or instrument drives. That stability is what the derived reading expects if the local value is one background plus a network-level calibration systematic — and not what a per-method fluke would produce. One boundary caveat is carried with this item (§43.2): the variant catalog certifies stability *within* the H0DN membership, which excludes the CCHP program and its 70.39 ± 1.94 ; the subsample channel of Part IV lives at exactly that boundary, and F-Hubble-10 is its dated resolution.
3. The orthogonal-path test retires the *single-method* unknown-systematic hypothesis. What it does not retire — and the paper leans on this honestly — is a network-level shared calibration bias, which is exactly where the predicted $\approx 10\%$ systematic is located, and exactly what the ladder-free siren decides (F-Hubble-2).

75.2 What is moderately supported

1. The S_8 tension is shared, not resolved (Section 39): the computed 0.825 sits between the CMB-normalized and shear-normalized amplitudes, 2.5σ from the latter, and the framework offers no environmental reading of the gap. Observational status is unsettled pending fuller analyses of KiDS+DES+HSC data.
2. The JWST early-galaxy excess is qualitatively consistent with the earlier, more efficient structure growth on the derived inventory; the quantitative match depends on uncertain baryon-feedback modeling at $z \sim 10$.
3. The decisive open test is the standard-siren H_0 , predicted near 67; a value near 73 would refute the derived reading.

75.3 What is not yet tested

1. No local-void lift is carried, and the under-density has not been measured on the ladder’s own weight function. A direct measurement — for instance by cross-correlating galaxy density with peculiar-velocity flow in a $\gtrsim 1$ Gpc volume — would fix whether any share of the ladder’s excess is environmental (the framework predicts none); and the predicted $\approx 10\%$ systematic is a claim about the measurement chain that dedicated network-level calibration audits can probe directly (Appendix G.5).
2. The Kaxion window has not been probed. One line, alone, is the sharpest confirming test; a second line at $\Delta\nu/\nu = 1/35$ the sharpest refutation of the count.
3. $w = -1$ exactly is confirmed at the 1σ level by DESI DR2 but remains to be tightened. A positive deviation at 3σ would activate kill criterion F-Hubble-5.

76 Significance and uniqueness, in ten points

The reading of this paper differs from every other response to the Hubble tension in a way that is worth stating as a single running argument, in ten numbered points. Each is a claim of *kind*, not of degree: not that this paper does the usual things better, but that it does things the usual responses do not attempt.

1. *A derived rate, not a fitted one.* Every other response to the tension begins by measuring H_0 and then explaining it; this paper derives it and then confronts the measurements. The derivation arrives twice over — Route I from the geometry’s own vacuum arithmetic, $498.3\sqrt{\pi/175} = 66.8 \pm 0.2$, with no astronomical input whatsoever, and Route II from the measured physical densities divided by the derived fractions, 66.6 ± 0.2 — and the two agree at 0.6σ (Part III). The number was in print, under these covers, before the H0DN precision existed to test it. In a field where “prediction” usually means “postdiction with a likelihood,” that ordering is itself the contribution.
2. *An expansion history, not a number.* The framework predicts the whole function $H(t)$, of which H_0 is one reading: the landmark chronology, the derived $H(z)$ at every redshift the curve-reading channels probe (Table 18), the drag-epoch ruler $r_d = 147.2$ Mpc that turns every BAO point into an absolute measurement (with $r_* = 144.7$ Mpc as the CMB landmark), the age 13.9 Gyr, the de Sitter floor $H_\infty = 55.2$, even the redshift drift at centimeters per second per decade. The tension is thereby demoted from a crisis about a constant to a question about the calibration of a single point on a known curve — and the curve, unlike the point, can be checked in a dozen places at once (§59).
3. *Forced, not installed.* The inventory behind the rate is a theorem on one Hopf fibration (T-Hubble-1): $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/2\pi^2$, three coupled readings of one geometric object. They cannot be moved independently, so the derived rate cannot be tuned toward either camp — there is no dial, and the anti-numerology discipline of the cohort (one ledger, one naming rule, no per-observable formula freedom) forecloses the search-until-it-fits escape. A reader who suspects numerology is invited to the codebook selection of §24.2, where the one remaining coefficient was chosen by structural rules and then pre-registered as a falsifier (F-Hubble-11) — the opposite of a fit.
4. *The early side needs no new physics.* The derived background lands on the CMB value with a coherent -0.9% π -offset; no early-dark-energy lobe, no extra radiation, no modified recombination is invoked to raise it. The claim is not merely that none is needed — it is that their absence is *predicted*: a tightening EDE null and a ΔN_{eff} bound closing on zero confirm this reading while they starve its rivals (F-Hubble-3, and the class-by-class ledger of Part V). The early universe, on this reading, was never broken; it was read correctly the first time.
5. *Over-determined, not merely consistent.* Fed the exact fractions with no background parameter tuned to the sky, the machinery reproduces the acoustic spectrum at $\chi^2 = 187$ with zero

background freedom, returns the damping scale $\ell_D = 1350 \pm 3$, and passes the h -free baryon-loading check at 1.7σ (§26). Seventeen observables miss coherently, in one direction, by one small step — the signature of a derivation displaced by a single named correction, not of a lucky fit, whose errors would scatter on both sides. Agreement one number at a time is cheap; agreement this correlated is the expensive kind.

6. *Honest where it cannot reach.* The whole $+10\%$ of the local excess is named, not derived; the one first-principles channel the framework owns ($U(1)_B$) is closed by the boundary-condition termination of GG-A3-BI6 and, independently, *excluded* by eight orders of magnitude, the void channel is priced and not carried, the averaging channel is bounded below 10^{-3} , and the backreaction escape is priced at a $> 12\sigma$ density fluctuation (§35). The candidate anatomy that remains is quoted from the field’s own literature at the field’s own error bars, graded Asserted-external, and priced in both directions — quadrature 4.3%/3.8%, stack 9.3%/8.3% — with the requirement sitting above both (§34.5). No other response to the tension prints the case against itself at this resolution.
7. *Cross-ledger exposure.* The integer ledger (1, 3, 5; 7, 35) that fixes the inventory also fixes the Standard Model’s twenty-eight parameters (GG-A9-SM) and the Kaxion spectrum at $8.1 \mu\text{eV}$ (GG-C1-Kaxion), so the Hubble reading is coupled to falsifiers in arenas that have nothing to do with cosmology: a haloscope, a collider, a proton-decay tank. Tuning any entry to rescue the rate would break a sibling in another paper. No fitted resolution of the tension can be shot at from three fields at once; this one can, and calls that a feature.
8. *Uniquely falsifiable, on a printed calendar.* Eleven named falsifiers with pre-declared numerical rules, four tiers deep, drawn on a dated decision calendar with the paper’s exposure at each date collected in one place (§56, §58.2); a single proton-decay event retires the entire program. Falsifiability here is not a philosophical credential but an operational schedule — the difference between “this could in principle be tested” and “here is what December 2, 2026 will do to this paper.”
9. *The landscape as an instrument.* Part V is not a literature review; it is a published measuring device: forty-five determinations, every value read from its primary record on a stated date, asserted in a released script, drawn in one chart and tabulated with identifiers (Fig. 20, Table 13). The device produced a finding no single-measurement argument could: the field has three bands, the middle one is filling, and every proposed resolution can now be asked one uniform question — *why three bands, and why is the middle one filling?* — that fitted responses struggle to answer and the derived reading answers structurally. The instrument outlives the paper: anyone can re-run the audit, add a value, and re-ask the question.
10. *The conduct as content.* The paper’s manner is part of its claim: values pre-registered before their tests; the referee’s own record printed against the forecast (one bright siren in nine years, §66); the forecast’s dates re-based on the empirical record while its thresholds — hostage to nothing — stand; confirmation statements printed with the same prominence as kill rules; and no relocation clauses anywhere. In a controversy that has consumed a decade partly because

every party could revise its story after each datum, a reading that *cannot* revise — structurally, not rhetorically — changes what kind of argument is possible. That, as much as any number, is what this paper offers the field.

77 Comparison with the leading responses

Part V read the six rival classes against the evidence, band by band, and rendered class-by-class verdicts (Table 15); that work is not repeated here. What belongs in a Discussion is the *explanatory trade-off* — what each habit of response does well, what it pays, and which missing layer this paper’s reading supplies. Read this section as theory choice, not scorekeeping.

“Measure better and wait.” The field’s default posture, and an honorable one: no new physics, no commitment, more data. Its strength is safety; its price is that it treats the tension as noise to be outlasted rather than information to be read — it asks no question the data could answer structurally, and after a decade it has none of the three bands explained. The derived reading keeps the data-first discipline (the calendar of Part VI *is* “measure better,” dated) but adds the missing layer: a stake, so that the measurements decide something.

Early-side fixes (EDE, modified recombination, dark radiation). Their strength is seriousness: they take the tension as real physics and modify the sound horizon to close it. Their structural price is now visible in the landscape: they need the true rate high, so the middle band starves them — each new corrected ladder or dark-siren value near 70 leaves them explaining a gap that is quietly closing from the other side; and the sound-horizon-free measurements (66.7–70.2) cap how much ruler-shortening the data will tolerate. The missing layer is a reason the early universe should be modified *at all* when a zero-freedom inventory already fits it at $\chi^2 = 187$. This paper’s early side is not a fix but an absence of one — and that absence is a two-sided prediction (F-Hubble-3).

Late-side dynamics (interacting dark sectors, modified gravity, evolving dark energy). Their strength is breadth: enough freedom to bend $H(z)$, growth, or the duality relation wherever the tension bites. The price is the freedom itself — each new sector carries potentials, couplings, and clustering choices whose number grows faster than the anomalies they explain, and the DESI-era $w(z)$ hints they lean on currently point in the *wrong direction* for raising the late rate (Part V). The missing layer is economy: a principle that says which of the infinitely many possible late-time modifications nature chose. The derived reading’s late side is one exact statement — $w = -1$, no evolution — and it is falsifiable this decade (F-Hubble-5).

Local-void explanations. The nearest relative, and the comparison the paper embraces: a real under-density really does lift locally measured rates. The habit’s price has been wanting the void to do *everything* — the deep-void readings needed to close the whole gap are now disfavored by field-level velocity modeling (§33.1). The missing layer is restraint with a derivation behind it: this paper derives the response law, finds the depth that would let a void carry any share of the excess unobserved, carries no lift, and hands the whole excess to a different mechanism rather than

stretching the void to cover it. The weakest viable member of the class, on purpose — and, since the parent’s withdrawal, not a member at all.

Systematics-only skepticism. The mirror image of the early-side fixes: the ladder is simply wrong somewhere, and the tension is bookkeeping. Its strength is that the middle band keeps feeding it — the published corrections are real and land near 70. Its price is shapelessness: without a derived background it cannot say what the corrected ladder *should* find, so even total victory leaves the rate unexplained. This paper contains a disciplined copy of the skeptic (the anatomy of Part IV) — but bolted to a derivation, so that every correction has a predicted endpoint (66.8 ± 0.2 ; the published chain reaches 69.5–70.2 and stops short) and the skepticism itself becomes falsifiable (F-Hubble-8/9/10).

The one-line summary. Each rival is strong exactly where it spends freedom, and weak exactly where it has no layer to constrain that freedom. The derived reading is the only entry that spends *no* freedom — its strength and its risk are the same fact, and the rest of this Part is an extended audit of that trade.

78 Historical precedents — including the cautionary ones

Pre-registered numerical predictions that hit $< 1\sigma$ on first test are rare in the history of physics, and it is worth placing this one among them honestly — which requires citing the failures as well as the triumphs.

General relativity and Mercury’s perihelion (1915). Einstein’s calculation produced the anomalous $43''/\text{century}$ from first principles, matching a long-standing measurement within $0.3''$, with nothing tuned. The parallel is the logical form: a number forced by a geometric theory, confronting a precision measurement that predated it. The disanalogy is candor about scope: Mercury was one number from a framework that went on to pass every test for a century; the derived 66.8 is one number from a framework whose siblings are still on the calendar.

The W and Z masses (1983). The electroweak theory predicted both masses in a narrow band once $\sin^2 \theta_W$ was measured; UA1/UA2 confirmed at $< 1\sigma$, and the confirmation was rightly decisive because the same structure had already organized dozens of other measurements. That is the standard the cohort aims at: the Hubble rate as one entry in a ledger that also carries the weak angle — literally, $\sin^2 \theta_W = 25/108$ at -0.07σ (Table 19) — so that the confirmation, if it comes, is correlated the way the W/Z confirmation was.

The cautionary precedents. The history of dramatic agreements also contains the OPERA superluminal-neutrino episode — a 6σ effect that was a loose cable — and the BICEP2 announcement, where a genuine signal was real dust; and, closer to this paper’s subject, a century of Hubble-constant history in which each era’s value carried confident error bars that excluded the next era’s value (Part II). Three lessons are taken, not merely noted. First, *coherence over significance*: the paper’s

CMB case rests on seventeen observables missing in one correlated direction, not on any single σ -count, because coherent structure is what cable faults and foregrounds do not produce. Second, *the referee must be independent*: the decisive test is assigned to an instrument (the siren) that shares no systematic with either discordant route — the exact discipline whose absence prolonged the historical episodes. Third, *pre-registration cuts both ways*: the same printed rules that would certify a confirmation make a refutation unarguable. The claim to precedent is therefore conditional, and the paper states the condition: it holds in full only if the siren lands on the derived rate and the residual is traced in the calibration chain. Until then the honest description is: the logical form of Mercury and the W/Z , with the verdict pending and the failure modes printed.

79 The problems resolved: PS-Hubble-1 through PS-Hubble-7

Significance is a claim about kind; accomplishment is a claim about score. Cast as the cohort’s problems-solved registry, here is what the reading has done to the standing puzzles of its subject — each entry with its verdict and its conditionality, because a score without its grading is advertising.

PS-Hubble-1: the absolute scale of H_0 . Why the expansion rate has the value it has was never even a question inside Λ CDM — the model measures h , it cannot ask about it. Within the framework the question closes twice independently: the geometry’s vacuum arithmetic returns $498.3\sqrt{\pi/175} = 66.8 \pm 0.2$ with no astronomical input, and the measured densities divided by the derived fractions return 66.6 ± 0.2 ; concordance 0.6σ (Part III). The physical picture is worth one sentence: the rate is not a free initial condition but the present-epoch reading of an inventory the geometry fixed — asking “why 66.8?” becomes asking “why π ?”, which is a question about the fibration, not about cosmic history. *Verdict: resolved, graded D* — conditional on the inventory conjecture inherited from GG-A10-Cosmos and on the codebook selection of the vacuum coefficient, the latter pre-registered as F-Hubble-11 with a closed two-sided window.*

PS-Hubble-2: the early/late split. The tension’s standard framing — two measurements of one constant disagree — presupposes there is one constant. There is not: there is one derived function $H(t)$, with one present value, read through two dictionaries on either side of the recombination switch at $z_* = e^7 \approx 1097$ — and the two readings *should* agree. The split is therefore not a contradiction to be patched but a diagnosis: the sound-horizon dictionary reads the derived rate to 0.9%, every ladder-free late probe is predicted to read it, and the calibrated ladder is the one instrument displaced — by a systematic the framework predicts and does not derive. (An earlier version of this entry said the two readings should differ by exactly a corridor the geometry prints, and called the split a projection effect Part III computes; the parent’s one-background theorem withdraws that, and so does this paper.) *Verdict: reframed — the early/late split is a statement about one instrument’s calibration; the size of that systematic is not computed here, and its confirmation is the ladder-free probes’ business.*

PS-Hubble-3: why the CMB gives ≈ 67 . In the fitted picture it is a brute output; here the acoustic sky is fit by the exact π -fractions with zero background freedom at $\chi^2 = 187$, and the

residual displacement is one *coherent* -0.9% π -offset across seventeen observables — a fingerprint, not a discrepancy. The CMB reading is therefore a derivation, and its agreement with the derived background is the paper’s deepest single piece of evidence. *Verdict: resolved, D* on the same inventory conditions; the zero-input framing (no background parameter tuned to the sky) is the load-bearing clause.*

PS-Hubble-4: the local excess. The one puzzle the paper refuses to claim. What the framework says about it is a prediction of location, not of size: on one background the excess must live in the calibrated rungs. The one environmental piece it could have credited — a local-void lift, sign forced by the openness theorem, response computed — is priced and not carried, because the depth it needs is not observed. The undischarged part is all of it: the $+10\%$ to H0DN’s 73.50 ± 0.81 , named as a network-level calibration systematic, priced against the ladder’s own error budget through three external channels (dust 2–4%, distance priors 0–3%, calibrator subsample 1–2%) that do not yet reach it, and carried *open* under dated kill tests F-Hubble-8/9/10. *Verdict: unresolved by design — predicted as a ladder systematic, not derived; the declaration is enforced by the registry, not by promises.*

PS-Hubble-5: new physics or systematic? The decade’s sterile dichotomy, reframed into a decidable question. The framework first cleans its own house: its one first-principles local channel ($U(1)_B$) is closed by the boundary-condition termination of GG-A3-B16 and, independently, excluded by eight orders of magnitude; cosmological averaging is bounded below 10^{-3} ; and buying the excess with inhomogeneity would cost a $\langle\delta\rangle \approx -0.56$ fluctuation on flow scales — $> 12\sigma$ in standard structure formation (§35). What remains is a binary with an assigned referee: a bias story with named channels predicts the ladder-free siren lands at 66.8; a real-73 story predicts 73. *Verdict: reframed — from a debate into a measurement, with the date on the calendar.*

PS-Hubble-6: the cosmic coincidence. Why should matter and vacuum balance *now*? On the derived history the order-unity window occupies $\sim 2\%$ of the logarithmic lifetime of the universe, and the observer’s epoch is tied to the same stellar clocks the inventory times — a factor-of-a-few timing question, not a fine-tuning; the full treatment is the parent’s (GG-A10-Cosmos). *Verdict: resolved at the level this paper needs; inherited in full from the parent recast.*

PS-Hubble-7: the consensus paradox. New to this registry, and forced by the 2026 landscape: how can a thirty-eight-expert, covariance-weighted 7σ consensus at 73.50 ± 0.81 and a growing literature of corrections landing at 70–71.5 both be excellent science? The paradox dissolves at the boundary of certification: a combination certifies the agreement of what its members *share* — and is silent on what its non-members dispute; the corrections live precisely in the shared layers (one dust law, one prior family, one calibrator convention) that internal cross-checks cannot see, and the excluded program (CCHP, 70.39 ± 1.94) marks where the membership boundary runs (§43.2). Nothing dishonest anywhere — just two instruments certifying different things, read as if they certified the same thing. *Verdict: dissolved by analysis; the dissolution is testable, because F-Hubble-10’s complete sample erases the boundary.*

The tally. Seven puzzles: three resolved (one twice over), one dissolved, two reframed into measurements, one — the systematic — deliberately left standing as a prediction of location, with its anatomy priced and its tests dated. The asymmetry between what is claimed and what is declared open is the registry’s point: a framework confident enough to derive the rate can afford to leave the one thing it cannot derive visibly unclaimed.

80 The structural innovations: I-Hubble-1 through I-Hubble-10

The register of what did not exist before, each entry in before-and-after form. None introduces a new continuous parameter; several introduce new *practices*, which the cohort regards as innovations of equal rank.

I-Hubble-1 (the derived rate). *Before:* H_0 a free parameter in every cosmology since Hubble — measured, debated, never derived; the tension a fight between two fits. *After:* $H_0^{(0)} = 66.8 \pm 0.2$ as the shadow of a π -locked inventory, arrived at twice (geometry alone; densities over fractions), concordant at 0.6σ , printed before the precision that now tests it.

I-Hubble-2 (the function, not the number). *Before:* “the Hubble constant,” one number per analysis, its redshift dependence a systematic to be marginalized. *After:* the Hubble *function* $H(t)$ as the object — landmarks, $H(z)$ rows, drift signs, de Sitter floor — with H_0 demoted to one reading and every other point on the curve a falsifier (Table 18).

I-Hubble-3 (the asymmetric reading). *Before:* the tension treated as a symmetric clash — maybe the early side, maybe the late side, presumably something in between. *After:* a structurally asymmetric statement: the CMB sits *on* the derived background (with a coherent, computed offset), only the ladder is displaced, and the displacement is read whole as a predicted calibration systematic. Symmetry restored only at the referee: the siren can kill either reading.

I-Hubble-4 (backreaction computed at the exact target). *Before:* inhomogeneity invoked qualitatively whenever convenient — “backreaction could be large” — with the required amplitude never priced. *After:* the escape computed at the exact excess: closing +10% needs $\langle\delta\rangle \approx -0.56$ on flow scales — $3.4/12/22\sigma$ at $40/100/150$ Mpc/ h — excluded in standard structure formation and invocable only through the fully nonlinear timescape loophole, which is priced as such (App. H.4).

I-Hubble-5 (the residual with its channel excluded). *Before:* unexplained excesses quietly absorbed into fitted nuisances, or attributed to “unknown systematics” with no further structure. *After:* the systematic named and located, its one first-principles channel ($U(1)_B$) closed by the termination of GG-A3-B16 and excluded by eight orders of magnitude, its void channel priced and not carried, its averaging channel bounded, and its candidate anatomy assembled from the field’s own literature — a declared debt with an itemized statement, not a hidden one.

I-Hubble-6 (the community invitation with dates). *Before:* alternative frameworks exhort the field to “reconsider Λ CDM” rhetorically. *After:* an explicit invitation to re-derive the

background (§40) operationalized as a dated program: a zero-parameter joint likelihood anyone can run, released scripts (Table 27), and a calendar on which each community can find its own experiment.

I-Hubble-7 (the audited landscape). *Before:* Hubble-tension reviews as narrative compilations — valuable, selective, unreproducible. *After:* the field’s state as a measuring instrument: forty-five determinations, each read from its primary record on a stated date, asserted value-by-value in a released script, charted and tabulated with identifiers — and the three-band structure it exposed elevated to a diagnostic every rival must answer (Part V). To our knowledge the first fully reproducible snapshot of the H_0 literature published *inside* a paper it tests.

I-Hubble-8 (the Asserted-external grade). *Before:* two bad options for contested external corrections: adopt them (and launder), or ignore them (and blind oneself). *After:* a third evidence class — quoted with the proposer’s caveats and the published rebuttals attached, priced in scenario brackets, never averaged into the framework’s own numbers, and guarded by symmetric kill tests. The laundering objection (O-Hubble-54) is answered by the grade’s mechanics, not by protest.

I-Hubble-9 (the decision calendar with a hostage ledger). *Before:* falsifiability asserted in principle; tests undated; rules renegotiable after the data. *After:* the registry drawn on a forward axis with dates, and — the new move — a commitments subsection collecting, appointment by appointment, the printed numbers that cannot move afterwards (§58.2). Pre-registration upgraded from a posture to a ledger.

I-Hubble-10 (bracket pricing). *Before:* systematic budgets quoted as single combined numbers, with the combination rule chosen tacitly — and favorably. *After:* both combination rules printed side by side (quadrature 4.3%/3.8%; stack 9.3%/8.3%) with the target between them, the correlation structure stated, and the conclusion graded by the bracket rather than by a point estimate. The scenario spread carried as the honesty, not hidden as the weakness (§34.5).

81 The honest boundary: proved, conditional, open

This is the concentrated honest frontier of the Part, stated plainly, because a paper that has mapped exactly where its proven territory ends loses nothing by naming the boundary — and forfeits everything by blurring it.

Proved, within the present scope. The existence of one background solution on the derived inventory and the switch redshift $z_* = e^7$ as an integer-ledger statement (GG-A10-Cosmos, Theorem T-Cosmos-16); the landmark chronology and the derived $H(z)$ table as arithmetic consequences of the inventory; the backreaction bound at the exact excess ($-0.56, > 12\sigma$); the $U(1)_B$ exclusion by eight orders, and its closure under the termination; the systematic’s two-scenario arithmetic and its significance-deflation chain, every number asserted in a released script (twenty-three checks, `c2_residual_arithmetic.py`); and the concordance of the two derivation routes at 0.6σ .

Conditional, on named conditions. The π -locked fractions themselves are D^* — derived conditional on the closed-packet conjecture of the parent GG-A10-Cosmos, and this paper must not be read as upgrading them. The vacuum coefficient $\zeta_\Lambda = \pi/175$ is codebook-selected, conditional on the selection lemma, pre-registered at F-Hubble-11 with window $[0.01784, 0.01806]$. The void channel is priced and not carried: its response law is derived, its depth is a survey input the surveys have not supplied, and the weight-function measurement of OQ-Hubble-3 can only confirm the withdrawal or re-open it. The inherited early-sector values ($\Delta N_{\text{eff}} \leq 0.03$, a derived null; $n_s = 0.9634$; $r = 0.0038$) carry their parent-paper grades unchanged.

Open, by declaration. The $+10\%$ systematic — the paper’s one standing debt, predicted in location, priced, and not paid in size; the zero-parameter global likelihood (OQ-Hubble-4), named as work not yet done; the shell-level $H_0(z_{\min})$ profile as a formal fitted-form prediction (OQ-Hubble-6); and the siren-depth refinement (OQ-Hubble-7). Every open item reappears in §86 with its closing condition, its actor, and its date.

Dependencies. The chain is not free-standing. It rests on the OGN-5 integer ledger, the BI-6 anomaly closure, the GDOS/OGGEP boundary semantics, the openness theorem of GG-A4-Euclid, and the one-background theorem of GG-A10-Cosmos (T-Cosmos-16); a decisive failure of any upstream pillar propagates here. The paper has deliberately *added* exposure of its own — F-Hubble-11’s window, the landscape’s band-2 wager, the calendar’s hostages — so that it can also fail alone, without taking the cohort down with it: local falsifiers first, structural ones behind them.

82 The tier ladder: what this paper computes, and what it does not

The cohort’s can/cannot discipline (GG-A5-GDOS, T9.9; applied to particles in GG-A9-SM and to cosmology in GG-A10-Cosmos) has a sharp Hubble instance, and stating it protects the paper from the two opposite misreadings — that it claims all of observational cosmology, or that it claims nothing checkable.

Tier 1 — computed. The background: $H_0^{(0)} = 66.8 \pm 0.2$, the full $H(z)$ curve and its landmark chronology, $r_s = 144.7$ Mpc, $z_* = e^7$, $t_0 = 13.9$ Gyr, $H_\infty = 55.2$, the drift values of Table 18, and $S_8 = 0.825$ on standard growth. These are outputs; each is a place to die.

Tier 2 — priced, not carried. The environmental layer: the void response law and the shell-level profile a lift would imprint on the locally read rate. Geometry fixes the form; the surveys have not supplied the depth, and the framework carries the amplitude at zero.

Tier 3 — priced, not computed. The measurement chain: the ladder’s dust, prior, subsample, and zero-point channels. This paper re-reduces no one’s photometry and re-runs no one’s pipeline;

it prices the field’s own published corrections in scenario brackets and assigns their kill tests. The distinction between pricing and computing is what the Asserted-external grade enforces.

Tier 4 — handed to the instruments. The precision likelihoods: the C_ℓ spectrum belongs to the Einstein–Boltzmann machinery (which, fed the exact fractions, returns the TT sky at $\chi^2 = 187$; §26); nonlinear growth and the JWST regime belong to simulation and to the parent GG-A10-Cosmos; and the residual’s final resolution belongs to the ladder community’s own instruments, under the registry’s dated tests.

This paper computes the expansion history and prices the measurement chain; it re-reduces no photometry and fits no likelihood. The ladder, Boltzmann, and survey pipelines remain the instruments it frames — and has staked itself against.

(52)

The tier ladder is the boundary the framework draws for itself, not a concession extracted under pressure: Tiers 1 and 2 are where a zero-parameter geometry can speak; Tiers 3 and 4 are where the irreducible instrumental and dynamical work lives. The two are complementary — the geometry supplies what the instruments assume, and the instruments supply the precision the geometry cannot.

83 The cost of being right: zero parameters, maximal exposure

A fitted cosmology meets a 7σ surprise by growing: a new sector, a new parameter, a re-weighted likelihood, and the surprise is absorbed. This construction cannot grow. Its inputs sit on four declared levels — the constitutional layer of GG-A1-Constitution; the geometric layer (the fibration, BI-6, the OGN-5 integers); the boundary baseline (four-dimensional GR and the FRW observable language); and the observational calibration layer, which for this paper is remarkably thin: the two CMB densities *as confirmations* (Route I needs neither) and the timing of the present epoch. What it refuses is the fifth level every rival keeps in reserve: the rescue knob. If the siren lands at 73, there is no parameter whose adjustment saves the reading; if the inventory drifts from $1/\pi$, no re-fit recovers it; if the Hosotani computation returns a coefficient outside $[0.01784, 0.01806]$, Route I demotes itself by its own printed rule.

The asymmetry with the rivals should be stated without embarrassment. A framework with adjustable sectors can never lose quickly — every adverse datum becomes a constraint on its parameters rather than a verdict on its existence — and after a decade the tension literature is a museum of readings that were never wrong, only updated. A zero-parameter reading inverts the economics: it is easy to kill, so its survival carries information; it cannot be updated, so its agreements cannot be manufactured. The residual is the test case of the discipline: the single place the framework runs out of derivation, it prints a debt rather than installing a mechanism — because one unforced mechanism would convert the whole paper from a derivation with a gap into a fit with extra steps. The gap, standing open under dated tests, is the price of everything else; the paper pays it in public.

84 The cumulative evidence base

Table 19 contains twenty-one quantitative benchmark rows. Under the conditional statuses of the companion codebook, most are consistent at the $\lesssim 1.5\sigma$ level; the baryon-sector entries carry a flagged few-percent residual in the companion audit, the early-versus-late S_8 gap is carried shared rather than resolved, and one entry (the Kaxion window) is pending first measurement — with the $\approx 10\%$ ladder systematic a predicted, not derived, quantity rather than a prediction of its size. No benchmark is contradicted at high significance, but several are conditional rather than closed.

One contingency deserves its own sentence in a section on cumulative evidence: if the landscape’s middle band continues to fill — if the corrected ladders, the sibling samples, and the dark sirens keep converging near 70 — the cumulative case does not collapse; it bifurcates, exactly as priced in §70. The derived rate is strained in that world ($1.4\text{--}2.4\sigma$) and says so here, in advance; there is no second reading of this paper’s for that world to vindicate instead. An evidence base that grades its own contingencies before the data arrive is the only kind that survives them.

This is the standard by which theories of fundamental physics are evaluated. Consider the Standard Model, which is consistent with the measured 28 parameters but does not derive them; consider string theory, which has produced few numerical predictions within experimental reach. GG-Theory carries 20+ numerical benchmarks, most at the percent level, and has now added the H0DN comparison to the list. The record is coherent, with its conditional entries and open problems stated alongside.

85 Is the GG-Theory reading necessary?

This report has argued that the GG-Theory reading of the H0DN result is *natural*, *independently formulated*, and *structurally distinctive*. We have not argued, and do not argue, that it is the *only* consistent reading of the data. Phenomenological EDE, interacting dark sectors, and various other proposals remain alive. What the H0DN result does is to raise the bar for all of them.

A theory of cosmology now must:

1. Account for the measured local band $[72.7, 74.3] \text{ km s}^{-1} \text{ Mpc}^{-1}$ — produce it, or locate, falsifiably, the bias that creates it;
2. Do so without worsening the CMB fit beyond viable;
3. Without worsening S_8 beyond current lensing measurements;
4. Without requiring more free parameters than are motivated outside cosmology;
5. Without breaking BBN, BAO, or solar-system tests;
6. And ideally with a single falsifier that retires the whole framework.

GG-Theory meets these criteria with one declared exception: it does not produce the local excess at all — it predicts that the $\approx 10\%$ lives in the ladder’s calibration and stakes the verdict on the ladder-free siren — so criterion 1 is met by location, not by production. Among the frameworks compared it remains the only one whose early-stage inputs are fixed outside cosmology and whose

failure modes are stated as single observations. It is therefore a candidate explanation whose remaining freedom is explicitly localized and declared.

86 Open questions, stated as a program

An open question is only as valuable as its closing condition. Each entry below therefore states what is open, what specific result closes it, who can produce that result, on what timescale, and which registry entry enforces it. Table 22 is the index; the paragraphs carry the substance.

Table 22: The open-question program of the derived Hubble reading. Each question carries its closing condition, its actor, its era, and the registry entry (falsifier or open-problem tag) that enforces it. The columns after the closing condition are what make this a program rather than a wish. A question with no actor and no era is not open; it is decorative. The two entries carrying no date are theory questions, and theory has no calendar; the rest are timed against instruments that already exist or are already funded.

OQ	Closes when	Actor / instrument	Era	Enforced by
OQ-Hubble-1	Hosotani/Wilson-line evaluation of ζ_Λ on S_θ^1 returns a value, either way	theory (GG-A5-GDOS O6)	undated	F-Hubble-11
OQ-Hubble-2	The three anatomy channels are individually settled	ladder community; JWST; simulation	2026–2028	F-Hubble-8/9/10
OQ-Hubble-3	Any void lift is measured on the ladder’s own weight function	Euclid/DESI counts + PV surveys	2027–2029	F-Hubble-4 (adjacent)
OQ-Hubble-4	A zero-background-parameter joint likelihood is run against six-parameter Λ CDM and EDE	any group; released inputs	open now	Priority 1
OQ-Hubble-5	(closed: GG-A10-Cosmos computed the amplitude on one background, $S_8 = 0.825$; there is no growth response to have a sign)	—	—	closed
OQ-Hubble-6	The shell-level $H_0(z_{\min})$ profile is published as a fitted form and confronted with binned data	this program + DESI DR3	2027–2029	(new; see text)
OQ-Hubble-7	Siren H_0 is reported in host-distance shells	LVK/XG catalogs	O5 and beyond	refines F-Hubble-2
OQ-Hubble-8	Any curve-reading channel reaches percent class	chronometers; ELT drift; calibrated BAO	2030s–2040s	Table 18
OQ-Hubble-9	Projection sourcing: why sub-lock vacuum loops do not re-source ρ_Λ	GDOS/OGGEP layer (GG-A5-GDOS, GG-A6-GRIP)	open (no date)	names the see-saw’s silent premise
OQ-Hubble-10	Epoch-offset relation: is the fitted-vs-exact Ω_m offset the 75-Myr epoch offset?	this program; sub-percent inventory era	2028–2030	explains one entry at most (§25)

OQ-Hubble-1: the vacuum coefficient. The single most consequential theoretical computation available to the cohort: the Wilson-line (Hosotani) evaluation of the vacuum residue on the compact angle (GG-A5-GDOS, Open Problem O6). A result inside $[0.01784, 0.01806]$ collapses Route I from a codebook-selected point to a computed one and closes the last conditional link in the derived rate; a result outside demotes Route I to the admitted stratum cluster, by the paper’s own printed rule. Either outcome is progress; only silence is not. No observation can substitute for it, and no date binds it — which is exactly why it is listed first, as the program’s standing theoretical debt.

OQ-Hubble-2: the anatomy verdicts. The residual’s three external channels are each one published result away from resolution: an injection–recovery campaign (software only, publishable in a season), the complete ~ 37 -host JWST sample under an agreed convention, and the NIR/sibling ladder. The paper’s exposure is already collected at §58.2; the open question is not whether these tests will run but whether the community will run them *as tests* — with the convention agreed before the reduction, so that the verdicts arrive unarguable.

OQ-Hubble-3: the void on the ladder’s weight function. The response law is derived, but no survey has yet measured the under-density weighted exactly as the ladder’s calibrator and flow samples weight it. Euclid/Roman-depth counts fix $\delta(r)$ outright; the peculiar-velocity programs read the outflow; together they turn the question of whether any share of the ladder’s excess is environmental into a number. A measured lift near zero confirms the parent’s withdrawal (F-Hubble-4’s territory, as re-registered); a measured lift at 2–3% would re-open the question this paper and its parent have closed.

OQ-Hubble-4: the zero-parameter joint likelihood. The over-determined CMB test of §26 promoted to full-likelihood level: background fixed at the π -locked inventory, only the seed and τ inherited, run head-to-head against six-parameter Λ CDM and EDE across CMB+BAO+SN+lensing with covariances kept. Every input is released; any group can run it; the paper’s Priority 1. If the derived chart’s χ^2 per *free* parameter beats the fitted alternatives, the case stops being philosophical — and if it does not, that too is a result the framework has nowhere to hide from.

OQ-Hubble-5: closed. An earlier version carried here the sign of an S_8 growth response, open pending the parent paper’s growth audit. The parent has since computed the amplitude on the one background from the declared A_s — $S_8 = 0.825 \pm 0.010$, sharing the tension with Λ CDM — and on one background there is no growth response to have a sign. The entry is kept so the numbering and the cohort’s cross-references hold; nothing remains open in it.

OQ-Hubble-6: the shell profile as a fitted form. The lift predicts not one number but a *profile*: the locally inferred rate must decay toward the background as low-redshift shells are cut away, at the shape Figure 17 draws. The open task is to publish that profile as an explicit fitted form — amplitude tied to the measured depth, shape to the derived history — so that DESI’s binned $H_0(z)$ analyses (which already report a descending trend at 6.4σ ; §33.1) can be confronted

with a curve rather than a qualitative expectation. This is the cheapest genuinely new confrontation available to the reading, and it is work this program owes the field.

OQ-Hubble-7: the siren-depth refinement. A refinement question the pre-registered rule deliberately brackets: F-Hubble-2 is a *background-rate* rule, and the dark-siren catalogs that will execute it sample far beyond the Hole, so they read the background. But the rare nearby bright events — GW170817’s host sat at forty megaparsecs, well inside — share the neighborhood, and if a lift were real they would read it, at an amplitude bounded by the depth the counts once credited ($\leq 2\text{--}3\%$, far below current per-event precision). The open question is catalog-level: report siren H_0 in host-distance shells, and the question of an environmental share acquires a third, fully independent reader — one that shares no systematic with either the counts or the flows. No number in this paper changes either way; the refinement sharpens where the confirmation would come from.

OQ-Hubble-8: percent-class curve reading. Every row of Table 18 is an open question with a printed answer: chronometers at percent class, BAO on the derived r_s , and — the purest — the ELT redshift drift at centimeters per second per decade. These closings are the slowest on the list and the most valuable per datum: each one tests the *curve*, not the point, and the derived history has already committed everywhere the curve can be read.

OQ-Hubble-9: projection sourcing (the seesaw’s silent premise). The vacuum-energy seesaw of Part III states what ρ_Λ *is*; it does not state what became of the Standard-Model zero-point contributions below M_{lock} , which naive effective field theory would add to the boundary source at up to 10^{55} times the answer. The framework’s structural reply — that the OGGEF projection registers the geometric attractor’s residue and not the flat-space loop sum — is at present a *property claimed of the GDOS layer, not a theorem proved in it*, and this entry exists so that no reader mistakes the one for the other. A degravitation-type statement, proved at the GG-A5-GDOS/GG-A6-GRIP level, would close it; until then the seesaw is billed as a derived scale with a conditional accounting, and the honesty triage carries it so.

OQ-Hubble-10: the epoch-offset relation. Landmark 7 records an arithmetic coincidence: evolving the exact crossing fraction $1/\pi$ forward 75 Myr lands on 0.315 — numerically the Planck-fitted value. If the fitted-versus-exact Ω_m displacement is, in whole or part, the epoch offset, one entry of the coherent π -offset pattern has a mundane explanation — but only one: the rate drifts just 0.24% over the same interval, an order short of the -0.9% h offset. A sub-percent inventory determination, read jointly with the epoch calibration, settles what share of the pattern the calendar owns.

87 Predictions and future directions

The paper’s forward predictions are collected, with their dates and rules, in three standing structures — the falsifier registry (Table 17), the decision calendar with its commitments ledger (§58), and the benchmark sheet (Table 19) — and they are not repeated here. What a Discussion

owes instead is the *assignments*: who can act, on what, and why it is in their interest regardless of how this paper fares. Five communities, five requests, each two-sided by construction.

To the distance-ladder community. Three requests, all of which strengthen the ladder whatever they find. Publish the injection–recovery campaign: it is the cheapest decisive test in the field, it settles the priors channel either way, and whoever runs it owns the verdict (F-Hubble-9). Complete the ~ 37 -host JWST sample under a tip/JAGB convention agreed *before* the reduction — the data will exist regardless; the agreement is what makes them a test (F-Hubble-10). And press the NIR ladder and the ZTF/Rubin sibling samples to a published H_0 : dust physics is the largest single allowance in anyone’s budget, and the near-infrared is where it goes to die (F-Hubble-8).

To the CMB community. The full SPT-3G survey and the Simons Observatory will tighten H_0^{ref} above the derived 66.8 ± 0.2 ; this paper fails at its own anchor if the coherent π -offset decoheres, and asks for exactly that scrutiny. The EDE lobe search (F-Hubble-3) and the ΔN_{eff} convergence are two-sided here: detections kill this reading; tightening nulls starve its rivals. And the zero-parameter joint likelihood (OQ-Hubble-4) is an afternoon’s configuration change for any Boltzmann-pipeline group — the inputs are released, and either outcome is publishable.

To the gravitational-wave community. The decisive verdict is yours, and the paper’s request is procedural, not scientific: report H_0 in host-distance shells as the catalogs deepen (OQ-Hubble-7), keep the counterpart-readiness infrastructure funded through O5 — the bottleneck is follow-up, not sensitivity (§66) — and continue publishing the dark-siren values with their mass-model and completeness systematics displayed, because the walk from 76.6 to 69.9 with shrinking error bars is, to date, the cleanest live demonstration in cosmology of a systematic budget averaging down in public.

To survey cosmology. Whether any share of the ladder’s excess is environmental is the one question in this reading that ordinary galaxy surveys can close by themselves (§33.1): Euclid/Roman-depth counts fix the depth profile, the PV programs read the outflow, and the binned $H_0(z_{\text{min}})$ analyses confront the flat shell profile this paper predicts (OQ-Hubble-6). None of these requires believing anything in this paper; each is a measurement the field wants anyway, and the reading has pre-committed to what they must find.

To theorists, inside and outside the cohort. The Hosotani computation (OQ-Hubble-1) is the standing debt, and it is stated as a falsifier, not a hope: the target is printed, the window is closed, and a computation that misses retires Route I’s point reading by the paper’s own rule. Behind it queue the growth audit (with GG-A10-Cosmos), the shell-profile fitted form (OQ-Hubble-6), and the global likelihood (OQ-Hubble-4). The division of labor with the companion papers is single-source-of-truth throughout: the inventory, the one-background theorem, and the why-now treatment live in GG-A10-Cosmos and are carried here at GG-A10-Cosmos’s values; the Kaxion window [1.59, 2.13] GHz and the count doctrine — one line, alone — live in GG-C1-Kaxion (F-Hubble-7 imports, never re-derives); the codebook grades and the O6 problem live in GG-A5-GDOS; the derived radiation

null $\Delta N_{\text{eff}} \leq 0.03$ lives in GG-A10-Cosmos; the integer ledger in GG-A4-Euclid and GG-A9-SM. One value, one home, every quotation synchronized — so a correction anywhere propagates everywhere, and no two papers can drift apart unnoticed.

The staged road, in one paragraph. December 2026: Gaia DR4 reads the ladder’s floor. Through 2028: the three channel verdicts arrive, or fail to be sought — itself informative. 2027–2029: the Hole census and the shell profile. 2028–2030: the early-side precision era and the O5 siren accumulation, with the dark-siren error bars passing through the gap they must adjudicate. The 2030s: the tensor verdict. Mid-2030s: the siren convergence, and with it the verdict on the central claim. Beyond: the drift era, when the curve itself is watched in real time. At every stage the paper’s number is already on the table; at no stage does it get to put down a new one. That is the whole design.

88 The nobility clause

Because the reading stakes a great deal on a single number, it binds itself to four commitments a reader may hold it to — each one checkable against the text, none of them rhetorical. *Pre-commitment*: every value was fixed before the H0DN measurement it is compared against; the early inputs have not been adjusted since the parent recast; and the additions of 2026 — the landscape audit, the anatomy pricing, the calendar — added exposure, never wiggle room. *Honest grading*: each claim carries its codebook grade (theorem, D*, Asserted-external, named-open), quoted identically wherever it appears, with the least flattering entries — the residual, the conditional fractions, the shared S_8 tension, the failing tilt — displayed in the abstract’s own register. *No retreat*: where the framework cannot derive a quantity, it says so and leaves the gap open rather than fill it with an unforced mechanism; where the empirical record embarrassed a forecast (one bright siren in nine years), the record was printed and the dates re-based, while the thresholds — hostage to nothing — stood. *Symmetric stakes*: every kill rule has its confirmation twin printed at equal prominence; the same siren that would confirm the reading would, at 73, destroy it; and the paper has arranged, in advance and in public, to be unable to escape that verdict. These commitments cost rhetorical safety and buy the only currency a decade-old controversy still respects: a participant who cannot move.

89 The bottom line

One sentence per movement, and then the verdict. The H0DN comparison supports the derived reading strongly where it is strong and is silent exactly where the paper says it is silent. The reading’s significance is not a better fit but a different kind of object: a rate derived twice, embedded in a derived history, exposed on a printed calendar. Its accomplishments are seven standing puzzles resolved, dissolved, or reframed, and ten structural moves the tension literature did not previously contain. Its honest boundary is drawn in tiers: it computes the history, prices the chain, and hands the likelihoods to the instruments — paying for that clarity with maximal exposure and one declared

debt. And its road ahead is a numbered program in which every open question carries its closing condition and almost every closing condition carries a date.

The verdict the Part offers is deliberately modest and deliberately sharp. Modest: nothing here claims the tension is solved — the residual stands open, the conditions are named, and the middle world of §70 is priced as the uncomfortable possibility it is. Sharp: the reading is now the only participant in the controversy that has specified, in advance, every observation that would end it — and the first appointments are months, not decades, away. Part IX closes the paper in a different register, asking what it would mean if the reading is right.

90 Part VIII in plain English: what the result is worth, what it is not, and where the debts are

This is the Part where the paper argues with itself.

What the comparison supports strongly: a background rate derived twice by disjoint routes, landing where the early-universe network already sits, with no continuous parameter anywhere in the chain. What it supports moderately: the local lift, whose response law is derived and whose depth is somebody else's disputed measurement. What it does not support at all, and says so in as many words: the seven-percent residual, which is named, priced, and left open.

The Part also does the comparisons a reader is entitled to. Ten points of significance set beside the leading alternative explanations. Seven standing problems that the reading resolves, dissolves, or reframes — and the three verbs are not interchangeable. Ten structural moves the tension literature did not previously contain. And the cautionary historical precedents alongside the flattering ones, because a framework that quotes only the cases where the outsider turned out to be right is not being honest about the base rate.

The boundary is drawn in tiers. This paper computes a history, prices a chain, and hands the likelihoods to the instruments. It does not run the likelihoods, and it should not be read as having done so. It pays for that clarity with maximal exposure — zero free parameters means nothing can be adjusted when a measurement moves — and it carries one declared debt, a coefficient whose first-principles computation has not been done.

The verdict offered is modest and sharp at the same time. Modest, because nothing here claims the tension is solved. Sharp, because this is now the only reading in the controversy that has published, in advance, every observation that would end it, and the first of those observations is months away rather than decades.

Part IX

Conclusions

The retirement of a measured constant — one derived history, a dozen ways the universe can answer, and what it means that we can ask

91 The summit thesis

For a century, the question “how fast is the universe expanding?” was answered by measuring. Better telescopes, better calibrations, better statistics — a heroic tradition, running from Hubble’s first straight line through thirty-eight specialists and a covariance-weighted consensus good to one percent. This paper has answered the same question a different way: by deriving. The expansion history was computed from a geometry with no dial to turn, and the measured values were then confronted with it — the early one landing on the derivation, the local one dissected against it, piece by honest piece.

The Hubble constant was never a constant, and — if the reading defended here is right — it was never merely a measurement either. It is one point on a derived curve: a curve that begins with a predicted inflation, runs on a predicted matter sector, and ends on a predicted vacuum. The century of measuring the universe is ending. The decade of answering it has a schedule.

That is the summit thesis, and the tension itself is its best witness. Measure one number two independent ways, to one percent each, and any hidden assumption anywhere in either chain is promoted to a seven-sigma crisis. The split was never a failure of the instruments — both are magnificent — and, on this reading, never a crisis in the laws. It was the measuring tradition running into the ceiling of what measurement alone can settle: precision had outrun understanding, and the surplus precision had nowhere to go but paradox. The resolution offered in these pages is not a better measurement. It is a change in what kind of thing the number is.

92 What this paper has established

The whole argument compresses into ten numbers, and they are worth walking in plain language before the table states them exactly — because they are not ten separate results. They are three acts of one story.

The curve (rows 1–8): the background rate today, 66.8 ± 0.2 , derived twice — once from the geometry’s own vacuum arithmetic with no astronomical input at all, once from the measured cosmic densities divided by the derived fractions — with the two routes agreeing at a fraction of

their stated errors; behind it the full inventory (matter $1/\pi$, dark energy $1 - 1/\pi$, atoms $1/2\pi^2$), the exact equation of state $w = -1$, the corridor switch at $z_* = e^7$ (the atomic-chain decoupling 0.4% below in e-folds), the age of the universe at 13.9 billion years, and the quiet floor, $H_\infty = 55.2$, that the expansion will hold forever.

The ladder (row 9): the framework predicts the calibrated distance ladder to read about ten percent high, offers no physical channel for any of it — the one environmental channel it could have credited is computed and found to need a depth the surveys do not show — and hands the verdict to the probes that use no rung. What lies between the derived 66.8 and the ladder’s 73.50 ± 0.81 is the one thing this paper refuses to claim: a named, priced, still-open systematic in the measurement chain, with its candidate anatomy quoted from the ladder literature itself, not yet sufficient, and its kill tests dated.

The referee (row 10): gravitational-wave standard sirens, which owe nothing to any ladder or any afterglow, are predicted to converge near 67. Near 73, this paper is wrong — publicly, quickly, and by its own printed rule.

Table 23: The paper in ten numbers: value, grade, location, and falsifier. Ten is the number of quantities a reader should be able to carry away, and the fourth column is why the list is that short. Every entry names a measurement that would retire it, and two of the ten are retired by the same standard-siren result. The grades are the honesty triage’s: D* marks a value the ledger derives subject to the measure lemma, and the three entries that are not D* say what they are instead.

Quantity	Value	Grade	Falsified by
Background rate $H_0^{(0)}$	66.8 ± 0.2	Derived (cond.)	a converged siren at 73
Matter fraction Ω_m	$1/\pi$	D*	Ω_m settling away from $1/\pi$
Dark energy Ω_Λ	$1 - 1/\pi$	D*	a measured $\Omega_\Lambda \neq 1 - 1/\pi$
Baryons Ω_b	$1/2\pi^2$	D*	$\Omega_b/\Omega_m \neq 1/2\pi$
Equation of state w	-1 exactly	Derived-cond.	a converged $w \neq -1$ at 3σ
Decoupling z_{dec}	1091.5	D*	a measured z_{dec} off the atomic chain by $> 1.4\%$
de Sitter floor H_∞	55.2	D*	follows from Ω_Λ
Present age t_0	13.9 Gyr	D*	inconsistent stellar/cosmic ages
Ladder excess ϵ_{syst}	$\approx 10\%$	Named, not derived	a ladder-free probe at 73.5
Siren prediction H_0^{GW}	$\rightarrow 67$	prediction	a converged siren at 73

Every row carries its honesty grade and its falsifier in the same breath, because that pairing — a claim and the observation that would kill it, printed together — is the discipline the whole twenty-two-paper series is built on. Nothing in the table is negotiable after its datum arrives; most of the data arrive within the decade.

93 The retirement of a measured constant

Step far enough back and this paper is one episode in a very old story: the story of constants that graduated.

The speed of light was, for two centuries, a measured number — Rømer timed it with Jupiter’s moons, Fizeau with a toothed wheel, ever more decimals, ever better apparatus. Then Maxwell

derived it from the constants of electricity and magnetism, and the meaning of the number changed forever; by 1983 the graduation was complete, and the speed of light is now a *definition*, a number no experiment can revise. The Rydberg constant was spectroscopy’s most precisely measured quantity — until Bohr computed it from the electron’s charge and mass, and a decimal-chasing industry became a test of quantum mechanics overnight. The W and Z masses were predictions before they were measurements, and the measuring of them was exciting precisely *because* a theory had already spoken. In each case the same event: a number crossed the line from *input* to *output*, from a fact about the world to a consequence of a law — and the field that owned it stopped doing natural history and started doing physics.

Cosmology has never had that moment. Its deepest numbers — the matter density, the vacuum energy, the expansion rate — have remained inputs for a century: things to be measured, tabulated, and accepted, with Λ CDM as the magnificent bookkeeping that turns six of them into the whole sky. H_0 is the oldest and most storied of them, and the proposal of this paper, stated in the plainest terms, is that it is next in the lineage: *the last measured constant of cosmology, due for retirement* — not from relevance, but from the status of input. On the reading defended here, H_0 is what the derived history happens to read today, exactly as the Rydberg is what Bohr’s atom happens to radiate.

And here is what the retirement buys, because a graduation is never a loss. When a constant becomes a consequence, disputing its last digit stops being the frontier. Nobody today measures the speed of light to adjudicate between theories; the excitement moved to the places where the law that fixed it could fail. The same movement is available to cosmology now, and it is overdue. The dispute over the last digit of H_0 — 73.5 or 73.2, 67.2 or 67.4 — has consumed a decade of the field’s finest instruments and sharpest people, and it cannot, even in principle, end by getting sharper: the tension is already seven times its stated errors, and another decimal on either side adds only another decimal of paradox. What ends it is structure: a derived curve that both endpoints must answer to, and a third instrument that owes neither of them anything. The question worth a career is no longer *what is the number?* It is *is the number a theorem?* — and that question, unlike the other one, has a scheduled answer.

94 One history, three acts, all predicted

What replaces the measured constant is not a cleverer fit. It is a complete, quantitative history of the universe in which every act is predicted in advance — the fuller framework of the parent paper GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos), read here at the one number this paper owns. It is worth telling once in plain words, because the Hubble constant is only the middle of a story whose beginning and end are staked just as sharply.

Act one: a predicted beginning. The universe opens not in an unexplained singularity but in a geometric launch — an inflationary plateau sited by the geometry itself, one e-fold below the grand unification edge. And the launch is not a flexible story: it comes with numbers. The tilt of the primordial spectrum is computed at $n_s = 0.9634$ — and here the paper prints its parent’s one currently failing test: the 2025 CMB-only value sits 1.5σ above it and the DESI-combined value

3.2σ , with no dial in the chain that moves the tilt up. The primordial gravitational-wave amplitude is fixed, two-sidedly, at $r = 0.0038$ — above the floor of the coming polarimeters, below the current bound, with *both* directions of failure fatal. No inflaton was invented for the occasion; the derivation is GG-A10-Cosmos's, and this paper simply inherits the launch its expansion history requires.

Act two: a predicted middle. The matter that clumped into galaxies is, on this reading, one specific particle with one specific address: the Kaxion, at $8.1 \mu\text{eV}$, waiting in a laboratory window — 1.59 to 2.13 GHz, the benchmark at 1.95 GHz — narrow enough to be opened or shut within the decade, as a single line standing alone: the count itself is the fingerprint, since the $1/35$ partner an earlier reading predicted is now the longitudinal Z'_B far above any receiver, and its absence is the prediction (GG-C1-Kaxion; Arisaka, 2026, C1-Kaxion). Dark matter with a mass, a frequency, and a deadline — not a parameter space, an appointment.

Act three: a predicted end. The dark energy driving the acceleration is the residual $1 - 1/\pi$ of a closed inventory, with an equation of state that is $w = -1$ *exactly* — not approximately, not on average, but exactly, because the geometric transition that produced it is complete. The expansion settles onto the de Sitter floor at $H_\infty = 55.2$ and holds it forever. A detected deviation of w from -1 , at the precision the coming surveys will reach, retires this act on the spot.

Three acts, each carrying numbers fixed before the instruments that will test them. The Hubble constant of this paper's title is where act two hands over to act three — one reading of one curve, flanked on both sides by predictions at least as sharp as itself. That is the sense in which the tension was always too small a question: the same geometry that says 66.8 says 0.0038, says $8.1 \mu\text{eV}$, says -1 exactly — and every one of those numbers can die independently.

95 The dozen ways the universe can answer

Here is the field this paper actually invites its readers into — not the defense of a decimal, but a decade in which the universe is asked twelve sharp questions, each with a printed answer it is free to contradict. They are walked here in the order the answers arrive; the formal registry, rules, and dates live in Part VI, and nothing below adds or alters a commitment. What each entry states is why the measurement is *meaningful* — what structure it kills or confirms — rather than merely precise.

One: the astrometric floor (December 2026). Gaia DR4 maps the parallax zero-point that every Milky-Way calibrator inherits. Meaningful because it reads the foundation of the entire ladder without a telescope pointed at a single Cepheid — pure geometry auditing a measurement chain.

Two: the ladder's own three verdicts (2026–2028). An injection–recovery simulation, the complete JWST calibrator sample, the near-infrared ladder. Meaningful because each settles a named few-percent allowance *either way* — the first time the tension's systematics have carried symmetric, dated tests rather than open-ended suspicion.

Three: the map of our neighborhood (2027–2029). Euclid- and DESI-class surveys count galaxies straight through the Local Hole and read its outflow. Meaningful because whether any share of the ladder's excess is environmental becomes a measured number — the framework predicts none — and one prediction of this paper is confirmed or retired by ordinary galaxy counting.

Four: the shell test. The locally inferred rate must decay toward the background, at a computable profile, as the nearby shells are cut away — a shape no constant calibration error can imitate. Meaningful because it distinguishes *where you live* from *how you measure* using the survey data alone; a descending trend is already reported and deepens with each data release.

Five: the early-side null (2028–2030). The coming microwave instruments reach the discovery floor for early dark energy and drive the radiation term toward the cohort’s null. Meaningful because this paper is the one participant that *needs nothing to be found*: tightening nulls confirm it while they starve every early-side fix.

Six: the equation of state (2028–2030). DESI, Euclid, and Roman push w toward the two-percent regime. Meaningful because $w = -1$ *exactly* is a statement about the completeness of a geometric transition — the sharpest possible contrast with every dynamical dark-energy sector, and one that cannot be softened if the data drift.

Seven: the dark-matter appointment (through ~ 2030). The funded haloscope programs sweep the Kaxion window. Meaningful because a single line at 1.95 GHz, with no partner, would not merely add a particle — it would read the count of a compact dimension’s light angles in a microwave receiver; and a swept, silent window retires the dark-matter claim of the entire series, this paper’s inventory included.

Eight: the referee (O5 to the mid-2030s). The standard sirens converge — on 67, on 73, or, most interestingly, between. Meaningful because the siren is the only instrument that reads the clock with no ladder and no sound horizon anywhere in its chain: the one measurement both camps agreed, in advance, to lose to.

Nine: the primordial whisper (the 2030s). LiteBIRD and the Simons Observatory read the tensor ratio at the 10^{-3} floor — realistically toward the decade’s end. Meaningful because $r = 0.0038$ is two-sided: a detection well above, or a null well below, each retires the predicted launch — the rare inflationary claim that can fail in both directions.

Ten: the growth ledger (this decade). The lensing surveys fix S_8 , where this paper’s inherited value is 0.825 ± 0.010 and the tension with cosmic shear is shared, not resolved. Meaningful because growth reads the same inventory a second way: matter that expands one way and clumps another would break the single ledger behind both.

Eleven: the curve, read directly (2030s). Cosmic chronometers approaching percent class, and baryon-acoustic points converted to absolute readings by the derived $r_s = 144.7$ Mpc, confront the predicted $H(z)$ row by row. Meaningful because a fitted model can absorb a surprising point by re-fitting; the derived curve cannot absorb anything, anywhere, ever.

Twelve: the expansion watched live (2040s). The extremely large telescopes attempt the redshift drift: +2.1 centimeters per second per decade at $z = 1$, falling through zero near $z \approx 1.9$. Meaningful because it is cosmology with no proxy at all — no candle, no ruler, no clock but the observer’s own — the expansion history watched, not inferred, at values this paper has already committed to.

Twelve questions; twelve printed answers; not one of them “what is the ninth decimal of H_0 ?” That is what a falsifiable, beyond- Λ CDM precision cosmology looks like — and every entry on the list is funded, scheduled, or already taking data.

96 What it would mean if it is right

The Hubble constant is, in the end, the rate at which everything recedes from everything else — the clock of the whole universe. For a century we treated it as a number to be measured ever more precisely, and the harder we measured, the more it split in two. The reading offered here is that the splitting was telling us we had mis-stated the question. There is no single number, because the expansion was never a constant: it is a function of time, $H(t)$, falling from an unimaginable early rate toward the quiet de Sitter floor it will hold forever, and H_0 is merely where that curve happens to cross the present moment.

If the geometry is right, the consequences run deeper than the tension. The rate of cosmic expansion is not a free fact about the universe but a consequence of what the universe is made of — and the densities that fix it are themselves fixed by π and a handful of integers. The composition of the cosmos joins the list of things that did not have to be measured to be known; the tension joins the list of paradoxes that were really promotions, the moment a field's precision outgrew its questions; and cosmology itself completes the graduation that physics made a century ago — from cataloging the world's numbers to deriving them, from natural history to law. A universe whose expansion rate is a theorem is a universe that can be *understood*, not merely surveyed; and the first generation to know that would be the one reading the instruments this decade.

And if the geometry is right there is a quieter consequence, easy to miss beside the grand one. The magnificent measurements do not become obsolete — they become *evidence*. Every decimal the ladder community fought for, every calibration audit, every variant in the consensus catalog turns into a precision record of how a superb instrument reads a derived universe through a biased chain — data about the measurement chain itself, of permanent value, whatever the verdict. Nothing in a century of measuring is wasted on this reading. It is re-purposed.

97 What it would mean if it is wrong

A conclusions Part that only imagines victory is advertising, so let failure be given its own section, written with the same care.

If the sirens converge at 73, the derived rate is wrong, and with it the reading of the inventory that produced it. That failure would be clean, public, and instructive in a way the tension's decade of stalemate has not been: it would say that the true background rate is high; that the CMB's densities cannot be divided by the geometric fractions to give it; that the coherent agreement of Part III was a coincidence of extraordinary depth — seventeen observables, one direction, percent scale — and coincidences of that depth are themselves data. Every rival would inherit a sharpened problem: the early side would then genuinely need new physics, and the landscape's three bands would need an explanation with the measurement-side reading struck off the list. The audit trail in these pages — which link broke, at which grade, under which condition — would show exactly where a zero-parameter program buckled, which is precisely the information the next such program needs. A wrong theory that dies cleanly pays for itself; the field has spent ten years learning how expensive the other kind is.

If the verdict lands *between* — the middle world of ≈ 70 — the outcome is subtler and is priced in advance: the derived rate is strained but not killed at the pre-registered level, and the open question sharpens into one this paper has already named. That world is uncomfortable for everyone, this reading included; the discomfort is printed, with its numbers, in Part VII.

And if the window simply refuses to close — if the Hosotani computation stays undone, the channel tests unsought, the sirens slow — then the paper’s contribution is the calendar itself: a public record that, at minimum, converts the field’s patience into appointments. Even the failure to be tested, on a printed schedule, is information.

98 The objective limitations, one last time

Honesty requires the caveats be gathered in one place, in the same plain language as the claims. The inventory is graded D^* : derived up to a single named measure lemma and the declared choice of observing epoch — a conditional, not a theorem, and never advertised otherwise. The vacuum coefficient $\zeta_\Lambda = \pi/175$ is selected by the theory’s own naming rules, conditional on one lemma, its closure pre-registered as a falsifier (F-Hubble-11) with a window that cannot be renegotiated. The $\approx 10\%$ local excess is not derived at all: it is named as a distance-ladder calibration systematic the framework predicts, its one geometric channel closed, its void channel priced and not carried, its averaging channel bounded, its candidate anatomy quoted from the ladder’s own literature at the ladder’s own error bars. The middle world — a true rate near 70 — would strain the pure derived value at the level Part VII computes, and the paper carries that possibility as live, not remote. And the whole reading is staked on a single prediction — $H_0^{\text{GW}} \rightarrow 67$ — that has not yet been measured. None of this is hidden; all of it is printed, because a reading that could be killed cleanly is worth more than one that cannot be killed at all.

99 How to check this paper from your desk

A last practical word before the closing sections, because this paper was built to be audited, not believed. Every quantitative figure and forecast regenerates from a short, documented, standalone script, released with the paper and indexed in Appendix M.2: the derived history, the landmark chronology, the residual arithmetic with its twenty-three asserted checks, the forty-five-value landscape with every entry pinned to its primary record, the siren forecast. Each script states its inputs, asserts its own arithmetic before writing anything, and runs in seconds. A skeptical reader with one evening can re-derive the paper’s numbers; a skeptical reader with one afternoon and a Boltzmann pipeline can run the zero-parameter confrontation of Part VI’s first priority. The invitation is genuine and the point is structural: a derivation, unlike a fit, has nothing behind the curtain — and a field-wide dispute conducted over released arithmetic is the kind that actually ends.

100 Why there is someone here to time the universe

One question remains that a paper about the universe’s clock cannot politely decline, though it can only sketch what the cohort’s capstone develops in full. Why is there anyone here to read the clock at all?

Notice, first, a quiet coincidence buried in the arithmetic of this paper. The expansion rate we measure is set by the same inventory that sets how long stars burn — the same fractions, the same integers, one ledger. The astronomers who measure H_0 are made of elements forged in stars whose lifetimes that ledger fixed, standing on a planet warmed by one such star midway through its allotted burning, at the one broad epoch when matter and vacuum share the budget within a factor of a few. We measure the universe’s clock at the moment our own clocks — stellar, biological, civilisational — happen to be running; and on the derived history that is no accident, because they are not different clocks. The why-now coincidence that puzzles Λ CDM is, on this reading, the statement that observers arrive when their kind of chemistry does, and their kind of chemistry is timed by the same geometry that times the expansion (the full account is [GG-A10-Cosmos’s](#)).

The deeper point belongs to the series rather than to this paper, but it deserves its plain-English sentence here. The geometry that fixes the expansion rate does not stop at galaxies: the same descent runs on through the stable atoms it licenses, the stars it schedules, the chemistry it makes inevitable, toward open systems that model their surroundings — and eventually toward systems that build interferometers and argue about decimals ([GG-A11-Origin](#); [Arisaka, 2026, A11-Origin](#)). On that reading, a measurement of H_0 is not an outsider’s act performed on the universe from nowhere. It is the universe, at one late rung of its own staircase, timing itself — and the century-long fight over the number is what it sounds like when the instrument first hears its own tick clearly. We are not the audience of the expansion. We are the part of it that has begun to keep the log.

101 A personal note, and a message to my colleagues

Papers rarely say plainly what their author expects, so let it be said. I do not know whether the geometry is right. What I know is that the background rate it forces was fixed before I compared it to anything, that it landed on the value the microwave sky already showed, and that it makes a prediction — a siren near 67, not 73 — that the coming decade will confirm or destroy. I have not adjusted a single early-universe input to make the numbers meet; where the framework cannot reach, in the local excess, I have said so and left the gap open rather than fill it with a mechanism I could not derive. If the sirens side with the ladder at 73, this reading dies in public, quickly and cleanly, and it will have done the one useful thing a wrong theory can do.

I spent fifty years as an experimental physicist — Kamiokande, kaon and B -meson CP violation, CMS at the LHC, Auger, XENON — measuring numbers the Standard Model declined to explain, and watching the era’s deepest discoveries arrive from the sky. Every number I helped measure in half a century was an input: something the world told us, and theory then accommodated. The number in this paper is the first one of my scientific life that was an *output* — written down by a geometry before any instrument of mine could weigh in — and I confess that the experience of

testing it has been of a different kind entirely. An input can only be refined. An output can be *wrong* — and therefore, for the first time, the measurement matters in the way discovery measurements matter: sharp enough to kill.

That experience is the substance of the message I want to leave with my colleagues in astronomy and cosmology, stated with the respect fifty years inside the measuring tradition should earn it, and with the bluntness the moment deserves. The dispute over the last digit of H_0 is over — not because either side lost, but because both sides won: the measurements are now so good that they have proven the question undecidable from inside. Another decade of sharpening 73.5 against 67.2 will produce another decade of seven-sigma stalemate, because no measurement of a number can settle what kind of number it is. Meanwhile, an alternative worth the field's instruments now exists: a beyond- Λ CDM precision cosmology that is not a deformation of the six parameters but a derivation of them — inflation with computed observables, dark energy with an exact equation of state, dark matter with a frequency and a deadline, an expansion history exposed at every redshift — falsifiable, near-term, and scheduled. The twelve appointments of this Part are not this paper's private wager; they are, I believe, where the actual discoveries of the next decade will be made, by the same instruments and the same people now spending their precision on one contested decimal. The tension was never the prize. It was the signpost — and it points away from itself. This is where the real excitement begins: not at the tenth measurement of an old constant, but at the first tests of whether the constants were ever free at all. I invite my colleagues — with their superb instruments, their hard-won calibrations, and their healthy skepticism fully intact — to spend the coming decade there. The universe has been asked the same question for a hundred years. It is time to ask it better ones.

And one sentence for the youngest readers, because an invitation owes its fine print. The appointments of this Part are mainstream observations — ladders, catalogs, surveys, sirens — that will count professionally whichever way this framework fares; a thesis built on *them* is safe in every world this paper prices. A thesis built on the framework itself is not: it is a proposal under test, unaccepted, and in the middle world it stays unresolved for a decade. Work the appointments; hold the framework to its dates — skepticism fully intact is not only permitted here, it is the job.

102 Part IX in plain English: what changes if a measured constant turns out to have been derivable

For a hundred years the Hubble constant has been a quantity to be measured. The claim of this paper is that it is a quantity to be derived, and that the measuring was always reading one point off a curve the geometry had already fixed.

If that is right, a great deal that has looked like coincidence stops being coincidence. The matter fraction, the baryon fraction, the epoch of equality, the sound horizon, the age of the universe, the calm the expansion eventually settles into — all of them come out of the same small set of integers, and not one of them is available to be adjusted when a new measurement arrives. The tension itself becomes a statement about where we happen to live rather than about what the universe is made of.

If it is wrong, it will go wrong quickly and in public. This paper names a dozen appointments,

most of them in this decade and several already scheduled. A converged standard siren near 73 ends it in one line. There is no version of this framework that survives that measurement, and the paper says so before the measurement rather than afterwards.

The limitations are restated one last time, because a conclusion that quietly drops them is not a conclusion. The measure lemma is conjectured and not proved. The seesaw coefficient is selected by a rulebook and not yet computed. The residual on the local ladder is named and open. Every one of those is written into the grade attached to every number this paper quotes.

What is left is an invitation, and the closing sections make it. Anyone with a laptop can check the central chain from the released scripts in an afternoon. Anyone with a telescope can attack the appointments. The framework asks that skepticism be kept fully intact — that is the job — and asks only that it be aimed at the dated commitments, where it will actually settle something.

Envoi

For a hundred years we asked how fast the universe expands, and measured the answer more and more precisely, until it split in two. Perhaps it was never one number to be measured, but one curve to be derived — a curve that began with a predicted burst, runs today on a predicted budget, and will end on a predicted calm — and the universe has been keeping that single clock since its first microsecond, at a rate we are only now, in this decade of appointments, learning how to read. The clock is old. The reading of it has just begun.

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 OGGE/ GDOS/Gi-4/Qi-4 Pillars), theoretical judgments, physical interpretations, and any remaining errors are the sole responsibility of the author.

Part X

Technical Appendices

A graduate course in the Hubble tension read through *GG-Theory* — reference dictionaries, a conceptual on-ramp, and from-scratch derivation courses that review each topic at textbook level, name what the standard treatment cannot explain, and derive the geometric resolution, with every step shown

How to use these appendices. These appendices are written for the reader we most hope picks this volume up: a graduate student entering cosmology, or a colleague from a neighboring field, who wants to understand the Hubble tension — and this paper’s reading of it — from the ground up, with nothing assumed and nothing skipped. Every topic is taught in the same three steps: first the *textbook*, built from scratch with every derivation shown and a physical intuition beside each result; then the *limitations* — what the standard treatment assumes and cannot explain, stated plainly; then the *geometric resolution*, graded by the honesty key of §A.4. The courses are self-contained and cross-referenced, so you can enter anywhere; if you are new to the subject, start with the primers of §B.1 and read the courses in order. Everything quantitative regenerates from the released scripts (§M.2), and the companion volume for the wider cosmology — inflation, nucleosynthesis, the microwave sky in full — is the appendix course of the parent paper (Arisaka, 2026, A10-Cosmos). Welcome; the coffee is on the left.

Table 24: The course map of these appendices. Each course reviews its topic at textbook level, names the standard treatment’s limits, and derives the geometric reading, with every step shown. The appendices are courses rather than a technical supplement, which is a deliberate choice about who this paper is for. A reader who already knows the material can take the headline outputs from the third column and return to the main text; a reader who does not can start at Course C and arrive at the same place without a second book open beside this one.

Course	Topic	Headline outputs
A–B	dictionaries; primers and questions	the vocabulary, the grades, the on-ramp
C	the Friedmann background	the equations; $t_0 = 13.9$ Gyr; the $H(z)$ rows
D	distances and the ladder	the five distances; the common-mode theorem; 0.208 mag
E	the CMB as an H_0 instrument	$r_s = 144.7$ Mpc; why the sky never measures h
F	vacuum energy and the seesaw	the 10^{120} dissolved; the closed form for H_0
G	the local universe	$\delta H/H = -(f/3)\bar{\delta}$; the response law; the -0.56
H	backreaction and averaging	Buchert’s $\mathcal{Q}_D$; the three walls
I	growth and S_8	the growth equation; $S_8 = 0.825$
J	standard sirens	the chirp mass; GW170817 end-to-end; $N = 22/61$
K	time and the curve-readers	the arcsinh clock; chronometers; the drift; the T clock
L	the statistics of a tension	why precision amplifies bias; the bracket theorem
M–N	reproducibility; registries	the scripts; the ledgers, audits, and roadmaps

A Reference dictionaries

This block is the reference desk: the naming convention and how the program's names nest, a one-page summary of every headline number, the symbol glossary, and the reader's key to the honesty grades used throughout. One acronym from outside the program belongs on this desk before any of ours: H0DN, the H_0 Distance Network — the 2026, thirty-eight-author network combination of the local distance ladder (Part II) — whose consensus 73.50 ± 0.81 is the late-side benchmark of the entire paper.

A.1 The cohort naming convention

Every named object in this report carries a cohort-global identifier of the form **Prefix-Suffix- N** , so that a claim can be traced across the twenty-two papers without ambiguity. The prefixes used here are **T** (theorem), **F** (falsifier / kill test), **O** (objection), and **I** (innovation); the suffix names the sector. Thus **F-Hubble-2** is the second falsifier of the Hubble reading (the standard siren, §56); **O-Hubble- N** tags the objections of the penultimate Part; and inherited results keep their home tags — **T-Cosmos- N** for the theorems imported from **GG-A10-Cosmos**, **T8.7** for the lock-scale theorem of **GG-A5-GDOS**. The convention is the program-wide standard of **GG-A11-Origin**, reproduced here so this paper is self-contained. Figure 23 shows how the names used in this paper nest inside the program's layers.




Figure 23: How this paper's names nest. Reading outward: the tags defined here (innermost) are instances of the cosmological stratum's objects, which are outputs of the one engine, which runs on the constitution and the integer ledger. A cross-cohort citation names the innermost layer that owns the object.

A.2 Comprehensive summary table: every headline number and where it is earned

Table 25 collects every quantitative claim of the report, its value, its honesty grade (§24), and the exact place the reader can check it. It is the one-page index to the paper's arithmetic.

Table 25: Every headline number of the derived Hubble reading, its grade, and where it is derived or defined. Grades as in the honesty triage (Table 5). This is the paper’s ledger and the place to check any number quoted elsewhere in it. Reading down the grade column is the fastest honest summary available: most of the headline values are conditional on a single lemma, one is measured rather than derived, and one is named as open. The last column gives the equation or the section where each is fixed, so nothing here has to be taken on the table’s word.

Quantity	Value	Grade	Where
Background rate $H_0^{(0)}$	66.8 ± 0.2 (two-road 66.6 ± 0.2)	Derived cond. / D*	§24, App. E.7
Baryon anchor	66.5	D*	Eq. (26)
Matter anchor	67.0	D*	Eq. (26)
Matter fraction Ω_m	$1/\pi = 0.3183$	D*	Eq. (16), Hopf ratio
Baryon fraction Ω_b	$1/2\pi^2 = 0.0507$	D*	Eq. (16)
Dark-energy fraction Ω_Λ	$1 - 1/\pi = 0.6817$	D*	Eq. (16)
Seesaw coefficient ζ_Λ	$\pi/175 = 0.01795$	Derived (cond.)	§24.2, Eq. (23), App. F.4
CMB reference	67.4	measured	§30
Ladder (predicted)	66.8 read $\approx 10\%$ high	named, not derived	Eq. (31)
Local-void channel ϵ_{Hole}	0 carried (response computed)	priced, not carried	App. G.5
Ladder systematic ϵ_{syst}	$\approx +10\%$ (+10.0% exact)	named, not derived	§30
Present age t_0	13.9 Gyr (today)	D*	Eq. (33)
Siren prediction H_0^{GW}	≈ 67	prediction	§57
Clustering S_8	0.825 ± 0.010	Derived-cond. (Lemma V')	§39
Corridor-switch z_*	$e^7 \approx 1097$	inherited	§30
Lock scale M_{lock}	3.85 TeV	Theorem	GG-A5-GDOS T8.7

A.3 Abbreviations and symbol glossary

For readers approaching this report without full familiarity with the Arisaka-GG preprint cohort, the following short glossary recapitulates the technical vocabulary used above. Fuller definitions are in the flagship summary paper (Arisaka, 2026, A1-Constitution; Arisaka, 2026, A9-SM; Arisaka, 2026, A5-GDOS; Arisaka, 2026, A10-Cosmos; Arisaka, 2026, A7-Qi4).

GG-Theory Grand Geometric Theory / Geometry-Gauge Theory; the integrated framework comprising GUT-10, GG-6, GDOS, Gi-4, and Qi-4.

GG-6 The six-dimensional bulk stage $M_4 \times I_r \times S_\theta^1$ carrying the unified gauge content $SU(5) \times U(1)_B \times U(1)_{Qi} \times \text{Gravity}$.

BI-6 Green-Schwarz-modified Bianchi identity $dH_3 = \text{tr } R^2 - \text{Tr } F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Qi}^2$, the anomaly-cancellation backbone.

OGN-5 Odd-integer Gauge Normalization ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ with $K = LJ = M^2 + N^2 + L^2$.

$U(1)_B$ The exact local baryonic $U(1)$ gauge factor; guarantees proton stability.

$U(1)_{Qi}$ The informational compact-angle $U(1)$ gauge factor on S_θ^1 .

TGG, CCI, GFF, GRIP The four Gi pillars of GG-A6-GRIP (Arisaka, 2026, A6-GRIP) (Total Geometry-Gauge conservation; Covariant Conservation Identity; Geometric Free-Energy Functional; Gauge-Runtime Invariance Pipeline).

DGI, GSSB, GA The three stages inside GRIP (Dynamic Gauge Invariance; Gauge-Selected Symmetry Breaking; Gauge Attractor).

Qi-4 The four pillars of the Qi runtime (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL).

GSSB lock crossing (throat cap) The infrared seam transition at M_{lock} , crossed during the radiation descent, that births the Kaxion and seeds ρ_Λ ; distinct from the end of inflation, which is the failure of slow roll at the edge of the primary plateau (GG-A10-Cosmos).

Arisaka Equation $J_E/Q_B = J_e = J_B = J_{Qi}$, the current-locking theorem of Qi-QUA.

Arisaka Ratio $(1, -1, -2/5, -2/5)$, the BI-6 coefficient vector.

Arisaka Constant $Q_B = (2/5)(E_{\text{Ryd}}/K) \approx 15.0 \text{ kJ/mol}$, the universal quantum of open-system operations.

f_{EDE} Fractional early-dark-energy contribution; the derived reading requires none, and a detected lobe refutes it (F-Hubble-3).

ΔN_{eff} Extra effective relativistic species; a derived null in the cohort, $\Delta N_{\text{eff}} \leq 0.03$ (GG-A10-Cosmos), with no source below the lock scale (the early side carries no radiation excess).

$\epsilon_{\text{local}}^{(B)}$ The total local excess, $\epsilon_{\text{local}}^{(B)} = 73.50/66.8 - 1 = 0.100$, carried whole as a predicted distance-ladder calibration systematic ϵ_{syst} ($\approx 10\%$); no void share is carried (ϵ_{Hole} is priced at the response law and not booked), and the $U(1)_B$ channel is closed by the termination and capped at $\sim 10^{-3}$ empirically.

ζ_B The screening ratio $\zeta_B \equiv (2/5) \delta\rho_B/\rho_B$ an earlier version’s graviphoton estimate used in a baryonic effective Newton constant (App. G.5, record only); the channel is closed under the boundary-condition termination of GG-A3-BI6. Written ζ_B (not η_B) to avoid collision with the cohort-wide baryon-to-photon ratio.

A.4 The honesty grades: a reader’s key

Every quantitative claim in this report carries a grade, quoted identically wherever the claim appears. The grades are the cohort’s codebook discipline (GG-A9-SM; GG-A10-Cosmos), and a reader who knows the key can tell at a glance how much any number is claiming.

Theorem. Proved within the framework’s stated axioms, with no named condition outstanding.

Example: the lock scale $M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K}$ (GG-A5-GDOS T8.7), pinned in the laboratory through the electroweak vev.

Derived (D*). Derived conditional on a single *named* lemma or conjecture, stated wherever the result is used. Example: the π -locked fractions (conditional on the closed-packet measure conjecture of GG-A10-Cosmos); the two-road background 66.6 ± 0.2 inherits this grade.

Derived (cond., MSMF/OGN). Selected by the codebook’s admissibility and naming rules, with the closing computation pre-registered as a falsifier. Example: $\zeta_\Lambda = \pi/175$ (F-Hubble-11).

Asserted. A declared input, flagged as such. The framework’s honesty requires these to be few; this paper’s Route I uses none.

Asserted-external. This paper’s own addition to the scale: a *cited allowance* from the field’s literature — quoted with the proposer’s caveats and the published rebuttals attached, priced in scenario brackets, never averaged into a GG-Theory value. Example: the three residual-anatomy channels of Part IV.

Named-open. A quantity the framework declares it cannot derive, carried as a debt with dated tests. Example: the +10% systematic itself.

The one-sentence rule for reading the paper: *a number’s grade travels with it, and no argument is allowed to borrow a stronger grade than its weakest input.*

B The conceptual on-ramp

Before the derivations: five primers that define the paper’s objects in plain words, then the questions a reader actually asks, from first meeting to seminar level.

B.1 Five primers: the objects of this paper, in plain words

B.1.1 The expansion rate $H(t)$

Take any two distant galaxies and divide their separation speed by their separation: the quotient is the same for every pair, and that quotient is H . It is not a speed — it is a *rate*, the fractional growth

of every cosmic distance per unit time, with units of inverse time; the astronomers' $\text{km s}^{-1} \text{Mpc}^{-1}$ is a historical dressing of $1/\text{time}$. Because the contents of the universe dilute as it grows, the rate falls with time: enormous in the hot beginning, 66.8 today, settling to 55.2 forever. Calling it the Hubble *constant* names one point on a falling curve — the point our era happens to read.

B.1.2 What “distance” means in an expanding universe

On an expanding stage there is no single distance. The distance light travel time suggests, the distance a ruler would measure today, the distance inferred from how faint a source looks (luminosity distance), and the distance inferred from how large it looks (angular-diameter distance) all differ, and they differ by known factors of $1+z$. Every H_0 measurement is a ratio of a *distance* to a *redshift*; every method's systematic risk lives in which distance it actually measures and what calibrates it. The ladder measures luminosity distances calibrated by geometry and stellar physics; the CMB measures one angular-diameter distance calibrated by a sound wave; the siren measures a luminosity distance calibrated by general relativity itself.

B.1.3 The sound horizon

For its first 380,000 years the universe was a plasma in which sound waves ran at over half the speed of light. The distance one such wave traveled before the plasma cleared — the sound horizon, $r_s = 144.7 \text{ Mpc}$ on the derived inventory — froze into the sky as a preferred spot size (the acoustic peaks) and into the galaxy distribution as a preferred separation (BAO). It is the early universe's built-in ruler: measure its apparent size, know its true size, and the geometry between is fixed. Every early-route H_0 is, at bottom, this one ruler read across thirteen billion years.

B.1.4 The standard siren

Two neutron stars spiraling together broadcast gravitational waves whose shape encodes the masses and whose loudness then encodes the distance — a candle whose absolute wattage is certified by the field equations, not by any calibration. The host galaxy supplies the redshift, and H_0 follows with no ladder and no sound horizon anywhere in the chain. One has been read (GW170817); the coming decade's catalogs make it a population. It is the referee of this paper: the one instrument both discordant routes agreed, in advance, to lose to.

B.1.5 The Local Hole

Galaxy counts show our few-hundred-megaparsec neighborhood is mildly under-dense. Matter flows outward from an under-density, so every expansion rate measured from inside reads slightly high — not because the universe expands faster here, but because the neighborhood adds an outflow on top of the expansion. The size of that lift is a computed response to a depth the surveys have not supplied: the depth the galaxy counts once credited would be an eleven-sigma fluctuation on the framework's own power spectrum, and the velocity data disfavor it, so the framework carries none of it. The survey programs of this decade will measure it directly, and the framework predicts they find nothing to carry.

B.2 Foundational questions about the Hubble tension

A reader meeting the derived reading for the first time will ask a short list of questions before the arithmetic means anything. This on-ramp answers them plainly, each with the physical picture first and the pointer to the rigor second.

B.2.1 What is the Hubble tension, physically?

Two ways of measuring how fast the universe expands today disagree. The early way reads the sound horizon imprinted on the microwave background and, under Λ CDM (Weinberg, 2008), infers $H_0 \approx 67$. The late way climbs a ladder of local distance indicators and measures $H_0 \approx 73.5$. The gap is now 7σ , too large to be a fluke and, after H0DN, not attributable to any single method. Either a systematic hides in one of them, or new physics separates the two epochs.

B.2.2 Why is the background rate 66.8 and not 73?

Because the rate is not measured against the ladder at all; it is *computed*. The geometry fixes the cosmic inventory in closed form ($\Omega_m = 1/\pi$, and the rest by the Hopf-fiber identification of §24); dividing those predicted fractions into the two physical densities the CMB measures returns $H_0^{(0)} = 66.6$ with no free parameter (App. E.7), and the geometry alone, with no sky input, returns 66.8 (§24). The number is the shadow of the inventory, not a fit to any distance.

B.2.3 If it is one rate, why do early and late disagree?

Because one of the two instruments reads the rate through calibrated rungs and the other does not. The CMB window lands on the background (the -0.9% is the inherited sound-horizon calibration, not a bias); the local ladder reads high by about ten percent, and the framework reads the whole of that as a calibration systematic in its rungs — predicted, not derived — to be tested by the probes that use no rung. The “tension” is the size of that systematic, not a contradiction in the rate.

B.2.4 What is ζ_Λ , and why is the seesaw “soft”?

ζ_Λ is the single dimensionless coefficient of the vacuum-energy seesaw, $H_0 = 498.3\sqrt{\zeta_\Lambda}$ (App. F.4). It is forced to be a ledger monomial in the OGN-5 integers; the bounded-cost cluster spans 0.0180–0.0204, and the codebook’s own admissibility and naming rules select $\pi/(LK) = \pi/175$ before any data enter (§24.2). The residual “softness” ($H_0 \propto M_{\text{lock}}^4$) is bounded in the laboratory: the same lock scale predicts the electroweak vev at $+0.045\%$, pinning the rate to $\pm 0.18\%$.

B.2.5 Why does a standard siren settle it when neither ladder nor CMB can?

Because a gravitational-wave standard siren reads the luminosity distance straight off the waveform and the redshift off the host galaxy — it uses *neither* a sound horizon *nor* a distance ladder. It therefore measures H_0 through a channel that shares no systematic with either discordant method. The derived reading predicts $H_0^{\text{GW}} \rightarrow 67$; a converged value at 73 refutes it outright (§57).

B.2.6 What would kill the derived reading?

Eleven named falsifiers, each with a pre-declared numerical rule (§56). The decisive one is the siren (F-Hubble-2); the single-event kill is proton decay (F-Hubble-1), which would retire the whole program. The reading is built to lose cleanly if it is wrong.

B.2.7 What exactly is a standard siren, and why is it immune to both routes?

A merging pair of neutron stars or black holes whose gravitational waveform is computed from first principles: the phasing gives the masses, the amplitude then gives the distance — a candle whose luminosity general relativity itself certifies. Its host galaxy gives the redshift. No sound horizon, no ladder, no calibration shared with either discordant route: the one clock-reader both camps must accept. The full physics, forecast, and budget are Part VII.

B.2.8 What is the Local Hole, and how sure are we it exists?

Galaxy counts show our neighborhood is mildly under-dense — matter flows out of it, so local expansion reads slightly high. That much is secure; the *depth and extent* are contested, from a shallow $\delta \approx -0.10$ to the disputed -0.3 over 300 Mpc; the depth the counts credit would be an eleven-sigma fluctuation on the framework's own power spectrum, and the velocity data disfavor it, which is why the paper adopts no lift and records the scenarios (§33). Full-sky infrared surveys will settle the profile this decade.

B.2.9 Why can't a void, or averaging effects, explain the whole tension?

Because the arithmetic is merciless: a +10% monopole needs a mean under-density of -0.56 across the supernova flow volume — excluded at $\gtrsim 12\sigma$ in standard structure formation — and the rigorous averaging (backreaction) term is bounded three orders of magnitude too small (App. H.4). The environment supplies percent-class effects; nothing in cosmological averaging supplies ten.

B.2.10 The acoustic angle is measured to four digits and the chart misses it by half a percent. How is that a success?

It is not a success by the standard a *fit* must meet, and the paper says so plainly (O-Hubble-32). A fit owns parameters whose job is to land inside the error bars; this chart owns none, claims percent-class derivation, and is displaced by one coherent, named systematic whose candidate origin is the theorem's unfinished measure lemma. The honest tests are the coherence of the displacement across all seventeen observables (Fig. 16) and the printed break condition: a converged inventory below $1/\pi$ kills the chart.

B.2.11 If Ω_m is exactly $1/\pi$, why do fits return 0.315?

Because a six-parameter fit distributes every small systematic across all six dials, and the fitted Ω_m is the fit's compromise, not a direct reading. The chart's claim is that the compromise sits one percent below the true value and drags every derived observable with it in one direction — which is

exactly the one-sided pattern the data show. Whether that pattern is the lemma’s residue or the chart’s epitaph is decided by the next generation of sub-percent inventory determinations.

B.2.12 What happens to the universe, on the derived curve?

Nothing dramatic, forever: $H(t)$ settles onto the de Sitter floor $H_\infty = 55.2$, the universe doubles in size every 12.3 Gyr, distant galaxies slide one by one across the horizon, and the sky slowly empties. The derived history has no turnaround, no rip, no cycle — with $w = -1$ exactly, the quiet ending is a theorem-grade consequence, guarded by F-Hubble-5. The clock that began at the first microsecond never stops; it just settles to its final rate.

B.2.13 What exactly does “Asserted-external” commit the paper to?

Three mechanics, all checkable. The value is quoted with the proposer’s own caveats and the published rebuttals attached (the priors channel is carried as “proposed, executed, and rebutted — undecided”); it is priced in scenario brackets rather than adopted (both combination rules, both scenarios); and it is never averaged into a GG-Theory number — 66.8 ± 0.2 contains no external channel. The grade exists so the paper can *use* the field’s contested corrections as evidence about the field without *claiming* them (O-Hubble-54 tests exactly this).

B.2.14 Why print two combination rules — quadrature and stack?

Because the right rule depends on correlations the field has not measured. Independent errors add in quadrature; fully common-mode errors add linearly; the three anatomy channels are neither cleanly independent (the prior recalibration re-treats supernova terms the dust channel also touches) nor fully common. So the paper prints both limits — 4.3%/3.8% and 9.3%/8.3% — and notes that the required +10% sits above the upper bound, which is the maximal honest statement: the published channels do not yet suffice, however they are correlated. A single quoted number would smuggle in a correlation assumption; the bracket is the honesty (§34.5).

B.2.15 Why is the recombination switch at e^7 , of all numbers?

Because in the derived chronology the radial coordinate $r = \ln(E_{\text{GUT}}/\mu)$ advances by integer-ledger steps, and the crossing from the plasma dictionary to the frozen geometric dictionary sits $J = 7$ e-folds of redshift above the present: $1 + z_* = e^7 \approx 1097$, against the measured $z_{\text{dec}} = 1089.9$ — a 0.6% agreement with an integer’s worth of freedom, i.e. none (§31). This is also where the ledger integer J enters the Hubble story; it is absent from the closed form for H_0 itself (§24.3).

B.2.16 How can $w = -1$ be *exact* when every fitted value has error bars?

Because in the derived reading the dark energy is not a dynamical fluid whose stiffness one measures; it is the residue of a geometric transition that is *complete*. A finished transition has no remaining dynamics to drag w off -1 : the value is exact the way a crystal’s ground-state symmetry is exact, not the way a measured coupling is precise. The measurable content is two-sided and live:

any converged $w \neq -1$ at 3σ retires the reading (F-Hubble-5), while every tightening of w toward -1 squeezes the dynamical-dark-energy rivals it competes with.

C Derivation course: the Friedmann background from the metric up

This is the first of the full derivation courses, and it sets their three-step contract. *Step one* is the textbook: the topic built from scratch, every component computed, at the level of a first-year graduate course — nothing assumed beyond special relativity, the equivalence principle, and the Einstein equation. *Step two* names what the standard treatment assumes and cannot explain. *Step three* derives the GG-Theory resolution and grades it. A reader who works this course with a pen can reproduce every background quantity in this paper — the equations, the age, the $H(z)$ rows, the de Sitter floor — and will know exactly which steps are century-old mathematics and which are this program’s claims.

C.1 The FRW metric and its conventions

Cosmology begins with an observation and a principle. The observation: averaged over scales above a few hundred megaparsecs, the universe looks the same in every direction. The principle (Copernican): we do not occupy a special place, so it looks that way from everywhere. Together they demand a spacetime that is spatially homogeneous and isotropic — and that symmetry is so restrictive that the metric is fixed up to one function of time and one discrete sign:

$$ds^2 = -dt^2 + a^2(t) \left[\frac{dr^2}{1-kr^2} + r^2 d\Omega_2^2 \right], \quad k \in \{-1, 0, +1\}, \quad (53)$$

in units $c = 1$, signature $(-, +, +, +)$ (conventions as in [Dodelson and Schmidt, 2020](#); [Weinberg, 2008](#)), with $d\Omega_2^2 = d\theta^2 + \sin^2 \theta d\phi^2$. Here t is proper time for observers at rest in the average matter (comoving observers); r, θ, ϕ are comoving coordinates, painted onto space once and carried along by the expansion; and $a(t)$ — the scale factor — is the single dynamical degree of freedom, the “size” of space. It is convenient to write the spatial part as $g_{ij} = a^2(t) \gamma_{ij}$, with γ_{ij} the comoving (time-independent) three-metric of constant curvature k .

Intuition. The metric says that the universe does not expand *into* anything: comoving galaxies keep their coordinates forever, and it is the ruler between them, $a(t)$, that grows. Every physical separation is (comoving label) $\times a(t)$; every question about cosmic history is a question about one function of one variable. That economy — an entire universe reduced to $a(t)$ and a sign — is the gift of symmetry, and everything in this course is the Einstein equation’s opinion about that one function.

C.2 Christoffel symbols: every non-vanishing component

The connection follows from the metric by the standard formula

$$\Gamma_{\mu\nu}^\lambda = \frac{1}{2} g^{\lambda\sigma} (\partial_\mu g_{\sigma\nu} + \partial_\nu g_{\sigma\mu} - \partial_\sigma g_{\mu\nu}). \quad (54)$$

The metric components are $g_{00} = -1$, $g_{ij} = a^2 \gamma_{ij}$; the inverse is $g^{00} = -1$, $g^{ij} = a^{-2} \gamma^{ij}$; and the only time dependence anywhere is the factor $a^2(t)$, so the only non-trivial time derivative is $\partial_0 g_{ij} = 2a\dot{a} \gamma_{ij}$. Now compute, term by term.

The time-space-space component. Set $\lambda = 0$, $\mu = i$, $\nu = j$ in Eq. (54). Since $g^{0\sigma}$ is non-zero only for $\sigma = 0$, and $g_{0i} = 0$,

$$\Gamma_{ij}^0 = \frac{1}{2} g^{00} (\partial_i g_{0j} + \partial_j g_{0i} - \partial_0 g_{ij}) = \frac{1}{2} (-1) (0 + 0 - 2a\dot{a} \gamma_{ij}) = a\dot{a} \gamma_{ij}.$$

The space-time-space component. Set $\lambda = i$, $\mu = 0$, $\nu = j$:

$$\Gamma_{0j}^i = \frac{1}{2} g^{ik} (\partial_0 g_{kj} + \partial_j g_{k0} - \partial_k g_{0j}) = \frac{1}{2} a^{-2} \gamma^{ik} (2a\dot{a} \gamma_{kj}) = \frac{\dot{a}}{a} \delta_j^i.$$

The purely spatial components. With all indices spatial, the factors of a^2 cancel between g^{ik} and ∂g_{kj} , leaving the connection $\hat{\Gamma}_{jk}^i$ of the comoving three-metric γ_{ij} alone (for $k = 0$ it vanishes in Cartesian comoving coordinates; for curved sections it is the standard constant-curvature connection). Every other component vanishes: $\Gamma_{00}^0 = \Gamma_{0i}^0 = \Gamma_{00}^i = 0$, because g_{00} is constant and the metric has no time-space cross terms.

Intuition. Two components carry all the physics. $\Gamma_{0j}^i = (\dot{a}/a) \delta_j^i = H \delta_j^i$ is the Hubble rate: it says that moving forward in time rescales every spatial vector by H per unit time — free particles separate, and peculiar velocities decay, at exactly this rate. $\Gamma_{ij}^0 = a\dot{a} \gamma_{ij}$ is its time-facing partner: stretching space costs time-direction curvature, which is why expansion can decelerate under gravity at all.

C.3 The Ricci tensor and the Ricci scalar

The Ricci tensor is $R_{\mu\nu} = \partial_\lambda \Gamma_{\mu\nu}^\lambda - \partial_\nu \Gamma_{\mu\lambda}^\lambda + \Gamma_{\lambda\sigma}^\lambda \Gamma_{\mu\nu}^\sigma - \Gamma_{\nu\sigma}^\lambda \Gamma_{\mu\lambda}^\sigma$.

The time-time component. Only the terms with Γ_{0j}^i survive:

$$R_{00} = -\partial_0 \Gamma_{0i}^i - \Gamma_{0j}^i \Gamma_{0i}^j = -\partial_0 \left(3 \frac{\dot{a}}{a} \right) - 3 \left(\frac{\dot{a}}{a} \right)^2 = -3 \left(\frac{\ddot{a}}{a} - \frac{\dot{a}^2}{a^2} \right) - 3 \frac{\dot{a}^2}{a^2} = -3 \frac{\ddot{a}}{a}.$$

Note the cancellation: the $\dot{a}^2/a^2$ pieces from the derivative and the quadratic term eat each other, leaving pure acceleration. The time-time Ricci component knows only about $\ddot{a}$ — this is the covariant ancestor of the Newtonian statement that tidal focusing measures the *acceleration* of test bodies, not their speed.

The space-space components. Assembling the same formula with the spatial connection included — and this time showing the three time-derivative terms the promise of this course owes, each a coefficient of γ_{ij} :

$$\underbrace{\partial_0 \Gamma_{ij}^0 = (\dot{a}^2 + a\ddot{a}) \gamma_{ij}}_{\text{term 1}}, \quad \underbrace{\Gamma_{\lambda 0}^\lambda \Gamma_{ij}^0 = 3\dot{a}^2 \gamma_{ij}}_{\text{term 2}}, \quad \underbrace{-2 \Gamma_{ik}^0 \Gamma_{0j}^k = -2\dot{a}^2 \gamma_{ij}}_{\text{term 3}},$$

which sum with the comoving three-space curvature (Ricci tensor $2k \gamma_{ij}$) to $(a\ddot{a} + 2\dot{a}^2 + 2k) \gamma_{ij}$, i.e.:

$$R_{ij} = \left(\frac{\ddot{a}}{a} + 2\frac{\dot{a}^2}{a^2} + 2\frac{k}{a^2} \right) g_{ij}.$$

Contracting, $R = g^{00} R_{00} + g^{ij} R_{ij} = 6 \left(\frac{\ddot{a}}{a} + \frac{\dot{a}^2}{a^2} + \frac{k}{a^2} \right).$

C.4 The Einstein tensor

With $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R g_{\mu\nu}$:

$$G_{00} = -3\frac{\ddot{a}}{a} - \frac{1}{2}(6) \left(\frac{\ddot{a}}{a} + \frac{\dot{a}^2}{a^2} + \frac{k}{a^2} \right) (-1) = 3 \left(\frac{\dot{a}^2}{a^2} + \frac{k}{a^2} \right),$$

the acceleration canceling exactly, and

$$G_{00} = 3\left(\frac{\dot{a}^2}{a^2} + \frac{k}{a^2}\right), \quad G_{ij} = -\left(2\frac{\ddot{a}}{a} + \frac{\dot{a}^2}{a^2} + \frac{k}{a^2}\right)g_{ij}. \quad (55)$$

Intuition. The cancellation in G_{00} is not an accident; it is the Bianchi identity at work. G_{00} is a *constraint*, not an evolution equation: it relates the expansion rate and spatial curvature to the energy present *now*, with no reference to acceleration — the general-relativistic analogue of energy conservation for the whole universe. G_{ij} carries the dynamics: how the expansion responds to pressure.

C.5 The perfect-fluid source

The same symmetry that fixed the metric fixes the source. Isotropy forbids any spatial vector (no momentum flux, $T_{0i} = 0$) and any anisotropic stress ($T_{ij} \propto g_{ij}$); homogeneity makes the two remaining scalars functions of time alone. Hence

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

a perfect fluid with energy density $\rho(t)$ and pressure $p(t)$, each species carrying an equation of state $w_i = p_i/\rho_i$: $w = 0$ for cold matter (pressureless dust), $w = 1/3$ for radiation (tracelessness of the electromagnetic stress tensor), $w = -1$ for vacuum energy (Lorentz invariance of the vacuum forces $T_{\mu\nu} \propto g_{\mu\nu}$).

C.6 The two Friedmann equations, the acceleration equation, and continuity

Insert Eq. (55) and the fluid into $G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G_N T_{\mu\nu}$. The time-time component gives the first Friedmann equation,

$$\boxed{H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G_N}{3} \rho - \frac{k}{a^2} + \frac{\Lambda}{3}} \quad (56)$$

— the speedometer reads the fuel gauge: energy content sets expansion rate, curvature and vacuum entering as bookkeeping terms of fixed dilution. The space-space component, after subtracting (56), gives the acceleration equation

$$\frac{\ddot{a}}{a} = -\frac{4\pi G_N}{3}(\rho + 3p) + \frac{\Lambda}{3}, \quad (57)$$

whose factor $\rho + 3p$ is the single most consequential combination in cosmology: *pressure gravitates*. Radiation ($p = \rho/3$) decelerates twice as hard per unit energy as dust; vacuum energy ($p = -\rho$) makes the combination negative and *accelerates* the expansion — not by pushing outward, but by gravitating with the opposite sign. Finally, conservation $\nabla_\mu T^\mu_0 = 0$ reads

$$\dot{\rho} = \underbrace{-\Gamma_{0i}^i \rho}_{\text{volume dilution}} - \underbrace{\Gamma_{0i}^i p}_{p dV \text{ work}} = -3\frac{\dot{a}}{a}(\rho + p), \quad \text{i.e.} \quad \dot{\rho} + 3H(\rho + p) = 0,$$

which is the first law of thermodynamics in disguise: for a comoving volume $V \propto a^3$ with energy $E = \rho V$, $dE = -p dV$ — expansion does $p dV$ work on each volume element. Integrated per species, $\rho_i \propto a^{-3(1+w_i)}$: matter dilutes as a^{-3} (number conservation), radiation as a^{-4} (number conservation *times* one redshifting wavelength), vacuum not at all. Only two of the three equations are independent — the Bianchi identity again — so the system is: one constraint, one evolution law, one bookkeeping rule.

Newtonian cross-check. Remarkably, Eq. (56) is older than relativity in disguise. Take a dust ball of radius $R = ar$, a test galaxy on its edge, and write energy conservation, $\frac{1}{2}\dot{R}^2 - G_N M/R = E$ with $M = \frac{4\pi}{3}\rho R^3$. Dividing by $R^2/2$: $\dot{a}^2/a^2 = \frac{8\pi G_N}{3}\rho + 2E/(r^2 a^2)$ — the Friedmann equation, with the curvature term unmasked as (twice) the Newtonian energy per unit mass: $k = -2Er^{-2}$. A bound universe ($E < 0$) is a closed one ($k > 0$); the critical density is the escape-velocity condition $E = 0$. Nothing about the expansion of the universe, at the background level, is more exotic than a stone thrown upward; general relativity's contributions are the *pressure* term of (57) and the license to apply the equation beyond the ball.

C.7 Curvature, the critical density, and the Ω bookkeeping

Define the critical density as the value that makes space flat today, $\rho_c = 3H_0^2/(8\pi G_N)$, and express every component as a fraction $\Omega_i = \rho_i/\rho_c$. Numerically, at the derived rate $H_0 = 66.8$: $\rho_c = 1.878 \times 10^{-29} (H_0/100)^2 \text{ g cm}^{-3} = 8.3 \times 10^{-30} \text{ g cm}^{-3}$ — about *five proton masses per cubic meter*. That the fate and geometry of the universe pivot on a density a laboratory vacuum exceeds by fifteen orders of magnitude is the sort of number worth carrying in one's head. Dividing Eq. (56) by H_0^2 and evaluating today,

$$1 = \Omega_r + \Omega_m + \Omega_\Lambda + \Omega_k, \quad \Omega_k \equiv -k/(a_0 H_0)^2,$$

and at general times, using the dilution laws,

$$H^2(a) = H_0^2 \left[\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_k a^{-2} + \Omega_\Lambda \right],$$

the master equation behind every row of Table 18. On the derived inventory the shares are not measured but *derived* (D*, the grade of Part III): $\Omega_m = 1/\pi = 0.3183$, $\Omega_\Lambda = 1 - 1/\pi = 0.6817$ (with the radiation share 9.32×10^{-5} computed from the measured CMB temperature and $\Omega_k = 0$ on the attractor), which is where this course hands back to the main text.

C.8 Worked numerics: the age, the acceleration epoch, and the $H(z)$ rows

Three computations a reader should do once by hand.

(a) *The age of the universe.* The age is the FRW integral

$$t_0 = \int_0^1 \frac{da}{a H(a)} = \frac{1}{H_0} \int_0^1 \frac{da}{a \sqrt{\Omega_m a^{-3} + \Omega_\Lambda + \Omega_r a^{-4}}}. \quad (58)$$

Neglecting the sub-permille radiation term, the matter-plus-vacuum integral has the closed form

$$t_0 = \frac{2}{3H_0\sqrt{\Omega_\Lambda}} \operatorname{arcsinh} \sqrt{\frac{\Omega_\Lambda}{\Omega_m}}. \quad (59)$$

Inserting the derived shares $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$ gives $\operatorname{arcsinh} \sqrt{\Omega_\Lambda/\Omega_m} = \operatorname{arcsinh}(1.464) = 1.175$ and prefactor $2/(3\sqrt{\Omega_\Lambda}) = 0.807$, so $t_0 = 0.948/H_0$. With the derived background rate $H_0 = 66.8$ (Hubble time $1/H_0 = 14.64$ Gyr),

$$t_0 = 0.948 \times 14.64 \text{ Gyr} = 13.9 \text{ Gyr}, \quad (60)$$

in agreement with the oldest-star and CMB determinations and consistent to a percent with the full numerical integral including radiation (13.88 Gyr). The age is thus not an input but a reading of the same inventory that fixes the rate: the de Sitter floor $H_\infty = H_0\sqrt{\Omega_\Lambda} = 55.2$ and the age 13.9 Gyr are two more shadows of $\Omega_\Lambda = 1 - 1/\pi$.

(b) *The acceleration epoch.* Setting $\ddot{a} = 0$ in Eq. (57) (matter era, vacuum included) gives $\rho_m = 2\rho_\Lambda$, i.e. $1 + z_{\text{acc}} = (2\Omega_\Lambda/\Omega_m)^{1/3} = 1.624$, $z_{\text{acc}} = 0.62$. The moment is prettier than it looks: at $\ddot{a} = 0$ the argument of the $\operatorname{arcsinh}$ in Eq. (59) is *exactly* $1/\sqrt{2}$ — independent of the shares — so $t_{\text{acc}} = \frac{2}{3} (1/\sqrt{\Omega_\Lambda}) \operatorname{arcsinh}(1/\sqrt{2}) H_0^{-1} = 7.8$ Gyr on the derived inventory. The universe decelerated for its first 7.8 billion years and has accelerated for the 6.1 billion since; the landmark table of Part III carries this as Landmark 5, and the exact $1/\sqrt{2}$ is a pleasant check that the closed form is being used correctly.

(c) *The $H(z)$ rows.* The master equation with the derived shares (radiation included) returns $H(0.5) = 88.6$, $H(1) = 120.1$, $H(2) = 203.5$, $H(4) = 425.3$ — the rows of Table 18 that the chronometer, BAO, and drift programs will confront. At the far end, $H \rightarrow H_0\sqrt{\Omega_\Lambda} = 55.2$: the de Sitter floor, on which the universe doubles in size every $\ln 2/H_\infty = 12.3$ Gyr, forever.

C.9 What the standard derivation assumes — and cannot explain

Everything above is a century old and beyond dispute — and it is worth listing exactly what it took as given, because the list is the tension's real cast of characters. (i) *The symmetry is*

assumed. The cosmological principle is imposed, then justified after the fact by the smoothness of the microwave sky; nothing in the standard derivation explains why the universe should have selected the maximally symmetric configuration. *(ii) The contents are measured, never derived.* The equations accept any ρ ; the actual shares — why 0.32, 0.68, 0.05, why not any other partition of unity — are inputs from the sky, with no principle behind them. *(iii) Λ is a catastrophe deferred.* The constant that dominates the budget is either left unexplained or mis-explained by one hundred and twenty orders of magnitude. *(iv) Flatness is an input dressed as an initial condition.* $\Omega_k = 0$ to a part in a hundred requires either fine-tuning or an inflationary prologue that is itself chosen, not derived. *(v) The present epoch is a coincidence.* Matter and vacuum share the budget within a factor of two only in a narrow window of cosmic history — and we happen to live in it. Five assumptions, five of the deepest open problems in physics. The standard course derives the *equations* magnificently and leaves the *solution the universe chose* entirely unexplained.

C.10 The geometric resolution: the Friedmann system as a theorem

The five assumptions of §C.9 are exactly where GG-Theory enters, and each is answered structurally rather than patched. In GG-Theory the Friedmann equations are not an assumed cosmological model but a *theorem*, forced from the boundary Einstein equation once the runtime has selected a homogeneous, isotropic attractor; the contents of ρ are read from the fixed integer ledger; Λ is the seesaw residue of §24.3; flatness is a property of the attractor; and the present epoch is the subject of the why-now analysis of the parent GG-A10-Cosmos. This subsection carries out the first of those answers — the theorem — because it is the rung on which the whole notion of a derived $H(t)$ stands.

From the boundary Einstein equation to a cosmology. The GDOS boundary law of GG-A2-GG6 delivers, on the four-dimensional stratum, an Einstein equation of the standard form

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G_N T_{\mu\nu},$$

valid for *any* boundary spacetime. A cosmology is the special solution in which the universe looks the same at every place and in every direction. In the textbook this cosmological 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 $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ of GG-A6-GRIP, run on the cosmological stratum, flows to the maximally symmetric configuration, so homogeneity and isotropy are *selected*, not posited — the classical solutions themselves being those of [Friedmann \(1922\)](#) and [Lemaître \(1927\)](#), and the parent’s appendix course carrying the full geometric derivation (Part III — the Friedmann-kinematics theorem block — and the boundary-law course of the Technical Appendices of [Arisaka, 2026, A10-Cosmos](#)). That the attractor is unique is a D^* result, resting on the attractor-uniqueness statement of GG-A6-GRIP; we state that dependence rather than bury it. Granting the attractor, the rest is the standard graduate computation.

The source: a perfect fluid from the inherited ledger. The same symmetry forces the boundary stress-energy to be that of a perfect fluid, $T^\mu{}_\nu = \text{diag}(-\rho, p, p, p)$: isotropy forbids

momentum flux or shear. Here is the one place **GG-Theory** departs from the textbook — not in the equation but in what may fill it. The density is not a free fluid to be fitted but the coarse-grained boundary projection of the bulk ledger, $\rho = \rho_r + \rho_b + \rho_{\text{dm}} + \rho_\Lambda + \dots$, each sector’s identity, equation of state $w_i = p_i/\rho_i$, and normalization fixed by the integer ledger $(1, 3, 5; 7, 35)$ of **GG-A4-Euclid** rather than chosen.

The equations, re-read. The computation of §C.2–§C.6 goes through verbatim on the attractor, and Eq. (56) is the object this paper calls $H(t)$: energy density sets the expansion rate. Nothing on the right is free — G_N is the throat-overlap integral of **GG-A2-GG6**, Λ carries the seesaw scale of §E.7, and the contents of ρ are read from the fixed ledger. This is the precise sense in which $H(t)$ is *derived*: it is the first integral of a boundary Einstein equation, restricted to a runtime-selected attractor, sourced by a ledger with no free dials.

The present value as one reading. Evaluating Eq. (56) today, with $k \rightarrow 0$ on the attractor and the π -locked shares $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, returns $H_0 = H(t_0)$ — a single point on a curve that runs from an unimaginable early rate to the de Sitter floor $H_\infty = H_0\sqrt{\Omega_\Lambda} = 55.2$. The Hubble “constant” is thereby exposed as a misnomer inherited from a century when only the present crossing could be measured: it is the value of a derived function at the moment we happen to read it, and the two dictionaries that disagree about it (§31) are reading one curve at one point, one of them through a calibrated ruler.

Status

The three-step contract, graded. Steps one and two (§C.1–§C.9) are standard mathematics and standard candor: nothing in them is a claim of this program. Step three carries two grades stated where they are used: the attractor uniqueness on which “theorem” rests is D^* (**GG-A6-GRIP**), and the shares that turn the equations into numbers are D^* under the parent’s measure conjecture. Every numerical output of this course — the age 13.9 Gyr, the acceleration epoch 7.8 Gyr, the $H(z)$ rows, the floor 55.2 — inherits exactly those conditions and no others, and every one of them is a falsifier: a robust stellar age above the FRW integral, or a measured $H(z)$ row off the master equation, kills the inventory that produced them.

D Derivation course: distances, the ladder, and the century of H_0

This is the paper’s signature course, because the subject of this paper is, at bottom, a dispute about *distances*. Step one builds the distance toolkit and the ladder from scratch, rung by rung, with the real error budget beside each rung; step two derives why the ladder’s own cross-checks are structurally blind to exactly one kind of error; step three reads the same budget through the derived background. A reader who works this course will know precisely what 73.50 ± 0.81 is a measurement *of* — which is the question the whole tension turns on.

D.1 The five distances of an expanding universe

On a static stage one number separates two points; on an expanding stage, five inequivalent ones, and confusing them costs factors of $(1+z)$. All descend from the comoving distance — the metric distance in Eq. (53) evaluated along a light ray ($ds = 0$):

$$D_C(z) = \int_{t_e}^{t_0} \frac{c dt}{a(t)} = \int_0^z \frac{c dz'}{H(z')}, \quad (61)$$

where the second form follows from $1+z = 1/a$ and $dz = -(1+z)H dt$. In a flat universe the transverse comoving distance D_M equals D_C . The two *observable* distances then follow from what an instrument actually registers.

Angular-diameter distance. An object of physical size ℓ at redshift z had that size when the universe was smaller by $1+z$; the angle it subtends today is $\theta = \ell/[a D_M] = \ell(1+z)/D_M$, so $d_A \equiv \ell/\theta = D_M/(1+z)$: distant objects look anomalously *large*, because we see them where they were.

Luminosity distance. Photons from a source of luminosity L arrive spread over the sphere $4\pi D_M^2$, but each photon is redshifted (energy down by $1+z$) and they arrive less often (rate down by $1+z$), so the flux is $F = L/[4\pi D_M^2(1+z)^2]$, i.e.

$$d_L = (1+z) D_M \quad \implies \quad d_L = (1+z)^2 d_A, \quad (62)$$

the second equality being the Etherington duality (Etherington, 1933) — true in *any* metric theory with photon-number conservation, which is why a violation of it would be bigger news than the tension. The astronomers' bookkeeping unit is the distance modulus $\mu = 5 \log_{10}(d_L/10 \text{ pc}) = m - M$; one magnitude is 46% in distance, a conversion this course will use repeatedly.

The low-redshift expansion. The ladder operates at $z \lesssim 0.15$, where Eq. (61) expands to

$$d_L(z) = \frac{cz}{H_0} \left[1 + \frac{1-q_0}{2} z + \mathcal{O}(z^2) \right], \quad q_0 \equiv -\frac{\ddot{a}a}{\dot{a}^2} \Big|_0 = \frac{1}{2}\Omega_m - \Omega_\Lambda = -0.523$$

on the derived inventory. The leading term is Hubble's law (Hubble, 1929); the correction, $0.761 z$, is 3.8% at $z = 0.05$ — not negligible at one-percent precision, and applied by every modern ladder through exactly this cosmography. *Intuition:* the ladder measures the local stretch of the curve of Course C; the tension is a dispute about whether that local stretch, correctly calibrated, sits on the global curve.

D.2 The ladder, rung by rung, with its error budget

Rung 0 — parallax. Geometry, and nothing else: a star at distance d sweeps an ellipse of semi-axis $p = 1 \text{ au}/d$ on the sky; $d[\text{pc}] = 1/p''$. The sensitivity is unforgiving, $\delta d/d = \delta p/p$, so a zero-point offset of tens of microarcseconds matters for every Milky-Way calibrator: Gaia's bright-star regime carries a measured offset check of $-38.9 \pm 10.3 \mu\text{as}$ with the sign of the induced H_0 shift unestablished (§34). The network dilutes the raw single-star sensitivity through its other anchors (NGC 4258's maser geometry, the LMC's detached eclipsing binaries), which is how the

paper’s budget carries this floor at 0.5–1% rather than several percent; Gaia DR4 (December 2, 2026) reads the entry either way.

Rung 1 — Cepheids. The physics is a driven acoustic resonance: a star’s fundamental pulsation period is its sound-crossing time, and since the mean sound speed scales as $c_s^2 \sim P_{\text{gas}}/\rho \sim GM/R$, the period obeys

$$\Pi \sim \frac{R}{c_s} \sim \sqrt{\frac{R^3}{GM}} \sim \frac{1}{\sqrt{G\rho}} \implies \Pi\sqrt{\rho} = \text{const},$$

the pulsation equation’s classical result. Brighter Cepheids are bigger, hence less dense, hence slower: luminosity follows period — the Leavitt law (Leavitt and Pickering, 1912), $M = \alpha \log_{10} \Pi + \beta$, calibrated on the geometric anchors of Rung 0 and applied in dust-hardened Wesenheit magnitudes $W = m_H - R(m_V - m_I)$, where the coefficient R is *assumed* from a dust law — the first shared convention on the ladder, and the ancestor of Part IV’s dust channel. The metallicity term is measured wavelength-independent from 0.5 to 4.5 μm — the chromatic test an artifact would fail — and carries $\sim 0.5\%$.

Rung 1’ — TRGB and JAGB. Two independent stellar candles cross-check the Cepheids. The tip of the red-giant branch (Freedman and Madore, 2010) is core physics: a degenerate helium core ignites at an essentially fixed core mass, so the brightest red giants in any old population reach the same luminosity ($M_I \approx -4.0$), visible as a sharp edge in the luminosity function. The J-region asymptotic giant branch uses carbon stars’ narrow luminosity distribution. Both bypass the Cepheid zero-point entirely — and *where the tip is measured* (which fields, which smoothing, which calibration of its absolute magnitude) is the convention dispute behind the 70.39 ± 1.94 vs 73.22 ± 2.06 split that Part IV prices as the subsample channel.

Rung 2 — Type Ia supernovae. The standardizable explosion: a white dwarf detonating near the Chandrasekhar scale, with the empirical Tripp standardization (the modern pipeline: Brout et al., 2022)

$$\mu = m_B - M_B + \alpha x_1 - \beta c,$$

stretch (x_1) and color (c) corrections with fitted global coefficients. The color coefficient is where dust hides: correcting flux for reddening requires a ratio of absorption to reddening (R_B), and the fitted $\beta \approx 3$ sits well below the Milky Way’s $R_B \approx 4.3$ — consistent only if the correction mixes intrinsic color with dust, or if calibrator and Hubble-flow supernovae see different dust. That fork *is* Part IV’s 2–4% dust channel, and the near-infrared ladder (F-Hubble-8) is its executioner: dust dies at long wavelengths.

Rung 3 — the Hubble flow. Finally, ~ 300 supernovae at $0.023 < z < 0.15$, far enough that peculiar velocities ($\sim 300 \text{ km s}^{-1}$, i.e. $\lesssim 4\%$ at the low edge, corrected with flow maps) and close enough that cosmography, not cosmology, carries the q_0 correction above. Here $H_0 = cz/d_L$ is finally read — through every convention the lower rungs installed.

D.3 Worked example: the common-mode algebra

Now the course’s central computation: what a shared calibration error does, in magnitudes. From $\mu = 5 \log_{10}(d_L/10 \text{ pc})$,

$$\delta\mu = \frac{5}{\ln 10} \frac{\delta d}{d} \iff \frac{\delta d}{d} = 0.4605 \delta\mu,$$

and since $H_0 = cz/d$, a calibration that makes every distance short by $x\%$ makes the rate high by $x\%$. The predicted systematic is therefore a statement about *one number*:

$$+10.0\% \text{ in } H_0 \iff \Delta\mu = 5 \log_{10}(73.50/66.8) = 0.208 \text{ mag}$$

shared across the calibrator set — and the field’s own contested corrections are of exactly this order (the rebutted prior-recalibration effect is $\sim 0.05 \text{ mag} = 2.3\%$; the dust reanalyses move 2–4%). *Why the ladder’s cross-checks cannot see it* is a five-line theorem. Let the network combine variant readings x_i with weights w_i , $\hat{y} = \sum_i w_i x_i$, and let every internal consistency statistic — leave-one-out shifts, variant scatter, anchor swaps — be some function $S(\{x_i - x_j\})$ of *differences*. Add a common mode: $x_i \rightarrow x_i + c$ for all i . Then every difference is unchanged, so S is unchanged — while $\hat{y} \rightarrow \hat{y} + c$. Internal audits are *exactly* orthogonal to common modes: the one error they cannot flag is the one that moves the answer. This is not a criticism of the ladder; it is a property of internal consistency itself, and it is why the paper assigns the verdict to instruments outside the network (the siren; the shell test) rather than to more cross-checking.

D.4 What the ladder cannot decide from inside

The limitations step, stated with the respect the measurement has earned. (i) *Shared layers are invisible* — the theorem above: one dust-law family, one prior family, one tip convention, one parallax floor, inherited by every variant at once. (ii) *A covariance-weighted consensus certifies its membership* — it measures, superbly, what the teams inside it share; it is silent on what its non-members dispute, and the one major program outside the membership reads 70.39 ± 1.94 (§43.2). (iii) *Precision cannot settle kind* — the ladder can measure its number to 1% forever without determining whether that number is the background rate or the background rate plus environment plus convention; that question needs either an external referee or a derived background to score against. These three limits define exactly the space in which Part IV operates.

D.5 The geometric reading of the same budget

Step three changes no measurement; it changes what the budget is scored against. Beneath: a derived background, 66.8 ± 0.2 (Course E and §24.3). On top of it: nothing carried — the environmental channel is priced at its response law and found to need a depth the surveys do not show (App. G.5). The whole +10.0% is then scored channel by channel against the rungs above: dust (2–4%, Rung 2’s R_B fork), priors (0–3%, the inference layer), subsample (1–2%, Rung 1’s convention dispute), plus the floor entries. The combination algebra is worth one derivation, because

it is where the honesty lives. For channels σ_i with pairwise correlations ρ_{ij} ,

$$\sigma_{\text{tot}}^2 = \sum_i \sigma_i^2 + 2 \sum_{i < j} \rho_{ij} \sigma_i \sigma_j,$$

whose limits are the paper’s two printed scenarios: all $\rho_{ij} = 0$ gives the quadrature 4.3%/3.8%; all $\rho_{ij} = 1$ gives the linear stack 9.3%/8.3%. The observed requirement +10.0% sits above the linear stack, so no uniform effective correlation reproduces it — solving the same formula backwards returns $\bar{\rho} > 1$ in both scenarios, which is the algebra’s way of saying the published channels are short however they are correlated — by under a percent if every one aligns, by four to five percent at their chained central values. That is an *illustration*, not a new claim: it quantifies the verdict section’s statement that the anatomy is real and not yet sufficient. Every channel carries its dated, two-sided test (F-Hubble-8/9/10), and either outcome of each is printed in advance.

Status

Step one is standard astronomy, presented without a single GG-Theory claim; the ladder it describes is among the great precision instruments of science. Step two is arithmetic (the common-mode theorem) and sociology-free bookkeeping (the certification boundary). Step three inherits its grades from the main text: the background Derived (cond.) and D*, the void channel priced and not carried, every channel Asserted-external, the systematic named-open, and the effective-correlation number an illustration on printed values. Nothing in this course re-reduces anyone’s photometry; that boundary is the tier ladder’s (§82).

E Derivation course: the microwave background as an H_0 instrument

The early side of the tension is one measurement chain read backwards through thirteen billion years, and this course builds it: the plasma physics, the ruler it froze, the angle that ruler subtends, and — the step most accounts blur — exactly what the microwave sky does and does not measure about the expansion rate. The punchline of step two is worth stating in advance: *the CMB never measures H_0* ; it measures two densities and one angle, and every quoted “CMB H_0 ” is an inversion conditional on a model. That fact is not a weakness of the CMB; it is the opening this paper’s Route II walks through.

E.1 The tightly coupled plasma and its sound speed

Before recombination, Thomson scattering couples photons and baryons into a single fluid: the scattering rate exceeds the expansion rate by many orders, so photons and baryons move together, and the fluid has a pressure (the photons’) and an inertia (both). Perturbations in this fluid obey a driven oscillator: gravity compresses a k -mode into potential wells, photon pressure resists, and the balance rings. The spring constant is the sound speed. Derive it: the pressure is the photons’,

$p = \rho_\gamma/3$; the inertia is both species', $\rho_\gamma + p_\gamma + \rho_b = \frac{4}{3}\rho_\gamma(1 + R)$ with the baryon loading

$$R \equiv \frac{3\rho_b}{4\rho_\gamma} = \frac{3\omega_b}{4\omega_\gamma} \frac{1}{1+z} \approx \frac{678}{1+z},$$

so

$$c_s^2 = \frac{\delta p}{\delta \rho} = \frac{\frac{1}{3}\delta\rho_\gamma}{\delta\rho_\gamma + \delta\rho_b} = \frac{c^2}{3(1+R)}.$$

Baryons add weight but no pressure: they slow the sound and deepen the compressions — which is precisely why the *relative heights* of odd (compression) and even (rarefaction) peaks read off ω_b , the first of the sky's two h -free density measurements. At the switch, $R(z_*) = 0.62$ and $c_s = 0.454c$: the early universe rang at almost half the speed of light.

E.2 The sound horizon, evaluated

The ruler is the distance such a wave travels before the plasma clears:

$$r_s = \int_{z_*}^{\infty} \frac{c_s(z)}{H(z)} dz, \quad c_s(z) = \frac{c}{\sqrt{3[1+R(z)]}}, \quad R = \frac{3\rho_b}{4\rho_\gamma} = \frac{3\Omega_b}{4\Omega_\gamma} \frac{1}{1+z}, \quad (63)$$

with $z_* = e^7$ on the derived chart. Evaluate it in pieces, so the anatomy is visible. Deep in the radiation era ($z \gg z_{\text{eq}} = 3400$), $H \propto (1+z)^2$ and $R \rightarrow 0$: the integrand falls as $(1+z)^{-2}$, so the early contribution converges — the ruler does not care about the first instants. Near equality, $R(z_{\text{eq}}) = 0.20$: baryon loading is already slowing the wave. From equality to the switch, matter drives $H \propto (1+z)^{3/2}$ while R grows to 0.62; this stretch contributes the bulk of the length. Carrying out the quadrature with the derived inventory at the carried $h = 0.6676$ gives 144.2 Mpc from this two-fluid integral. The last 0.5 Mpc deserves its own ledger, because it is the one step of this course we do not display line by line: (i) the detailed radiation inventory near equality (neutrinos in the pre-recombination plasma), +0.6; (ii) for calibration of the scale of such conventions, the choice of lower limit alone — the ledger's switch $z_* = e^7 = 1096.6$ versus the measured centroid 1089.9 — is worth another ≈ 0.6 ($c_s/H \approx 0.09$ Mpc per unit redshift at the epoch); half-percent conventions, all of them, and the reason r_s carries no sub-percent claim anywhere in this paper. The full evaluation returns the landmark value

$$r_s(z_*) = 144.7 \text{ Mpc}$$

carried throughout this paper (landmark 4; the released Boltzmann pipeline of GG-A10-Cosmos at the carried rate), the sub-percent difference being exactly the level at which the simple two-fluid picture stops being the whole story. One disambiguation the literature often blurs: baryons remain loaded on the photons slightly *after* the photons decouple, so the ruler frozen into the *galaxy* distribution is the slightly longer drag-epoch horizon. On the derived background the same two-fluid quadrature, evaluated at the standard drag epoch $z_d \approx 1060$, gives $r_d \simeq 147.2$ Mpc (the integrator reproduces the concordance $r_d = 147.1$ to 0.1% on Planck inputs, so the derived-chart shift is real, not numerical) — the BAO ruler. The CMB's acoustic angle uses $r_s(z_*) = 144.7$; the galaxy surveys use $r_d = 147.2$; this paper's landmark is the former, and every BAO sentence in this paper spends the latter.

E.3 The acoustic angle and the peak series

The observable is the angle the ruler subtends across the universe:

$$\theta_* = \frac{r_s}{D_M(z_*)}, \quad D_M(z_*) = \int_0^{z_*} \frac{c dz}{H(z)}, \quad (64)$$

measured to four significant figures, $\theta_* = (1.0411 \pm 0.0003) \times 10^{-2}$. On the derived chart the denominator evaluates to $D_M(z_*) = 13,940$ Mpc, giving $100\theta_* = 1.038$ — the released pipeline’s 1.0381 at the carried rate, 0.3% below the measured value, displaced *coherently* with the sixteen other observables of the offset audit, and the paper scores that displacement honestly rather than claiming a bullseye (O-Hubble-32; Fig. 16). The peak multipoles follow as a harmonic series with a gravitational phase shift,

$$\ell_m \simeq (m - \phi_m) \frac{\pi}{\theta_*}, \quad \frac{\pi}{\theta_*} = 303,$$

with $\phi_1 \approx 0.27$ from the decay of the potentials at horizon crossing: $\ell_1 \approx 220$. The precise version of this computation is the Boltzmann audit of §26, which returns 221/539/816 against the observed 220/536/813 on the exact derived fractions at the carried rate.

E.4 Why the microwave sky never measures h

Now the course’s load-bearing derivation. What the peak pattern measures, cleanly, is three things: ω_b (odd/even height ratio — baryon loading), ω_m (the height envelope through the equality epoch and the potential-decay driving), and θ_* (the positions). Notice what is *not* on the list. The physical densities $\omega_i = \Omega_i h^2$ carry no separate h ; and θ_* constrains h only through the late-time integral $D_M = \int c dz/H(z)$, where $H(z)$ depends on h and on the late-time budget. Hold ω_b and ω_m fixed (the early universe unchanged, r_s unchanged) and vary h : flatness re-partitions the budget, $\Omega_\Lambda = 1 - \omega_m/h^2$, and D_M changes — so *given* flatness and $w = -1$, the measured θ_* picks out one h . Relax either assumption and the same θ_* admits a family of (h, w, Ω_k) combinations: the geometric degeneracy. Hence the precise statement the tension literature often compresses: *the “CMB value of H_0 ” is an inversion of two densities and one angle through an assumed late-time model*. Standard Λ CDM performs that inversion by fitting Ω_i ; Route II of this paper performs it by *deriving* them — same sky, same densities, opposite direction of inference.

E.5 The audit: zero background freedom against the whole spectrum

The derived chart does not stop at three summary statistics; it is scored against the full temperature spectrum with zero background freedom: the exact fractions, the carried rate, nothing tuned, $\chi^2 = 187$ across the 83 band powers; on this background χ^2 is monotone in the tilt and does not select the computed value (§26). A fit this rigid cannot hide a mistake in one parameter inside another — there are no others — which is why the residual pattern matters more than its size: seventeen observables, one direction, percent scale, the signature of a single coherent displacement rather than scattered errors. Step two of this course, and the precedent lessons of §78, meet exactly here: coherence is the statistic that cable faults and foregrounds do not fake.

E.6 The sound-horizon lever, and its cap

The limitations step, aimed at the early-side fixes. At fixed measured θ_* and fixed late-time physics, D_M must scale with r_s ; and since D_M scales roughly as $1/H_0$ for small changes, the inferred rate obeys $\delta H_0/H_0 \approx -\delta r_s/r_s$ — the lever of Eq. (35) that every early-universe “fix” pulls. To lift the inference from 67.4 to the ladder’s 73.5 requires $r_s : 145 \rightarrow 133$ Mpc, an eight-percent amputation of the ruler in the last few e-folds before recombination — which is precisely what early dark energy and modified recombination are engineered to do. The cap is now empirical from two directions: the joint EDE searches find no lobe of the required size (F-Hubble-3), and the sound-horizon-*free* measurements — which would be untouched by any ruler surgery — span 66.7–70.2, closing the road to 73 from below (§43.3). The lever exists; the room to pull it does not.

E.7 The geometric resolution: the rate from two densities, and the coherent offset

The derived reading pulls no lever: r_s stays at its derived length, the recombination history stays standard, and the rate is obtained by running the degeneracy argument of §E.4 in the opposite direction — fractions derived, densities measured, h forced. In full arithmetic:

E.7.1 Inputs

Two model-independent physical densities from the microwave background, neither carrying h :

$$\omega_b \equiv \Omega_b h^2 = 0.02237 \pm 0.00015, \quad (65)$$

$$\omega_m \equiv \Omega_m h^2 = 0.1430 \pm 0.0011, \quad (66)$$

and the two GG-predicted fractions,

$$\Omega_b = \frac{1}{2\pi^2} = 0.050661\dots, \quad (67)$$

$$\Omega_m = \frac{1}{\pi} = 0.318310\dots \quad (68)$$

E.7.2 The two anchors

Since $\omega_i = \Omega_i h^2$, each pairing gives h directly:

$$h_b = \sqrt{\omega_b/\Omega_b} = \sqrt{0.02237/0.050661} = 0.6645, \quad (69)$$

$$h_m = \sqrt{\omega_m/\Omega_m} = \sqrt{0.1430/0.318310} = 0.6703, \quad (70)$$

i.e. $H_0 = 66.45$ (baryon anchor) and 67.03 (matter anchor).

E.7.3 Combination and error

The two anchors rest on different spectral statistics and agree to 0.9%. Let us do the bookkeeping in the open, because a referee will. Their naive inverse-variance combination is 66.69 ± 0.17 ; the

sky-free BBN variant below reads 66.4 ± 0.5 ; the geometric Route I gives 66.8 ± 0.2 ; the closed form evaluates to 66.7. The three density routes are not equally clean — the two *baryon* routes are the model-cleanest (ω_b is the one density measured two independent ways, CMB and BBN, agreeing at 66.4–66.5), while ω_m at ± 0.0011 precision is a concordance-chain output — so a formal weighting would itself be a modeling choice. The two-road value this paper’s companions carry is therefore

$$\boxed{H_0^{(0)}|_{\text{two-road}} = 66.6 \pm 0.2 \text{ km s}^{-1} \text{ Mpc}^{-1}}, \quad (71)$$

the *center of the route span with a structural, spread-inflated error* — an adopted center, not a formal inverse-variance mean, and we say so plainly. The rate this paper carries, with the parent GG-A10-Cosmos, is the geometric 66.8 ± 0.2 of Route I; every conclusion of this paper survives any choice in the 66.4–66.8 span.

Error propagation, explicitly. Since $h = \sqrt{\omega_i/\Omega_i}$ with Ω_i exact (ledger-fixed, no measurement error), $\sigma_h/h = \frac{1}{2} \sigma_{\omega_i}/\omega_i$. The baryon route carries $\frac{1}{2} \times 0.67\% = 0.34\%$, i.e. ± 0.22 ; the matter route $\frac{1}{2} \times 0.77\% = 0.38\%$, i.e. ± 0.26 . The two routes rest on different spectral statistics (baryon loading versus the peak-height envelope), so their errors are treated as independent; their central-value spread (0.5) is comparable to either error, which is why the adopted error is structural rather than formal. The honest statement is that the multi-route *consistency*, not any single weighting, carries the weight — and every conclusion in this paper survives any choice in the 66.4–66.8 span the combined routes bracket: the coherent π -offset moves between -0.9% and -1.5% , and nothing downstream changes sign or verdict. The geometric route of Eq. (20) gives $H_0 = 498.3\sqrt{\pi/175} = 66.8$ at the codebook-selected $\zeta_\Lambda = \pi/175 = 0.01795$ (§24.2), concordant with the inventory value at 0.6σ . No quantity in this calculation is adjusted to the local distance-ladder measurement — or, on the geometric route, to any astronomical measurement at all.

The variant that removes the microwave sky entirely — the BBN deuterium abundance with laboratory reaction rates, $\omega_b^{\text{BBN}} = 0.02233 \pm 0.00036$, hence $H_0^{\text{BBN}} = 66.4 \pm 0.5$ — closes the loop: two densities, three routes, one rate, no lever. What remains between the derived chart and the fitted sky is the coherent -0.9% π -offset across the seventeen-observable audit — carried as a named, one-directional displacement whose candidate origin is the measure lemma, and priced as such in the honesty triage. On this reading the early side of the tension was never broken, and Course D’s ladder is measuring exactly what Part IV says it measures.

Status

Step one is textbook plasma physics and geometry; the evaluated numbers ($R_* = 0.62$, $c_s = 0.454c$, $D_M = 13,970$ Mpc, $\phi_1 \approx 0.27$) are reproducible with the released scripts and a quadrature routine. Step two — the never-measures- h derivation and the lever algebra — is model bookkeeping, not GG-Theory. Step three inherits the main text’s grades: fractions D^* , the audit’s zero-freedom framing as printed, the half-percent acoustic displacement carried open at the O-Hubble-32 standard. The landmark 144.7 and the drag ruler 147 are now explicitly distinguished; nothing else in the paper changes.

F Derivation course: the vacuum energy of the universe and the seesaw

The deepest number in this paper is not the Hubble rate; it is the energy density of empty space that sets it. Step one of this course states the vacuum-energy problem the way quantum field theory actually produces it — with the integral written out, so the notorious “ 10^{120} ” is a computation rather than a slogan; step two shows why every standard mechanism misses the observed scale by tens of orders; step three derives the geometric seesaw and collapses it into the closed form of §24.3 (mirrored at Route I of the parent’s §6.4; [Arisaka, 2026, A10-Cosmos](#)), whose two corollaries then *dissolve* the accounting problem of step one into two powers of one hierarchy.

F.1 The vacuum energy of quantum fields, computed honestly

Every mode of every quantum field carries zero-point energy $\frac{1}{2}\hbar\omega_k$. Summing over modes up to a cutoff Λ_c (natural units),

$$\rho_{\text{vac}} = \int^{\Lambda_c} \frac{d^3k}{(2\pi)^3} \frac{k}{2} = \frac{1}{4\pi^2} \int_0^{\Lambda_c} k^3 dk = \frac{\Lambda_c^4}{16\pi^2},$$

one line of quadrature (the classic accounting is [Weinberg, 1989](#)) and the origin of the crisis: the answer is the *fourth power of whatever scale one trusts the theory to*. Trust it to the reduced Planck scale and $\rho_{\text{vac}} \sim (2.4 \times 10^{18} \text{ GeV})^4/16\pi^2$; the measured dark-energy density is $\rho_\Lambda^{1/4} = 2.24 \text{ meV}$, i.e.

$$\frac{M_{\text{Pl,red}}^4}{\rho_\Lambda} = \left(\frac{2.435 \times 10^{18} \text{ GeV}}{2.24 \times 10^{-12} \text{ GeV}} \right)^4 = 10^{120.2},$$

the worst prediction in the history of physics, computed in two steps. *Intuition:* the fourth power is the whole cruelty — halving the cutoff buys barely one of the hundred and twenty orders, so no modest re-scaling helps; either almost-exact cancellation or a structurally tiny input scale is required.

F.2 Seesaw structures: how physics makes small numbers

Nature’s standard trick for a small scale without a small input is the seesaw. Diagonalize the two-state mass matrix

$$\mathcal{M} = \begin{pmatrix} 0 & m \\ m & M \end{pmatrix}, \quad m \ll M \implies \lambda_\pm = \frac{M \pm \sqrt{M^2 + 4m^2}}{2} \simeq M, -\frac{m^2}{M} :$$

a heavy state near M , and a light one at the *geometric combination* m^2/M — pushed down exactly as far below m as M sits above it, whence the name. The neutrino sector is the classic instance ($m_\nu \sim v^2/M_{\text{GUT}}$); the pattern is general: to manufacture a tiny scale, couple an intermediate scale to a heavy one. For the vacuum, the required output is meV; a seesaw delivers it if some intermediate scale m satisfies $m^2/M_{\text{Pl,red}} \sim \text{meV}$, i.e. $m \sim \sqrt{(2.4 \times 10^{18})(2 \times 10^{-12})} \sim \text{a few TeV}$.

Read that sentence twice: *the vacuum energy of the universe points at the TeV scale* — the one scale particle physics independently cares about. The observation is old; what has always been missing is a derivation of *which* TeV quantity enters and *what* dimensionless coefficient dresses it.

F.3 What the standard mechanisms cannot do

The limitations step, mechanism by mechanism. *Cancellation by symmetry*: unbroken supersymmetry cancels the zero-point sum exactly — but supersymmetry is broken at or above the TeV scale, leaving $\rho_{\text{vac}} \sim (\text{TeV})^4$, still sixty orders too large; the symmetry route buys half the exponent and stalls. *Dynamical relaxation (quintessence)*: trades the constant for a field, at the price of a potential whose mass scale ($\sim 10^{-33}$ eV) and flatness are new fine-tunings, plus fifth-force constraints; the problem is relocated, not solved. *Anthropic selection*: makes the smallness an observation about observers, at the price of a multiverse whose measure is undefined — an accounting of last resort, and the philosophical opposite of this series’ program. The scorecard: no standard mechanism *derives* the meV scale; the best either cancel too much (to zero) or too little (to TeV^4).

F.4 The geometric resolution: the seesaw, derived and closed

The dark-energy sector delivers H_0 a second time, structurally unrelated to the density ratio. Folding the geometric seesaw vacuum density and the flat late-time budget into the first Friedmann equation gives

$$H_0 = \frac{1}{M_{\text{Pl,red}}} \left(\frac{M_{\text{lock}}^2}{M_{\text{Pl,red}}} \right)^2 \sqrt{\frac{\zeta_\Lambda}{3\Omega_\Lambda}}, \quad (72)$$

in which every symbol is inherited: $M_{\text{lock}} \approx 3.85$ TeV is the lock scale (a theorem of the scale chain, GG-A5-GDOS T8.7), $M_{\text{Pl,red}}$ the reduced Planck mass from the warp overlap, and $\Omega_\Lambda = 1 - 1/\pi$ the inventory value. Numerically this reduces to $H_0 = 498.3\sqrt{\zeta_\Lambda} \text{ km s}^{-1} \text{ Mpc}^{-1}$. The coefficient ζ_Λ is not free: MSMF forces it to be a bounded-cost monomial in the OGN-5 integers $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ with the transcendentals $\{\pi, e\}$; the admitted stratum cluster spans 0.0143–0.0359 at cost ≤ 5 with one member in the observed decade (§24.1), and the codebook’s stratum-availability and structural-naming rules select $\zeta_\Lambda = \pi/(LK) = \pi/175$ with no reference to data (§24.2), so $H_0 = 66.8 \pm 0.2$. The route’s M_{lock}^4 sensitivity is pinned in the laboratory through the electroweak vev (App. B.2); its role is primary: the rate from geometry alone, which the density inversion then confirms.

Two remarks complete the promotion of this compact chain to course level. *First, the bookkeeping*. In reduced units the lock theorem reads $M_{\text{lock}} = \sqrt{2\pi} M_{\text{Pl,red}} e^{-K}$, so the seesaw mass is $m_\Lambda = M_{\text{lock}}^2/M_{\text{Pl,red}} = 2\pi M_{\text{Pl,red}} e^{-2K} = 6.08 \text{ meV}$ — the intermediate scale of §F.2 is not postulated at the TeV scale; it *is* the TeV lock, a theorem with a laboratory pin, and the coefficient that dresses m_Λ^4 is the codebook-selected $\zeta_\Lambda = \pi/175$, pre-registered at F-Hubble-11. Substituted into the Friedmann equation, every power of the Planck mass cancels except one overall rate, and the whole course collapses to the boxed line of §24.3: $H_0 = e^{-4K} \sqrt{2\pi^5 c^5 / [3LK(\pi - 1)\hbar G_N]} = 66.7$. *Second, the dissolution of step one*. The ratio that opened this course is, on the derived chain, $\rho_\Lambda/\rho_{\text{Pl,red}} = \zeta_\Lambda(2\pi)^4 e^{-8K} = 10^{-120.2}$: the hundred and twenty orders are $8K$ e-folds of one derived

descent, corrected by 1.4 orders of geometry. Nothing cancels; nothing is tuned; the smallness is the fourth power of a hierarchy the same ledger uses at the collider. The error budget is as printed in §24: the M_{lock}^4 sensitivity is pinned through the electroweak vev ($v = M_{\text{lock}}(2/5)^3 = 246.33$ vs measured 246.22 GeV, a +0.045% agreement mapping to $\pm 0.18\%$ on the rate), and the selection lemma's window $[0.01784, 0.01806]$ is the one conditional link, cashed by the Hosotani computation.

Status

Step one is textbook quantum field theory; step two is the field's own candid literature. Step three carries the printed grades: the lock scale a theorem (T8.7) with a laboratory pin; the seesaw identification and the π -locked Ω_Λ D*; the coefficient Derived (cond.) under the selection lemma, two-sidedly falsifiable. The course adds no new claim; it adds the derivation chain at a depth a graduate student can audit line by line.

G Derivation course: the local universe — peculiar velocities, the Hole, and the lift a void would give

The tension's late side is measured from inside a particular neighborhood, and this course derives what a neighborhood does to a rate measurement. Step one builds linear peculiar-velocity theory from the continuity equation and lands on a one-line master formula; step two applies it to the measured Local Hole and shows — from the same line — both what an honest environment *can* supply and what it *cannot*; step three states what the framework supplies (the response law, not the depth — so no lift is carried), records the framework's one exotic channel as closed, and states the shell-level prediction the surveys can test.

G.1 Linear peculiar velocities from the continuity equation

Perturb the matter density, $\rho = \bar{\rho}(1 + \delta)$, and let $\mathbf{v}$ be the peculiar velocity (motion over and above the Hubble flow). Mass conservation in comoving coordinates reads $\dot{\delta} + \nabla \cdot \mathbf{v}/a = 0$: where matter converges, density grows. In the linear growing mode every Fourier amplitude grows as the growth factor $D(a)$, so $\dot{\delta} = \delta d \ln D/dt = \delta fH$, with

$$f \equiv \frac{d \ln D}{d \ln a} \simeq \Omega_m^{0.55} = 0.533$$

on the derived inventory — the growth rate, the one number linear theory needs. Substituting back, $\nabla \cdot \mathbf{v} = -afH\delta$: velocities are sourced by density contrasts, pointing from voids toward walls. *Intuition*: gravity is a book-keeper — whatever density history a region has, its velocity field must account for it, which is why measured flows can audit claimed voids.

G.2 The master formula: what an under-density does to a local rate

Now the course's central line. Apply Gauss's theorem to a sphere of radius r containing mean interior contrast $\bar{\delta}(< r)$: the outflow through its surface is $v(r) = -\frac{1}{3}fH\bar{\delta}(< r)r$ (radial, by

symmetry). An observer inside, measuring recession speeds of tracers at r , reads $v_{\text{obs}} = Hr + v(r)$, i.e. an effective local rate

$$\boxed{\frac{\delta H}{H} = -\frac{f}{3} \bar{\delta}(< r) = -0.178 \bar{\delta}(< r)} \quad (73)$$

on the derived inventory. One line, three consequences, and the whole of Part IV’s environmental arithmetic:

(i) *The response.* The surveyed Local-Hole depths, weighted over the ladder’s volume, span $\bar{\delta} \approx -0.11$ to -0.17 in exactly the bracket the counts report (§33.1); Eq. (73) converts them to $\delta H/H = +2\text{--}3\%$ — the lift a void of that depth would give, as arithmetic, which the flow data do not permit (§33.1) and the framework does not carry.

(ii) *The impossibility.* Ask the same line to produce the *whole* excess, $+10\%$, and it answers with $\bar{\delta} = -0.100/0.178 = -0.56$ across the supernova flow volume — the deep-void requirement of §35, here revealed as the same formula read at a different target. The two numbers the paper prices — the modest response it computes and does not carry, and the -0.56 it excludes — are one equation at two inputs.

(iii) *The signature.* Because $\bar{\delta}(< r) \rightarrow 0$ as the sphere grows past the Hole, the locally read rate must *decay toward the background* as low shells are cut away: $H_0(z_{\text{min}}) = H_0[1 - \frac{f}{3} \bar{\delta}(< r(z_{\text{min}}))]$ — a shape no constant calibration error can imitate, drawn on the measured profiles in Fig. 17 and confronted with DESI’s binned analyses in §33.1.

G.3 The measured Hole, and the velocity audit

Step one’s formula needs a depth, and the depth is measured two independent ways. *Counts:* galaxy number counts across optical and near-infrared surveys report a mild under-density over a few hundred megaparsecs, with the profile’s depth the contested quantity — from a shallow $\delta \approx -0.10$ to the disputed -0.3 (§33). *Flows:* the same velocities that Eq. (73) predicts are directly measurable as bulk flows and as the smoothness of the supernova Hubble diagram — and the field-level velocity modeling now disfavors the deep reading: a -0.3 hole over 300 Mpc would drag the local diagram visibly, and it does not (§33.1). The paper adopts no lift: the counts license a few percent and the flows do not permit it, and the intersection of the two audits is empty at the precision in hand — which is what “priced and not carried” means.

G.4 What the environment cannot explain

The limitations step, in three sentences. A local environment can shift a local rate by $\sim (f/3) \times (\text{depth})$ — percent-class for any depth the sky allows, and the required -0.56 is a $> 12\sigma$ fluctuation on the flow scale (Course H prices it exactly). An environment cannot touch the early side at all: the CMB’s densities are read at $z \sim 1100$, oblivious to our suburb. And an environment predicts a *profile* — the decay of (iii) above — so it is not an infinitely flexible excuse: the shell test can kill any lift tomorrow, and this paper predicts a flat profile. Percent-class, one-sided, and falsifiable: exactly the three properties the paper’s usage respects.

G.5 The geometric resolution: the response law, and the closed exotic channel

Step three has two parts. The response law is not exotic: it is Eq. (73) evaluated on measured depths, with GG-Theory supplying the *sign and existence* argument (an open-boundary neighborhood must register its environment; the openness theorem of GG-A4-Euclid) and the derived inventory supplying f ; what the framework does not supply is a depth, and the surveys have not supplied one either, so no lift is carried. What *is* specifically GG-6 is the second part: the framework once owned one first-principles local channel — a light $U(1)_B$ vector — and under the boundary-condition termination of GG-A3-BI6 it no longer does: no vector survives below the lock scale, and there is no long-range baryonic force at any strength. The record of the earlier estimate follows, kept because the exclusion it reached empirically is now also structural, and because an earlier version of this paper declared its amplitude a research target, which it is not:

G.5.1 Setting

Unique to GG-6 is the *exact local* $U(1)_B$ gauge factor, which forbids proton decay ($\tau_p = \infty$) and gives $\bar{\theta} \lesssim 4.7 \times 10^{-18}$ naturally. An earlier reading treated the gauge boson of $U(1)_B$, in the standard gauge-theory sense (Peskin and Schroeder, 1995), as an ultra-weakly coupled long-range vector (“graviphoton”) coupling to baryon number rather than to mass or charge, and asked what such a vector would do to the expansion history in regions of nonzero baryonic overdensity. Under the boundary-condition termination of GG-A3-BI6 the vector is eaten at the lock scale and there is no long-range force; what follows is the record of the estimate made under the earlier reading.

G.5.2 Effective local Friedmann equation

At any point in the universe one could write, following the legacy appendix of Arisaka (2026, A10-Cosmos) (kept there as record),

$$H_{\text{loc}}^2 = H_{\text{bg}}^2 + \frac{8\pi G_N}{3} \delta\rho_B^{\text{eff}}, \quad (74)$$

where $\delta\rho_B^{\text{eff}}$ is the effective baryonic correction from the $U(1)_B$ graviphoton pressure. The graviphoton force between baryonic overdensities, at distances small compared to its Compton wavelength (which is of cosmological scale), takes the form

$$F_{U(1)_B}(r) = \frac{g_B^2}{4\pi} \frac{q_{B,1} q_{B,2}}{r^2}, \quad g_B^2/4\pi \sim \frac{2}{5} G_N m_p^2, \quad (75)$$

with the coefficient $2/5$ forced by the same BI-6 entry that gives the Arisaka Ratio. The net gravitational-plus- $U(1)_B$ “effective Newton” for baryons is thus

$$G_{\text{eff}}^{(B)} = G_N (1 - \zeta_B), \quad \zeta_B \equiv \frac{2}{5} \delta\rho_B / \rho_B, \quad (76)$$

i.e. the graviphoton is *repulsive*, slightly inflating mildly underdense regions on baryonic overdensity contrast.

G.5.3 Magnitude estimate: the first-pass result and the sixty-fold gap

The local universe at the $\sim$ Gpc scale relevant to the distance ladder has a baryonic density contrast typically in the range $|\delta_B| \sim 0.05\text{--}0.08$ (from redshift surveys and peculiar-velocity reconstructions). Plugging into Eq. (74):

$$\frac{H_{\text{loc}} - H_{\text{bg}}}{H_{\text{bg}}} \simeq \frac{1}{2} \Omega_B \zeta_B \approx \frac{1}{2} (0.049) \frac{2}{5} |\delta_B| \approx 0.0098 |\delta_B|, \quad (77)$$

so the observed contrast yields the first-pass estimate

$$\epsilon_{\text{local}}^{(B), \text{first-pass}} \sim 5\text{--}8 \times 10^{-4}. \quad (78)$$

This first-pass estimate fell roughly an order of magnitude short of even a few-percent lift when it was made, and the channel is now closed structurally. The honest statement of record is therefore: *no first-principles channel of the framework reaches any part of the local excess*. The excess that separates the derived background $66.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$ from the observed local 73.5,

$$\epsilon_{\text{local}}^{(B)} = \frac{73.50}{66.8} - 1 = 0.100, \quad (79)$$

is carried whole as a distance-ladder calibration systematic the framework predicts. There is no environmental amplitude to derive from the $U(1)_B$ sector, and the research target an earlier version of this paper declared here is withdrawn.

G.5.4 Why this is not a large local void

The distinction between the GG-6 $U(1)_B$ graviphoton mechanism and the phenomenological “large local void” model is important. A pure void model requires $|\delta| \gtrsim 0.2$ at several hundred Mpc, which is in serious tension with observed large-scale isotropy. The environmental channel this paper prices is the ordinary Local-Hole monopole of §32, which needs a depth the surveys do not report and is therefore not carried. The graviphoton channel computed above is recorded for honesty, not invoked: its coupling was fixed (the same $2/5$ that appears in BI-6 as the coefficient of F_B^2), its first-pass magnitude (78) fell an order short of even a few-percent lift, no amplification mechanism for it exists in the theory, and under the termination there is no vector to carry it. It is a channel closed by the termination and, independently, excluded empirically, and the closure is part of the argument that the excess is not cosmological.

Status

Step one is classical linear theory; every number in it ($f = 0.533$, the 0.178 coefficient, the depth conversions) is checkable on one line. Step two is survey phenomenology with the audits cited. Step three carries the printed grades: the response law derived and the depth not supplied, so no lift carried; the $U(1)_B$ channel closed by the termination of GG-A3-BI6 (its earlier estimate, $5\text{--}8 \times 10^{-4}$, kept as record); the excess named-open. The master formula Eq. (73) is this course’s export: Part IV’s void pricing, Part V’s void-class verdicts, and Course H’s rarity pricing all read

from it.

H Derivation course: backreaction and averaging

“Maybe the average of an inhomogeneous universe simply expands differently” is the tension’s most seductive escape, because it sounds like it requires no new physics at all. This course does what the slogan never does: derives the averaging formalism from scratch, computes the effect on the derived inventory, and prices the escape at its exact target.

H.1 Buchert averaging, from the definition up

Following [Buchert \(2000\)](#), average any scalar f over a comoving spatial domain D with volume $V_D = \int_D \sqrt{h} d^3x$: $\langle f \rangle_D = V_D^{-1} \int_D f \sqrt{h} d^3x$. The volume element evolves as $\partial_t \sqrt{h} = \theta \sqrt{h}$, with θ the local expansion rate; differentiating the definition gives the *commutation rule* — the formalism’s one genuinely new object:

$$\partial_t \langle f \rangle_D - \langle \partial_t f \rangle_D = \langle f \theta \rangle_D - \langle f \rangle_D \langle \theta \rangle_D,$$

time derivative and average fail to commute by exactly the *covariance of f with the expansion*: faster-expanding regions grow their volumes and vote more heavily in tomorrow’s average. Apply this to Raychaudhuri’s equation for irrotational dust and to the Hamiltonian constraint, define the averaged scale factor by $a_D \propto V_D^{1/3}$, and the averaged equations close up to one new term,

$$3 \left(\frac{\dot{a}_D}{a_D} \right)^2 = 8\pi G \langle \rho \rangle_D - \frac{1}{2} \langle \mathcal{R} \rangle_D - \frac{1}{2} \mathcal{Q}_D, \quad \mathcal{Q}_D = \frac{2}{3} (\langle \theta^2 \rangle_D - \langle \theta \rangle_D^2) - 2 \langle \sigma^2 \rangle_D,$$

the kinematical backreaction: variance of expansion, minus shear (Eq. (42) of the main text). *Intuition*: $\mathcal{Q}_D$ is the statement that a mixture of fast and slow expanders does not average to the mean — but note the minus sign: shear, which inhomogeneity also produces, *fights* the variance. The two are comparable in linear theory, which is why the net effect will prove tiny.

H.2 The size of the terms on the derived inventory

Linear theory turns both terms into integrals over the power spectrum. From Course G’s master relation, the local expansion perturbation is $\delta\theta/3H = -(f/3)\delta$, so $\text{Var}(\theta) = H^2 f^2 \sigma_\delta^2(R)$, and the linear tidal shear evaluates to $\langle \sigma^2 \rangle = \frac{2}{15} H^2 f^2 \sigma_\delta^2$; hence

$$\frac{\mathcal{Q}_D}{H^2} = \left(\frac{2}{3} - \frac{4}{15} \right) f^2 \sigma_\delta^2 = 0.4 f^2 \sigma_\delta^2, \quad \frac{\delta H_D}{H} = -\frac{\mathcal{Q}_D}{12H^2},$$

with $\sigma_\delta(R)$ the top-hat density variance on the averaging scale. On the derived inventory ($\sigma_8 = 0.80$, $f = 0.533$) the released script’s linear machinery evaluates $\sigma_\delta = 0.163, 0.048, 0.026$ at $R = 40, 100, 150 \text{ Mpc}/h$, and the anatomy of the escape is over before it starts: at $R = 40$, $\mathcal{Q}_D/H^2 = 0.4 \times 0.284 \times 0.0266 = 3.0 \times 10^{-3}$ and the induced $|\delta H_D/H| \approx 3 \times 10^{-4}$ — *two hundredths of a percent*, and negative.

H.3 What averaging cannot buy

Three separate walls, in decreasing order of generality. *The systematic wall*: the averaged shift is $\mathcal{Q}_D/12H^2 \lesssim 10^{-3}$ on every relevant scale — three orders below the residual — because variance and shear nearly cancel and both are quadratic in small contrasts. *The stochastic wall*: one can be *lucky* rather than systematic — live in an unusual patch — and the price is computed from the same $\sigma_\delta(R)$: the $\bar{\delta} \approx -0.56$ that Course G showed the full excess requires is a 3.4σ fluctuation at 40 Mpc/h, but 12σ at 100 and 22σ at 150 — and the ladder’s supernovae live in the 100–150 shell. Excluded, in the precise sense that standard structure formation does not make such patches. *The loophole, priced*: the fully non-linear “timescape” reading rejects the perturbative frame entirely; the paper carries it as the class’s honest extreme (§44) rather than pretending the linear bound closes it — but notes that it must then also re-derive everything the linear frame gets right.

H.4 The record: the bound at the exact target

The compact record follows — the numbers behind Fig. 19, evaluated at the parent’s $\sigma_8 = 0.80$ and the whole +10% excess on the carried background, every row printed by the released script:

This course substantiates §35: that the GDOS-coarse-grained backreaction cannot supply the local excess. In Buchert’s spatial average of the inhomogeneous dust cosmology, the volume-averaged scale factor a_D obeys $3(\dot{a}_D/a_D)^2 = 8\pi G\langle\rho\rangle_D - \frac{1}{2}\langle R\rangle_D - \frac{1}{2}\mathcal{Q}_D$, with $\mathcal{Q}_D$ of Eq. (42). In the linear growing mode the local expansion perturbation is $\delta H/H = -(f/3)\delta$ (continuity), so $\text{Var}(\theta) = H^2 f^2 \sigma_\delta^2$ and the linear tidal shear is $\langle\sigma^2\rangle = \frac{2}{15}H^2 f^2 \sigma_\delta^2$; the induced shift of the averaged rate is $\delta H_D/H = -\mathcal{Q}_D/(12H^2)$.

Evaluating on the exact derived background ($\Omega_m = 1/\pi = 0.3183$, $\Omega_b = 1/2\pi^2$, $h = 0.6676$, $n_s = 0.9634$, $\sigma_8 = 0.80$ from the declared A_s , GG-A10-Cosmos; growth rate $f = \Omega_m^{0.55} = 0.533$) with an Eisenstein–Hu linear $P(k)$ (Eisenstein and Hu, 1998) normalized to $\sigma_8 = 0.80$:

- *Systematic backreaction*: $\langle\mathcal{Q}_D\rangle/H^2$ runs from $\sim 10^{-2}$ at $R = 20$ Mpc/h to $\sim 10^{-4}$ at 150 Mpc/h; the induced $|\delta H_D/H| \lesssim 0.1\%$, and negative.
- *Cosmic variance*: $\sigma_H/H = (f/3)\sigma_\delta(R)$ is 2.9% at 40 Mpc/h, 0.85% at 100, 0.29% at 200.
- *Rarity of a +10% void*: $\langle\delta\rangle \approx -0.56$ is 3.4σ at 40 Mpc/h but 12σ at 100 and 22σ at 150 — excluded on the ladder scale. (At the earlier inputs, $\sigma_8 = 0.737$ and a +7.3% target, the same rows read -0.41 and $2.7/9.3/17\sigma$.)

The open-source script that produces these numbers and Figure 19 (`c2_backreaction_pilocked.py`) is released with the paper, and recomputing at the exact $1/\pi$ inventory reproduces every quoted value at the stated precision. The conclusion is robust to the $O(1)$ shear coefficient: the systematic averaging bias is sub-percent and cannot close a ten-percent gap.

Status

Step one is the standard averaging formalism, derived rather than cited; step two is linear-theory quadrature on the derived inventory, every value reproduced by `c2_backreaction_pilocked.py`;

step three’s exclusions are as printed in the main text, with the timescape loophole carried open at its stated price. This course closes the environmental trilogy: Course G derived what a neighborhood can do, this course what an average cannot, and between them the residual’s non-cosmological character is not an assumption but an audit.

I Derivation course: the growth of structure and S_8

The inventory that sets the expansion rate also sets how fast matter clumps, so the growth of structure is the derived chart’s second, independent report card. Step one builds linear perturbation theory to the point where the two survey observables — σ_8 and S_8 — are defined by integrals a student can evaluate; step two states the S_8 tension and what the standard framework cannot say about it; step three reads the computed amplitude $S_8 = 0.825$ off the same ledger as everything else, with its one open flag displayed.

I.1 The perturbed universe: three equations

Write the metric with weak potentials (conformal-Newtonian gauge), $ds^2 = -(1 + 2\Phi)dt^2 + a^2(1 - 2\Phi)d\mathbf{x}^2$, and perturb the matter fluid: density contrast δ , peculiar velocity $\mathbf{v}$. Three classical statements close the system on sub-horizon scales. *Continuity* (mass conservation): $\dot{\delta} = -\nabla \cdot \mathbf{v}/a$ — convergence feeds density (Course G used exactly this). *Euler* (momentum): $\dot{\mathbf{v}} + H\mathbf{v} = -\nabla\Phi/a$ — gravity accelerates, expansion drags: the $H\mathbf{v}$ term is Hubble friction, the same brake that freezes fields in inflation and axion cosmology. *Poisson* (gravity): $\nabla^2\Phi = 4\pi G_N a^2 \bar{\rho}_m \delta$ — only the *contrast* gravitates. Differentiate continuity, insert Euler and Poisson, and the system collapses to one line:

$$\ddot{\delta} + 2H\dot{\delta} = 4\pi G_N \bar{\rho}_m \delta, \quad (80)$$

a damped, driven oscillator with gravity as negative stiffness: growth is a competition between self-attraction (right side) and Hubble friction (the $2H\dot{\delta}$).

I.2 The growth factor, era by era

Solve Eq. (80) in each era. *Matter domination* ($H^2 = \frac{8\pi G}{3}\bar{\rho}_m$, $a \propto t^{2/3}$): try $\delta \propto t^p$; the equation gives $p(p-1) + \frac{4}{3}p = \frac{2}{3}$, i.e. $p = 2/3$ or -1 : the growing mode is $\delta \propto t^{2/3} \propto a$ — contrast grows exactly as fast as the universe expands, no faster. *Radiation domination*: the driving term is negligible next to the expansion rate and growth stalls logarithmically — the Mészáros effect (Mészáros, 1974); structure marks time until equality ($z_{\text{eq}} = 3400$ on the derived chart), which is why ω_m , by setting equality, imprints itself on the power spectrum’s turnover. *Vacuum domination*: the driving term dilutes away, $\dot{\delta} \rightarrow \delta \times 0$: growth freezes. Collecting all eras, the growing mode has the closed integral form of Eq. (81) below, and its logarithmic rate today defines the growth rate $f \equiv d \ln D / d \ln a \simeq \Omega_m^{0.55} = 0.533$ — the number Courses G and H imported. *Intuition for the exponent*: 0.55 is not mysterious; it is the statement that growth responds to the *instantaneous* matter share slightly sub-linearly, because friction remembers the recent expansion history.

I.3 From modes to observables: σ_8 and S_8

The linear field is a Gaussian random process, specified completely by its power spectrum $P(k) = A k^{n_s} T^2(k) D^2(a)$: primordial tilt, transfer function (the frozen record of the radiation-era stall), growth. Surveys compress it to one number: the rms contrast in a sphere of radius $8 h^{-1} \text{Mpc}$,

$$\sigma_8^2 = \int_0^\infty \frac{dk k^2}{2\pi^2} P(k) W^2(kR_8), \quad W(x) = \frac{3(\sin x - x \cos x)}{x^3},$$

the top-hat window. Weak lensing measures the combination $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$, because lensing responds to *amount of matter times its lumpiness*; the 0.3 is a historical normalization. Everything in these integrals is either primordial (seed, tilt), thermal-history (transfer), or inventory (growth): there is no independent “clustering dial.”

I.4 What the standard framework cannot say about its own tension

The lensing generation of surveys reads $S_8 \approx 0.76$; the CMB, extrapolated through fitted ΛCDM , prefers ≈ 0.83 — a milder sibling of the Hubble tension, and the standard framework can only fit it, not interrogate it: with six dials, a low lensing amplitude can be bought by shifting several combinations, and the literature duly offers baryonic feedback, neutrino mass, intrinsic alignments, and new physics in any mixture. The structural blind spot is the same one Course D found in the ladder: a fitted framework cannot say what the amplitude *should* be, so it cannot say which measurement is reading truth. And the two tensions cannot be discussed jointly with integrity unless the same parameters move both — which, in a fit, they are free not to do.

I.5 The geometric reading: one ledger, both amplitudes

On the derived chart the growth sector inherits everything and chooses nothing, and the compact chain is worth preserving verbatim:

The same inventory that fixes H_0 fixes the clustering amplitude, and the chain has three steps, none free.

Step 1 — the growth factor. In the linear regime every Fourier mode of the matter field grows by one universal factor $D(a)$, governed by the competition between gravity and expansion; for a ΛCDM -like history the growing solution has the standard closed form ([Dodelson and Schmidt, 2020](#))

$$D(a) \propto \frac{5\Omega_m}{2} E(a) \int_0^a \frac{da'}{[a' E(a')]^3}, \quad E(a) = \sqrt{\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_\Lambda}, \quad (81)$$

with $E(a)$ the derived history of Eq. (85) — so the growth factor inherits the inventory with no new freedom.

Step 2 — the amplitude. The overall normalization traces back to the primordial seed: the declared $A_s = 2.10 \times 10^{-9}$ at $k = 0.05 \text{ Mpc}^{-1}$ (the launch sector’s one declared input), the *derived* tilt $n_s = 0.9634$, and the Eisenstein–Hu transfer function on the π -locked background at $h = 0.6676$, propagated through Eq. (81). That chain is carried by the parent GG-A10-Cosmos (its released `sigma8_gg.py`) and reported in three numbers so a reader can stop where they prefer: 0.823 with

massless neutrinos; 0.808 with the inherited $\Sigma m_\nu = 59$ meV; 0.801 with the fit’s own $+0.9\%$ bias against a Boltzmann integration, measured on Planck’s parameters, divided out. The codebook carries $\sigma_8 = 0.80 \pm 0.01$, and the backreaction course of App. H.4 uses the same value.

Step 3 — the lensing combination. Weak-lensing surveys report $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$; with the derived shares,

$$S_8 = 0.80 \times \sqrt{\frac{1/\pi}{0.3}} = 0.80 \times 1.030 = 0.825 \pm 0.010, \quad (82)$$

against the KiDS-1000 determination $0.759^{+0.024}_{-0.021} - +8.7\%, +2.5\sigma$ — and the CMB- Λ CDM 0.832 ± 0.013 (-0.4σ). The amplitude shares the S_8 tension (§39). Because the amplitude and the rate come from one inventory, they cannot be moved independently — the same structural signature that separates a derivation from a per-observable fit — and the tension is inherited, not claimed away. (An earlier version of this course printed $\sigma_8 = 0.737$ and $S_8 = 0.759$, attributed to the parent; the parent’s chain returns 0.80, and the earlier number had no generator.)

The course-level remarks. *First*, the coincidence to notice: the computed $S_8 = 0.825$ lands on the *CMB* side of the split, for the same structural reason the derived $H_0 = 66.8$ lands on the *CMB* side — both are read off the inventory, not fitted to either camp; the pattern “agree with the cleaner instrument, dissect the messier chain” repeats across the two tensions, and in the growth sector the messier chain is not this paper’s to dissect. *Second*, the joint rigidity: σ_8 , f , Ω_m , and H_0 move together or not at all, so the growth sector is a falsifier of the Hubble sector and vice versa — the redshift-space-distortion combination $f\sigma_8(z)$ is fixed at every redshift with no remaining freedom. *Third*, what is *not* open: an earlier version flagged the sign of a growth response as pending the parent’s audit (OQ-Hubble-5); on one background there is no growth response to have a sign, and the flag is closed.

Status

Step one is the standard linear theory, derived; step two is the field’s own accounting. Step three: $\sigma_8 = 0.80$ inherits the parent’s amplitude chain (seed A_s declared, tilt derived), $S_8 = 0.825$ is arithmetic on it, the tension with cosmic shear is shared and said so, and there is no growth sign to leave open. One ledger, one amplitude, every step auditable.

J Derivation course: standard sirens from the quadrupole formula up

The referee of this paper deserves its own course: the reader should see, derived and not asserted, why a binary inspiral is an absolute distance meter with no ladder anywhere in it. Step one builds the siren from linearized general relativity to the GW170817 numbers; step two derives the degeneracies and floors that limit it; step three re-derives the forecast algebra this paper stakes its verdict on.

J.1 Gravitational waves in one page

Linearize gravity: $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, $|h| \ll 1$ (full treatments: [Maggiore, 2008](#)). In the transverse-traceless gauge the vacuum Einstein equations reduce to the wave equation $\square h_{\mu\nu} = 0$: ripples of spacetime, two polarizations $(+, \times)$, propagating at c . Coupling to matter, the leading radiation is *quadrupole* — monopole radiation is forbidden by mass conservation, dipole by momentum conservation (the gravitational “charge” cannot oscillate against itself), so the first radiating moment is the quadrupole Q_{ij} , and the far-field strain is

$$h_{ij} \sim \frac{2G_N}{c^4} \frac{\ddot{Q}_{ij}}{d},$$

with the $1/d$ of any wave and the ferocious G/c^4 that makes spacetime the stiffest medium there is. *Dimensional intuition*: $G/c^4 \sim 10^{-44} \text{ s}^2\text{kg}^{-1}\text{m}^{-1}$ — it takes solar masses at relativistic speeds to make $h \sim 10^{-21}$, and that stiffness is precisely why the source physics is clean: nothing else in the universe can fake it.

J.2 The chirp: one mass from the phasing

Specialize to a circular binary, masses m_1, m_2 , separation a , orbital frequency ω (Kepler: $\omega^2 a^3 = GM$, $M = m_1 + m_2$). The quadrupole moment rotates at ω , so the wave comes out at $f = 2f_{\text{orb}}$; the radiated power is the Einstein quadrupole luminosity $\mathcal{L} \propto (G^4/c^5) \mu^2 M^3/a^5$, and energy balance — the orbit paying the radiation bill from its binding energy $E = -G\mu M/2a$ — turns loss into acceleration. Carrying the algebra through (each substitution is Kepler or the two lines above), every mass dependence collapses into one combination:

$$\dot{f} = \frac{96}{5} \pi^{8/3} \left(\frac{G_N \mathcal{M}_c}{c^3} \right)^{5/3} f^{11/3}, \quad \mathcal{M}_c \equiv \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \quad (83)$$

the *chirp mass* — the observable [Schutz \(1986\)](#) identified as cosmology’s gift. Read the miracle carefully: f and $\dot{f}$ are direct observables of the phasing, so Eq. (83) hands over $\mathcal{M}_c$ with no astronomy at all — for two $1.36 M_\odot$ neutron stars, $\mathcal{M}_c = 1.18 M_\odot$, and GW170817’s phasing measured exactly this to four digits.

J.3 The amplitude: an absolute distance

Insert the binary’s quadrupole into the strain formula; the same chirp mass appears again,

$$h \sim \frac{4}{d_L} \frac{(G_N \mathcal{M}_c)^{5/3}}{c^4} (\pi f)^{2/3},$$

and the logic closes like a trap: the phasing has *already* fixed $\mathcal{M}_c$ and f is measured, so the only unknown in the amplitude is d_L . The source’s absolute wattage is certified by the field equations — this is what “standard siren” means, and why Course D’s entire ladder of calibrations simply does not appear. (Cosmology enters only as the redshifting $\mathcal{M}_c \rightarrow (1+z)\mathcal{M}_c$, which is why the waveform alone cannot supply z and a host galaxy must.)

J.4 Worked example: GW170817, end to end

The one bright siren, in four measured numbers. The phasing gave $\mathcal{M}_c = 1.188 M_\odot$; the amplitude then gave $d_L \approx 40$ Mpc; the host, NGC 4993, supplied the Hubble-flow velocity 3017 ± 166 km s⁻¹; and the quotient, with the inclination degeneracy marginalized, gave $H_0 = 70_{-8}^{+12}$ (central value 3017/43.1; the skewed error is the inclination’s signature) — fifteen-percent precision from one event, straddling both camps, exactly as Part VII tells it. Nothing in the chain touched a Cepheid, a supernova standardization, or a sound horizon.

J.5 What limits the referee

The limitations step, derived rather than listed. *Inclination–distance*: the two polarizations carry $h_+ \propto (1 + \cos^2 \iota)/2$ and $h_\times \propto \cos \iota$; a face-on binary is louder, so a nearby edge-on and a distant face-on system produce similar strain — the d – ι degeneracy behind GW170817’s asymmetric error, broken partly by the polarization ratio and decisively by jet-afterglow imaging. *The calibration floor*: strain is measured against laboratory photon calibrators, currently at the one-to-two-percent level — the one *common-mode* term in the siren budget (Course D taught why common modes matter), an engineering quantity with no coupling to either discordant route. *Dark-siren statistics*: without a counterpart the redshift comes from weighting every galaxy in the localization volume; catalog incompleteness and the assumed black-hole mass spectrum are the floors, visible in the wild as the $76.6 \rightarrow 71.0 \rightarrow 69.9$ walk of the report card (§66). All three limits shrink with catalog size or engineering; none couples to the tension’s contested chains — which is the property that makes the siren the referee.

J.6 The geometric stake: the forecast algebra, re-derived

Step three is this paper’s own arithmetic, now derivable by the reader in two lines. N unbiased events of fractional error σ_{ev} combine to $\sigma(N) = \sigma_{\text{ev}} \times 70/\sqrt{N}$; demanding k -sigma separation of the two camps ($\Delta = 73.5 - 66.8 = 6.7$) gives

$$N_{k\sigma} = \left(\frac{70 \sigma_{\text{ev}} k}{\Delta} \right)^2 \implies N_{3\sigma} = 22, \quad N_{5\sigma} = 61$$

at the GW170817-like $\sigma_{\text{ev}} = 15\%$ — the thresholds of Part VII, hostage to no rate assumption. The three-worlds pricing then follows from the same Gaussian algebra: a converged catalog at 70.0, read between the verdict thresholds ($\sigma_{\text{cat}} = 2.23 \rightarrow 1.34$), sits 1.4 – 2.4σ from the derived 66.8 ± 0.2 (§70) — every number in this paragraph recomputable from this course’s two boxed formulas and a table of quadratures. The prediction itself is unchanged since Part III: $H_0^{\text{GW}} \rightarrow 67$, not 73, with both outcomes’ meanings printed in advance.

Status

Steps one and two are classical relativity and instrument physics; the derivations here are the standard ones, compressed but complete in their logic, and every number matches the released `c2_siren_forecast.py`. Step three carries the paper’s grades unchanged: the forecast thresholds

are arithmetic; the dates are empirical (the report card); the prediction is the pre-registered F-Hubble-2. This course exists so that the reader who will watch the verdict arrive can check, personally, that the instrument owes nothing to either side.

K Derivation course: time, age, and the curve-reading channels

The tension is about one point on a curve; this closing course is about the curve as a function of *time* — the ages, the lookbacks, and the three channels that read $H(z)$ directly. Step one derives the clock integrals in closed form; step two shows what a one-number cosmology cannot do; step three states what the derived chronology commits to, row by row.

K.1 Cosmic time, derived in closed form

Course C quoted the age formula; here is its derivation, worth doing once because the same substitution recurs everywhere. From the Friedmann equation in the matter-plus-vacuum era,

$$t(a) = \int_0^a \frac{da'}{a' H(a')} = \frac{1}{H_0} \int_0^a \frac{a'^{1/2} da'}{\sqrt{\Omega_m + \Omega_\Lambda a'^3}}.$$

Substitute $u = a'^{3/2}$, $du = \frac{3}{2} a'^{1/2} da'$: the integral becomes elementary,

$$t(a) = \frac{2}{3H_0} \int_0^{a^{3/2}} \frac{du}{\sqrt{\Omega_m + \Omega_\Lambda u^2}} = \frac{2}{3H_0 \sqrt{\Omega_\Lambda}} \operatorname{arcsinh} \left(\sqrt{\frac{\Omega_\Lambda}{\Omega_m}} a^{3/2} \right),$$

Eq. (59) of Course C, now earned rather than quoted. Three limits deserve their intuitions: at early times $\operatorname{arcsinh} x \rightarrow x$ and $t \propto a^{3/2}$ (matter-era clock); at late times $\operatorname{arcsinh} x \rightarrow \ln 2x$ and $a \propto e^{H_\infty t}$ (the de Sitter clock); and at $a = 1$ the formula delivers $t_0 = 13.88$ Gyr. The lookback ladder on the derived chart: at $z = 1$ the clock read 5.87 Gyr (lookback 8.01); the acceleration — the discovery of [Riess et al. \(1998\)](#) and [Perlmutter et al. \(1999\)](#) — began at $t = 7.8$ Gyr ($z = 0.62$, Course C's exact $1/\sqrt{2}$); at $z = 2$ the clock read 3.28 Gyr. Every one of these is a falsifiable statement about stellar populations: no object at $z = 2$ may be older than 3.3 Gyr on this chart.

K.2 Chronometers and the drift, from first principles

Two channels read the curve with no ruler and no candle, and their master relations take four lines each. *Cosmic chronometers* ([Jimenez and Loeb, 2002](#); [Moresco et al., 2022](#)): from $1 + z = a_0/a(t_e)$, differentiate with respect to the *emission* time of a passively evolving galaxy population: $dz/dt_e = -(1+z)\dot{a}/a = -(1+z)H(z)$, so $H(z) = -[1/(1+z)] dz/dt_e$ — Eq. (47): measure the differential age of old stellar populations across a small redshift step and the curve appears, no cosmology assumed. *The redshift drift* ([Loeb, 1998](#); [Sandage, 1962](#)): watch a single comoving source as the *observer's* clock advances by Δt_0 . Both endpoints of $1 + z = a(t_0)/a(t_e)$ move, and the chain rule gives $\dot{z} = (1+z)H_0 - H(z)$ — Eq. (48): the competition between today's stretch of the whole interval and the source-epoch stretch of the emitted wave. Convert to velocity, $\Delta v = c \Delta z/(1+z)$, and the derived chart commits to numbers: +2.4, +2.1, −0.3, −5.6 cm s^{−1} per

decade at $z = 0.5, 1, 2, 4$, with the sign change at $z \approx 1.9$ — the redshift at which $H(z)/(1+z)$ crosses today's rate, i.e. where the universe's deceleration-era imprint balances its acceleration-era one. A drift map is cosmology with no proxy at all; on the derived chart it is also cosmology with no freedom at all.

K.3 What a one-number cosmology cannot do

The limitations step, and the closing statement of the whole appendix chain. A cosmology organized around measuring H_0 ever better can, at saturation, deliver exactly one datum — and the tension is the proof that one datum cannot arbitrate its own systematics. The curve is different in kind: ages bound it at every redshift where an old object is found; chronometers sample its derivative; BAO points, calibrated by the derived $r_s = 144.7$ Mpc, read it absolutely; the drift will watch it move. A fitted model meets each new curve-point by re-fitting; a derived curve meets it with a number already printed — which is why Table 18 is not a summary but an exposure schedule.

K.4 The geometric chronology, as commitments

Step three is a recital, because by now every entry has been derived in some course of this appendix: $t_0 = 13.88$ Gyr (this course); $t_{\text{acc}} = 7.8$ Gyr at the exact $\text{arcsinh}(1/\sqrt{2})$ (Course C); $z_* = e^7$ and the two rulers $r_s = 144.7 / r_d \approx 147$ Mpc (Course E); $H_\infty = 55.2$ and the 12.3 Gyr doubling time (Course C); the $H(z)$ rows and drift values above; and the five clocks — t , a , z , T , and the radial coordinate r — kept synchronized by the one ledger (the annex below walks the T clock). Each is graded as its source course grades it; each is a place to die; and together they are what this paper means when it says the constant was never the point — the *history* is the object, and the history is now fully on the record.

K.5 Annex: the thermal clock — scaling laws, thermal history, and freeze-out

The Friedmann equation of §C.10 runs the universe forward; the ledger tells it what to carry. This course collects the standard results on which every cosmological milestone quoted in the body rests, read through the GG-6/GDOS source ledger.

Exact scaling law. For a component with constant $w_i = p_i/\rho_i$, the continuity equation $\dot{\rho}_i + 3H(1+w_i)\rho_i = 0$ integrates immediately to

$$\rho_i(a) = \rho_{i,0} a^{-3(1+w_i)}, \quad (84)$$

so radiation ($w = \frac{1}{3}$) dilutes as a^{-4} , non-relativistic matter ($w = 0$) as a^{-3} , and the vacuum ($w = -1$, the derived value of GG-A10-Cosmos) stays constant. The first Friedmann equation therefore becomes a sum over the inventory,

$$H(a) = H_0 \sqrt{\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_\Lambda}, \quad \Omega_m = \frac{1}{\pi}, \quad \Omega_\Lambda = 1 - \frac{1}{\pi}, \quad \Omega_r \simeq 9.32 \times 10^{-5}, \quad (85)$$

the curve the body calls $H(t)$, with every coefficient fixed by the ledger and only the radiation share carried at its measured photon-plus-neutrino value.

Matter–radiation equality. Radiation and matter scale as a^{-4} and a^{-3} , so they cross at $a_{\text{eq}} = \Omega_r/\Omega_m$, i.e.

$$1 + z_{\text{eq}} = \frac{\Omega_m}{\Omega_r} = \frac{1/\pi}{9.32 \times 10^{-5}} \simeq 3.4 \times 10^3, \quad (86)$$

the epoch after which structure can grow; the derived $\Omega_m = 1/\pi$ fixes it with no free parameter.

Cooling and the temperature–time dictionary. Comoving entropy conservation, $g_*(T) T^3 a^3 = \text{const}$, gives the cooling law $T \propto g_*^{-1/3} a^{-1}$. In the radiation era $H^2 = \frac{8\pi G_N}{3} \rho_r$ with $\rho_r = \frac{\pi^2}{30} g_*(T) T^4$, and $t = 1/2H$, so

$$t \simeq 2.42 g_*(T)^{-1/2} \left(\frac{T}{1 \text{ MeV}} \right)^{-2} \text{ s}, \quad (87)$$

the dictionary that dates every early milestone. A species freezes out when its interaction rate falls below expansion, $\Gamma(T) \sim H(T)$ — the universal competition that sets every relic abundance, including the dark-matter density that the GG-C1-Kaxion Kaxion misalignment fixes to the ledger value.

Recombination as the switch. At $z_* = e^7 \approx 1097$ (Prop. 31.1) the plasma recombines and the microwave background is released; this is the epoch whose sound horizon calibrates the early-side distance scale of the next course. Everything to the past of it is read through the acoustic dictionary, everything to its future by the late probes — and the one late probe with calibrated rungs is the one that disagrees.

The fuller treatment of everything in this annex — entropy conservation, g_* bookkeeping, and the freeze-out calculus — is the classic text of Kolb & Turner (1990) and the thermal-history course of the parent’s appendices (Arisaka, 2026, A10-Cosmos).

Status

Every closed form in this course is derived above it; every number is asserted in the build script against the arcsinh forms and the drift relation; the grades are inherited unchanged. With this course, the appendix chain covers the full arc of the paper: background (C), measurement (D, E), vacuum (F), environment (G, H), growth (I), referee (J), and time (K).

L Derivation course: the statistics of a tension

The last course is about the tension itself as a statistical object — what a “ 7σ discrepancy” does and does not mean, why a decade of improving measurements made it worse rather than better, and what combination rules honestly do when correlations are unknown. Nothing in this course is cosmology; all of it is the bookkeeping that the cosmology of Courses C–K gets scored by, and it is taught last because every example is now familiar.

L.1 What a sigma means when systematics exist

Two measurements of one quantity, $x_1 \pm \sigma_1$ and $x_2 \pm \sigma_2$, are in tension at

$$T = \frac{|x_1 - x_2|}{\sqrt{\sigma_1^2 + \sigma_2^2}}$$

— *if* the quoted errors are complete and independent. Every word of that proviso works hard. If a systematic b biases x_1 and is absent from its error bar, the numerator is really $|\Delta_{\text{true}} + b|$ and T measures the bias as confidently as it would measure physics; if a systematic is *shared* by both measurements, it cancels in the difference and T never sees it at all — both values are wrong together, consistently (Course D’s common-mode theorem, met again from the other side). The practical moral for the tension: a large T proves the two chains disagree; it cannot, even in principle, say *which* chain, or whether the disagreement is nature or bookkeeping. That adjudication needs structure from outside the two chains — a third instrument, or a derived expectation.

L.2 Why precision amplifies bias: the decade’s paradox, derived

Now the one-line theorem that explains the last ten years. Hold any unresolved bias b fixed and let both teams improve, $\sigma \rightarrow 0$:

$$T = \frac{|\Delta_{\text{true}} + b|}{\sqrt{\sigma_1^2 + \sigma_2^2}} \xrightarrow{\sigma \rightarrow 0} \infty.$$

Precision does not resolve a bias; it *amplifies its significance*. If the true backgrounds agree ($\Delta_{\text{true}} = 0$) and a percent-level calibration bias exists anywhere, then better measurements must produce an ever-growing sigma count — which is exactly the observed history: the tension climbed from 2σ to 7σ precisely as both chains reached percent-class precision. The sigma count, by itself, cannot distinguish this world from one with genuinely discordant physics; the *growth pattern* of T with precision is identical in both. This is the statistical face of the statement the paper makes throughout: the dispute over the last digit cannot settle what kind of number is being measured, because $T \propto b/\sigma$ rewards the digits indiscriminately.

L.3 The bracket theorem: combining channels with unknown correlations

Course D derived the channel-combination formula $\sigma_{\text{tot}}^2 = \sum_i \sigma_i^2 + 2 \sum_{i < j} \rho_{ij} \sigma_i \sigma_j$; here is the statement that makes the paper’s two printed scenarios rigorous rather than rhetorical. For any set of non-negative channel sizes, σ_{tot} is monotonically increasing in every ρ_{ij} , so over the whole physically allowed family $0 \leq \rho_{ij} \leq 1$ — the restriction to *non-negative* correlations being itself an assumption we argue rather than smuggle: these channels share calibrations, compilations, and anchors, mechanisms that push *together*, and no proposed mechanism anticorrelates them; under a negative correlation σ_{tot} would drop below quadrature and the bracket’s floor would only strengthen

$$\underbrace{\left(\sum_i \sigma_i^2\right)^{1/2}}_{\text{quadrature}} \leq \sigma_{\text{tot}} \leq \underbrace{\sum_i \sigma_i}_{\text{linear stack}} :$$

the two rules the paper prints are not two opinions; they are the *sharp bounds* of the correlation-agnostic problem, attained at $\rho = 0$ and $\rho = 1$. Quoting any single interior value would smuggle in a correlation matrix nobody has measured; quoting the bounds is the maximal honest statement — and the observation that the required +10% falls above the upper bound — the maximal honest statement: the published channels do not yet suffice (§34.5; effective $\bar{\rho} \approx 0.5\text{--}0.7$, Course D) is therefore a consistency, not a construction.

L.4 What a covariance-weighted consensus certifies

The optimal combination of correlated readings x_i with covariance C is the generalized least-squares mean, $\hat{y} = (\mathbf{1}^\top C^{-1} \mathbf{1})^{-1} \mathbf{1}^\top C^{-1} \mathbf{x}$, with variance $(\mathbf{1}^\top C^{-1} \mathbf{1})^{-1}$ — the machinery behind every modern consensus value, and it is genuinely optimal for what it models. Now add a common mode: $\mathbf{x} \rightarrow \mathbf{x} + c \mathbf{1}$. Because the weights sum to one, $\hat{y} \rightarrow \hat{y} + c$ — while C , the internal residuals, and every leave-one-out statistic are untouched. In one line: a consensus certifies the *covariance it models*, and transmits unmodeled common modes at full strength with undiminished confidence. This is not a criticism of any collaboration; it is a property of weighted means, and it delimits precisely what 73.50 ± 0.81 establishes: the stability of the network’s shared frame — Part II’s genuine achievement — and nothing about the frame itself. The boundary of the certification is where Part IV’s channels live.

L.5 Reading a landscape: bands, not rows

The forty-five values of Part V invite naive statistics — “average them all” — and the course closes its textbook step by deriving why that is wrong twice over. First, the values are not independent: methods share calibrations, catalogs, and even authors, so the effective number of independent readings is far smaller than forty-five, and an unweighted mean double-counts every shared systematic. Second, the field’s own meta-analysis finds the literature’s scatter exceeds its quoted errors — direct evidence of unmodeled (and partially shared) systematics — so *clustering*, not averaging, is the information-bearing statistic: the eye should weight bands, not count rows, which is exactly how Part V instructs its reader. A band structure with a filling middle is the signature of a population of measurements sharing a few large common modes that are being found and removed one at a time — and that reading is testable, band by band, as the corrections land.

L.6 The geometric posture: statistics with the freedom removed

Step three is the paper’s own statistical conduct, now assembled in one place. *Pre-registration*: every decision rule of the falsifier registry was fixed numerically before its datum (§56), so the trials factor of each confrontation is exactly one — the look-elsewhere tax that haunts post-hoc tension statistics is paid in advance, a discipline the cohort’s dark-matter companion has already executed live (Arisaka, 2026, C1-Kaxion). *Bracket pricing*: unknown correlations are carried as the sharp

bounds of §L.3, never as a chosen interior value. *Deflation, stated symmetrically*: when the field’s own published corrections are applied at face value, the significance of the residual tension deflates from the headline 7σ class to the 2σ class (§34.5) — and the paper prints that arithmetic while explicitly *not* adopting the corrections, because the same bracket honesty cuts both ways. *Released arithmetic*: all of it — chains, brackets, deflations — regenerates from `c2_residual_arithmetic.py`’s twenty-three asserted checks (§M.2). And beneath all four habits, the structural point this course has been building toward: statistics can referee the tension only when something in the problem is *not* free. A derived background is exactly that fixed point — it converts the sigma count from an amplifier of unknown bias into a test with a printed answer.

Status

Everything in this course is elementary probability, derived in place; its cosmological content is zero and its bookkeeping content is the whole paper. The reader who retains one line should retain the theorem of §L.2: $T \propto b/\sigma$ — precision amplifies unresolved bias — and its corollary, which is this paper’s methodological thesis: the cure for a tension is never another digit; it is a fixed point outside both chains.

M Reproducibility and the released pipeline

How the decisive observation will be reproduced, and how every number in this report can be regenerated from the released scripts: the input tables with their justifications, the scripts index, the siren recipe, and the H0DN variant ladder.

M.1 Reproducibility: inputs, justifications, and the systematic budget

The derived rate is reproducible in the strict sense: every input is a published number or a ledger-fixed constant, and the propagation is elementary. This course collects the inputs, states where each comes from, and assembles the systematic budget on the background rate, so that an independent worker can regenerate the two-road $H_0^{(0)} = 66.6$ and bound its uncertainty (the carried 66.8 is Route I’s, and needs no astronomical input at all).

The inventory route, step by step. The first Friedmann equation at the present epoch reads $H_0^2 = (8\pi G_N/3)\rho_{\text{crit}}$, and each measured *physical* density $\omega_i = \Omega_i h^2$ is the product of the derived share and h^2 . Dividing the measured ω_i by the derived Ω_i therefore returns h^2 with no free parameter:

$$h = \sqrt{\frac{\omega_m}{\Omega_m}} = \sqrt{\frac{0.1430}{0.31831}} = 0.670 \Rightarrow H_0 = 67.0, \quad h = \sqrt{\frac{\omega_b}{\Omega_b}} = \sqrt{\frac{0.02237}{0.05066}} = 0.664 \Rightarrow H_0 = 66.5. \quad (88)$$

The two anchors bracket the two-road background $H_0^{(0)} = 66.6$, and their spread (0.5) is the leading systematic. Propagating the Planck error on ω_m (0.8%) gives ± 0.17 ; the two-route consistency is the dominant check, not the formal error.

Table 26: Inputs to the derived background rate, their provenance, and their contribution to the systematic budget. Ledger-fixed shares carry no measurement error; the uncertainty is entirely in the measured physical densities and the route-to-route spread. The last column is the one that matters for reproducing the result. The two ledger-fixed shares are exact, so everything the error bar contains enters through the two measured densities and the structural spread on the seesaw coefficient. A reader who wants to redo the calculation needs no input beyond this table.

Input	Value	Source	Contribution to $\sigma(H_0)$
Ω_m	$1/\pi = 0.31831$	GG-A10-Cosmos (ledger, exact)	0 (no measurement error)
Ω_b	$1/2\pi^2 = 0.05066$	GG-A10-Cosmos (ledger, exact)	0
$\omega_m = \Omega_m h^2$	0.1430 ± 0.0011	Planck 2018	± 0.26 (matter route)
$\omega_b = \Omega_b h^2$	0.02237 ± 0.00015	Planck 2018	± 0.22 (baryon route)
ζ_Λ	$\pi/175 = 0.01795$	codebook selection (§24.2)	± 0.2 structural (geometric route)
M_{lock}	3.85 TeV	GG-A5-GDOS T8.7 (theorem)	folded into ζ_Λ

The seesaw route, as the geometric determination. The vacuum-energy seesaw of §F.4 gives $H_0 = 498.3\sqrt{\zeta_\Lambda}$ with ζ_Λ selected by the codebook as $\pi/(LK) = \pi/175$ (§24.2), returning $H_0 = 66.8 \pm 0.2$ with no astronomical input. The $H_0 \propto M_{\text{lock}}^4$ sensitivity is pinned in the laboratory: the same lock scale predicts the electroweak vev at +0.045%, i.e. $\pm 0.18\%$ on the rate. That two structurally unrelated routes (a vacuum seesaw and a density ratio) land within 0.5 of one another, and both on the CMB value, is the reproducible core of the claim.

What an independent worker must supply. Only three continuous inputs enter — ω_m and ω_b (for the confirmation route only) and the lock scale M_{lock} — each published; the geometric route consumes nothing astronomical, and everything else is ledger-fixed. The computation is a pocket-calculator exercise, deliberately: the content is in the derivation of the shares Ω_i (GG-A10-Cosmos) and the seesaw normalization (GG-A5-GDOS), not in the arithmetic that turns them into a rate. This is the sense in which the background rate is not fitted — there is nothing to fit.

M.2 Scripts released with this paper

Every quantitative figure and forecast in this report regenerates from a short, documented, standalone script; Table 27 is the index. Each script states its inputs in its header, asserts its own arithmetic before writing anything, and runs in seconds on a laptop. The fastest way to referee this paper is to run them.

M.3 Reproducing the decisive test: the standard-siren H_0

The decisive observation (§57) is reproducible in principle by any group with a gravitational-wave catalog. A standard siren yields H_0 from two independently measured quantities: the luminosity distance d_L , read from the amplitude and chirp of the waveform, and the redshift z of the host (bright sirens with an electromagnetic counterpart) or of the redshift distribution of candidate hosts (dark sirens). With a catalog of N events the combined constraint scales as $\sigma_{H_0}/H_0 \sim (0.1\text{--}0.15)/\sqrt{N}$ for well-localized bright sirens, so a few tens of events reach the $\lesssim 1\%$ needed to separate 67 from 73.

Table 27: The open-source scripts released with this paper and what each reproduces. Each script asserts its own printed numbers, so a script that runs without failing has reproduced the figure or table it names rather than merely produced one. The last two entries do this most strictly: the residual arithmetic asserts every number in its section, and the landscape script asserts all forty-five audited values against the primary-record audit. No figure in this paper is drawn from a stored data file that is not regenerated here.

Script	Reproduces
<code>c2_fig_H_of_t_pilocked.py</code>	the derived expansion history (Fig. 13)
<code>c2_fig_a_of_t.py</code>	the derived scale factor (Fig. 11)
<code>c2_fig_omega_budget.py</code>	the running budget $\Omega_i(t)$ (Fig. 10)
<code>c2_fig_landmarks.py</code>	the landmarks walk and its numbers (Fig. 14, Table 6)
<code>c2_showcase_offsets.py</code>	the coherent-offset chart (Fig. 16)
<code>c2_fig_two_routes.py</code>	the two-route convergence chart, with the chain re-derived and asserted (Fig. 12)
<code>c2_fig_h0_history.py</code>	the century-of- H_0 chart (Fig. 4)
<code>c2_fig_hole.py</code>	the Local-Hole profiles and the lift each would give (Fig. 17)
<code>c2_fig_bridge.py</code>	the reconciliation bridge (Fig. 18)
<code>c2_backreaction_pilocked.py</code>	the backreaction and cosmic-variance bounds (Fig. 19; App. H.4)
<code>c2_siren_forecast.py</code>	the events-to-verdict forecast (Fig. 22, Table 20)
<code>c2_residual_arithmetic.py</code>	The two-scenario channel arithmetic, central-value chains, and significance deflation of §34.5; every printed number asserted.
<code>c2_showcase_landscape.py</code>	The 2026 H_0 landscape figure (Figure 20): all 45 audited values asserted in-script against the primary-record audit of July 12, 2026.

No sound horizon and no distance ladder enters at any step, which is precisely why the channel is decisive: its systematics are disjoint from both discordant methods. The pre-declared verdict is fixed in advance (F-Hubble-2): a converged $H_0^{\text{GW}} = 73 \pm 1$ refutes; 67 ± 1 confirms. The full forecast — events-to-verdict arithmetic, systematic budget, and the simulated catalog — is Part VII, and its open-source script (`c2_siren_forecast.py`) is released with the paper.

M.4 The H_0 DN variant ladder: the reproducibility table

All 22 H_0 DN variants lie within $\pm 1\sigma$ of the baseline 73.50; the largest deviations are V08 (no Cepheids, -0.76σ) and V99 (everything, $+0.70\sigma$), and the median offset is $\approx 0.2\sigma$. This is the stability pattern expected if every variant measures the same underlying local rate — no single method, anchor, or instrument drives it. The derived reading interprets that stable local value as the derived background 66.8 read high by a distance-ladder calibration systematic (Section 30).

N Registries, audits, and roadmaps

N.1 Internal source ledger: where each object lives

For the reader navigating between the main text and these courses, the load-bearing objects and their home addresses:

Table 28: Full list of H0DN baseline and variant values from [Casertano et al. \(2026\)](#). The derived reading takes this stable network consensus at face value as the measured local rate; its interpretation — background 66.8, read high by a predicted ladder systematic, plus a named open calibration residual — is Section 30 (Appendix G.5).

Variant	H_0 [km s ⁻¹ Mpc ⁻¹]	Description	Deviation from baseline
V00	73.50 ± 0.81	Baseline	0.0σ
V01	73.49 ± 0.81	+JAGB	-0.01σ
V02	73.51 ± 0.81	+Miras	$+0.01\sigma$
V03	73.60 ± 0.80	+DESI FP	$+0.13\sigma$
V04	73.52 ± 0.81	+SNe II (SCM)	$+0.02\sigma$
V05	73.72 ± 0.74	+SNe II (EPM)	$+0.30\sigma$
V07	73.31 ± 0.79	+SMC anchor	-0.24σ
V08	72.51 ± 1.30	no Cepheids	-0.76σ
V09	73.53 ± 0.88	no TRGB/SBF	$+0.03\sigma$
V10	73.61 ± 0.93	no Gaia/MW	$+0.12\sigma$
V11	73.34 ± 0.85	no DEB/LMC	-0.19σ
V12	73.08 ± 0.92	no NGC 4258	-0.46σ
V13	73.43 ± 1.80	no SNe Ia	-0.04σ
V14	73.33 ± 0.83	no SBF	-0.20σ
V15	73.63 ± 0.84	no masers (HF)	$+0.16\sigma$
V16	73.75 ± 1.33	no HST/LMC/MW	$+0.19\sigma$
V17	73.00 ± 0.86	no JWST	-0.58σ
V18	73.69 ± 0.82	no pre-1994 SNe	$+0.23\sigma$
V99	73.99 ± 0.70	everything	$+0.70\sigma$
V99a	73.66 ± 0.71	everything except TF	$+0.23\sigma$
O1	73.11 ± 0.92	orthogonal path 1	-0.42σ
O2	73.43 ± 1.80	orthogonal path 2	-0.04σ

Table 29: The internal source ledger: each headline object’s home in the main text, and the appendix course that teaches it. The ledger exists because the main text and the appendices were written to be read in either order, and a reader who enters through a course needs to know where its result is used. Reading the second column down is also a fair table of contents for the argument: the background first, then the local corrections, then the audits, then what is still open.

Object	Home in the main text	Course
Derived background 66.8 ± 0.2 (two-road 66.6)	§24 (Routes I and II)	E, F
The closed form for H_0	§24.3	F
The switch and the two dictionaries	§31	C, E
The local void, priced and not carried	§32, §33	G
The residual anatomy and verdict	§34, §34.5	D, L
The backreaction bound	§35	H
The 2026 landscape (45 values)	Part V (Fig. 20, Table 13)	L
The falsifier registry and calendar	§56, §58	—
The siren forecast and stake	Part VII	J
Growth and S_8	§39	I
The chronology and the drift	§59	K

N.2 The falsification catalog, compact

The complete falsification surface of this paper, in three layers, each detailed elsewhere and indexed here. *The registry*: F-Hubble-1 through F-Hubble-11, four tiers, each with a pre-declared numerical rule and its confirmation twin (Table 17); the channel-level trio is also carried cohort-wide as global ledger rows F26–F28. *The calendar*: the dated appointments and the commitments ledger (§58.1–§58.2). *The courses*: every derivation course above ends by naming its own falsifiable outputs — the $H(z)$ rows and ages (C, K), the shell profile (G), the growth rigidity $f\sigma_8(z)$ (I), the siren thresholds (J) — so the appendix chain is itself part of the paper’s exposure, not commentary on it.

What the reading inherits, the theorems it stands on, and how its open entries would be promoted.

N.3 Cross-paper inheritance table

This report derives nothing new about the cosmology; it reads the H0DN measurement through results imported from the cohort. Table 30 lists exactly what is inherited, from where, and what this paper does with it.

Table 30: What the derived Hubble reading inherits from the cohort. Each entry is imported at the grade its home paper assigns; this report adds only the reading of the H0DN measurement against them. Nothing in the first column is derived here, and that is the reason for listing it. This paper’s own contribution is the reading in the third column and nothing else, so a reader who rejects an inherited object should follow it back to the paper named beside it rather than argue with this one. Two of the entries are conditional at home and stay conditional here.

Inherited object	From	Use here
π -locked inventory ($\Omega_m = 1/\pi, \dots$)	GG-A10-Cosmos	Divided into CMB densities to give $H_0^{(0)}$
Hopf-fiber identification (Conjecture VIII)	GG-A10-Cosmos	The measure lemma
Two-route H_0 (66.8 geometry; 66.6 densities)	GG-A10-Cosmos	the D* grade rests on
Vacuum-energy seesaw, ζ_Λ	GG-A10-Cosmos/GG-A5-GDOS	The background the
Lock scale $M_{\text{lock}} = 3.85 \text{ TeV}$ (T8.7)	GG-A5-GDOS	tension splits around
OGN-5 integer ledger (1, 3, 5; 7, 35)	GG-A4-Euclid/GG-A2-GG6	Route I: the geometric
GDOS open-system boundary law	GG-A5-GDOS	determination (§24.2)
Derived radiation null ($\Delta N_{\text{eff}} \leq 0.03$)	GG-A10-Cosmos	Sets the seesaw nor-
Exact local $U(1)_B$ (proton stability)	GG-A2-GG6/GG-A3-BI6/GG-A9-SM	malization
Kaxion count: one line in 1.59–2.13 GHz	GG-C1-Kaxion	Fixes ζ_Λ and $z_* = e^7$
		Makes the two H_0 val-
		ues registered observ-
		ables
		The early side carries
		no excess
		F-Hubble-1; under GG-
		A3-BI6’s termination it
		sources no local force
		F-Hubble-7

N.4 Theorem-and-conjecture deck

The reading rests on one imported theorem and one named conjecture, with the rest of its content graded against them (Table 31).

Table 31: The named theorems and conjectures the derived Hubble reading depends on. “Home” gives the paper that proves or states each; this report uses them without re-deriving. Five items is a short deck for a cosmological result, and the shortness is part of the claim. One of the five is open rather than proved, and it is named as open with its target pre-registered, so that the eventual computation cannot be read afterwards as a fit. If any of the other four fails at home, this paper fails with it and no repair made here would help.

Tag	Home	Statement / role
T-Cosmos-17	GG-A10-Cosmos	$\Omega_m = 1/\pi$ as the Hopf fiber-volume ratio (conditional on the measure lemma)
Conjecture VIII	GG-A10-Cosmos	The measure-normalization lemma weighting the Hopf circle; the D^* condition
T-Cosmos-DE ($w = -1$)	GG-A10-Cosmos	The vacuum equation of state is exactly -1 ; falsifier F-Hubble-5
Scale chain T8.7	GG-A5-GDOS	$M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K} = 3.85 \text{ TeV}$, a theorem
Hosotani lemma (O6)	GG-A5-GDOS	Open: closes the selection lemma; target pre-registered, $\zeta_\Lambda = \pi/175$ (F-Hubble-11)

N.5 Promotion roadmap

The reading is honest about what would raise each of its graded claims to the next tier (Table 32). None of these is required for the falsifiable content — the background rate and the siren prediction — to stand.

N.6 Notation

H_0 : present expansion rate ($\text{km s}^{-1} \text{Mpc}^{-1}$). $H_0^{(0)}$: the *derived* background rate, 66.8 (two-road reading 66.6). H_0^{GW} : the standard-siren value. $\Omega_m, \Omega_b, \Omega_\Lambda$: present matter, baryon, and dark-energy fractions. $\omega_i = \Omega_i h^2$: the corresponding physical densities (h -free). ζ_Λ : the vacuum-energy seesaw coefficient. $\epsilon_{\text{local}} = \epsilon_{\text{Hole}} + \epsilon_{\text{syst}}$: the total local excess, read whole as a predicted distance-ladder calibration systematic (no void share carried). $z_* = e^7$: the corridor-switch redshift. M_{lock} : the lock scale, 3.85 TeV. S_8 : the clustering amplitude. ΔN_{eff} : extra effective relativistic species (≤ 0.03 here, a derived null). Cohort acronyms (GG-6, BI-6, OGN-5, GDOS, GRIP) are expanded in App. A.1 and the glossary.

Table 32: What would promote each conditional or asserted entry of the derived Hubble reading. A conditional grade is a promise about what would remove the condition, and this table is that promise written out. Two of the five close on the same measure lemma, one on a Wilson-line computation somebody could start today, and one on an amplitude nobody has yet written down. The fifth is not promoted at all but retired, by a ladder analysis converging on the derived value, and it is listed here because leaving it out would make the column look tidier than it is.

Entry (grade)	What would promote it
π -inventory (D*)	Exhibiting the measure lemma (Conjecture VIII / VIII-b'') closes the Hopf-fiber identification and promotes the fractions to Derived
Background rate 66.8 (Derived cond.; two-road 66.6, D*)	Inherits promotion from the inventory and from the selection lemma; no separate step
Seesaw ζ_Λ (Derived cond.)	The Hosotani/Kaluza–Klein Wilson-line computation on S_θ^1 (GG-A5-GDOS O6) promotes it to Derived (theorem); pre-registered target $\pi/175$ (F-Hubble-11)
Ladder systematic (named, not derived)	Nothing to promote: the framework predicts the systematic’s location and declines to derive its size; a controlled distance-ladder analysis converging on 66.8 retires it, and a converged ladder-free probe at 73 (F-Hubble-4, as re-registered) retires the reading

Author’s Note

This synthesis report is authored as a companion to **GG-A10-Cosmos** (the cosmology recast) and the flagship paper **GG-A11-Origin** (the unified summary of **GG-Theory**). No new cosmological observation is reported here. The content consists of (i) a theory-neutral summary of the H0DN collaboration paper of [Casertano et al. \(2026\)](#); (ii) the derived expansion history $H(t)$ imported from [Arisaka \(2026, A10-Cosmos\)](#) and its confrontation with the microwave sky and the cosmological landmarks; (iii) the reconciliation of the local value — the void channel priced and not carried, and the predicted systematic priced against the ladder’s error budget; (iv) a comparison with six classes of alternative frameworks and the siren discriminator; (v) the standard-siren forecast, with its systematic budget and released scripts; and (vi) an explicit set of kill criteria, observational priorities, and seventy-nine answered objections.

Part XI

Common Objections and Responses

Twenty big questions for the curious reader, plainly answered; then the hardest questions a skeptical cosmologist would ask of the derived-rate reading, answered in the framework’s own terms

This Part answers the questions this paper provokes, in two registers. It *opens* for the curious reader — twenty big questions a thoughtful non-specialist actually carries into a book about the universe’s expansion, answered plainly (§O) — and then hardens, theme by theme, into the objections a skeptical professional would raise, answered in the framework’s own terms. All seventy-nine are tagged in the cohort’s O-Hubble- N registry (the curious-reader set carries the tags O-Hubble-55–74; the third-exchange entries carry 75–79; tags number the ledger, not the reading order), and the Part is closed by a dedicated seven-persona hostile review (§Y) in which a distance-ladder cosmologist, a standard-siren cosmologist, a CMB theorist, a tension statistician, a dark-energy phenomenologist, a philosopher of science, and an internal cohort auditor each put their sharpest question — and then, because a single exchange flatters the respondent, return for a second round against the answers themselves (§Z) — and then, because an answer is itself a claim and claims get audited, a five-reviewer panel puts the sharpest question that survives both rounds; its five, with our answers, close the Part as the third exchange (§Z.1). None is softened, and where the honest answer is “open,” it says so.

O For the curious reader: the big questions, plainly answered

Before the specialist referees take the floor, the questions a thoughtful non-specialist actually brings to a book about the expanding universe — about balloons and edges, mistakes and π , dark things and endings — answered plainly, with pointers to where each answer is earned. Popular science has made some of these feel settled in directions this framework does not take, so we meet them first, and in the open; the device follows the parent paper’s example (Arisaka, 2026, A10-Cosmos). Where an answer rests on a still-open step, it says so.

O-Hubble-55 (expanding into what?). “Everyone says the universe is expanding. Expanding *into* what?”

Response. Into nothing — because expansion is not the universe’s edge advancing through a surrounding space; it is the distances *inside* the universe growing. The mathematics of Course C says it precisely: galaxies keep fixed coordinate labels forever, and the ruler between labels, the scale factor $a(t)$, grows. A useful picture is the surface of an inflating balloon as seen by ants living *on* it: every ant sees every other recede, no ant is at the center, and the question “what is the rubber expanding into?” never arises for creatures whose whole world is the rubber. The

balloon needs a room only because we drew it embedded in one; the universe’s geometry is defined intrinsically, with no room required. What this paper adds to the old picture is only this: the *rate* of that stretching, on our reading, was never a free fact — it is a theorem of the same geometry doing the stretching.

O-Hubble-56 (do we expand too?). “If space itself expands, do atoms, the Earth, and we ourselves expand with it? And if rulers grow too, how could we ever notice?”

Response. No — and the reason is a genuine insight, not a technicality. Expansion is a property of the *averaged, force-free* universe: it is what test particles do when nothing binds them. Atoms are held by electromagnetism, the Earth by chemistry and gravity, the Solar System and the Galaxy by gravity: within every bound system the binding forces long ago overwhelmed the gentle cosmic stretch and settled the system at its own equilibrium size. Quantitatively the mismatch is absurd: the expansion attempts a fractional stretch of order 10^{-18} per second across an atom, while the atom’s own dynamics restores its size 10^{33} times faster. Only across tens of megaparsecs — between galaxies not bound to each other — does the stretch win. That is also why the measurement is honest: our rulers and interferometers are bound systems and do *not* grow, so the recession of distant galaxies is measured against rulers that hold still.

O-Hubble-57 (surely a mistake?). “Two teams measure the same universe and get different answers. Isn’t the boring explanation — somebody made a mistake — overwhelmingly likely? Why has nobody just found it?”

Response. That was everyone’s first hypothesis, and its failure is the real story of the decade. Both chains have been audited at a level almost without precedent in astronomy — the local ladder by thirty-eight specialists across thirty structured variants, every method swapped against every other; the microwave side by three independent instruments on separate continents and calibrations (Part II). No single mistake survived. What *can* survive that kind of cross-checking is the one thing cross-checking is mathematically blind to: an error *shared* by every variant — a common calibration layer — which is a five-line theorem of Course D, not a conspiracy. So the honest menu narrowed to exactly two items: a shared systematic in one chain, or new physics between the epochs. This paper’s contribution is to make the menu decidable: it derives what the rate should be, prices the shared-layer candidates against the ladder’s own literature (Part IV), and hands the verdict to an instrument outside both chains (Part VII).

O-Hubble-58 (faster than light?). “I have read that distant galaxies recede faster than light. Doesn’t that violate relativity?”

Response. No. Special relativity’s speed limit governs motion *through* space — two objects passing each other can never clock more than c . The recession of a distant galaxy is not motion through space; it is the growth of the space between, and general relativity places no limit on that. Beyond the Hubble radius $c/H_0 \approx 4,500$ megaparsecs on the derived rate, galaxies do recede superluminally — and nothing paradoxical follows: no signal, no object, no information overtakes a light beam anywhere, locally, ever. The distinction is the same one Course C’s

metric makes formally: coordinates label, the scale factor stretches, and the speed limit lives in the local light-cones, which expansion tilts but never violates.

O-Hubble-59 (why call it a constant?). “If the expansion rate has been falling since the Big Bang, why does everyone call it the Hubble *constant*?”

Response. Historical inertia, and the paper argues it is time to retire more than the name. H is constant across *space* at any moment — the same everywhere, which is what Hubble’s straight line showed — but it falls steeply across *time*: from unimaginable early values, through 66.8 today, settling toward 55.2 forever (Part III). The name froze in the decades when only the present-day value could be measured, so the one accessible point got mistaken for the whole function. Part IX makes the fuller argument: the “constant” is one reading of a derived curve, and the century of measuring it ever more finely was bound to end in paradox once precision outran the question — as Course L shows with one line of algebra.

O-Hubble-60 (was Einstein wrong?). “If this theory changes cosmology, was Einstein wrong?”

Response. Not by one comma. Every equation this paper integrates is Einstein’s: the Friedmann equations of Course C are his field equations applied to a symmetric universe, exactly as in 1922 (Friedmann, 1922), and the standard siren of Part VII is his prediction doing the refereeing. What changes is not the machinery but *where its inputs come from*. Standard cosmology feeds Einstein’s equations measured contents; this framework claims the contents themselves — how much matter, how much vacuum — were never free, but are fixed by a deeper geometry, and Einstein’s equations then run exactly as before. If anything, the reading is maximally conservative about gravity: no modification, no screening, no extra fields — general relativity inherited whole, with the freedom removed one level *above* it.

O-Hubble-61 (why 95 percent dark?). “Why is the universe 95% dark? It seems absurd that we can see only 5% of everything.”

Response. On the standard account it is simply an awkward fact: two unseen substances dominate, detected only through gravity. On the derived inventory the proportions stop being awkward and become arithmetic: the budget is $\Omega_m = 1/\pi \approx 32\%$ matter, $1 - 1/\pi \approx 68\%$ vacuum, and within the matter only $1/2\pi^2 \approx 5\%$ is ordinary atoms — three readings of one geometric object (Part III). The darkness itself then has a structural reading, which the parent paper states plainly: light — the gauge ledger our instruments are built from — illuminates only the slice of the inventory that carries ordinary charge, so “95% dark” is not a statement about strange substances hiding but about what electromagnetic instruments were ever built to see (Arisaka, 2026, A10-Cosmos). And the claim is not left cosy: the dark matter is named (the Kaxion), massed ($8.1 \mu\text{eV}$), and given a laboratory appointment this decade (Arisaka, 2026, C1-Kaxion).

O-Hubble-62 (pi as numerology?). “Isn’t it suspicious that π shows up in the amount of matter in the universe? That smells like numerology.”

Response. The suspicion is healthy, and the paper’s whole discipline is built to face it. Numerol-

ogy is what you get when you *search*: scan enough elegant expressions against enough measured values and matches are guaranteed. The framework’s numbers run the other way — *forced, not searched*: $1/\pi$ is not one of many candidates that happened to fit; it is the output of a specific geometric construction (the Hopf-fiber volume ratio of §24), computed before comparison, with no alternative on the menu and cohort-wide naming rules that forbid inventing a new formula per observable. The test that separates the two is exposure: a numerologist’s match retreats when the data sharpen; this claim tightens its own noose — the same π must survive every future inventory determination (F-Hubble registry), and the one remaining selected coefficient is pre-registered against a computation that does not yet exist (F-Hubble-11). Numerology never volunteers a way to die.

O-Hubble-63 (the 10-to-the-120 scandal). “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 fix that?”

Response. By deriving the small number instead of canceling the big one. Course F walks the embarrassment honestly: quantum field theory prices the vacuum at the fourth power of whatever scale it trusts, and trusting it to the Planck scale misses the measured value by 10^{120} — the worst prediction in science (Weinberg, 1989). The geometric seesaw replaces the missing cancellation with a hierarchy: the vacuum scale is $M_{\text{lock}}^2/M_{\text{Pl,red}} \approx 6$ meV, built from the laboratory-pinned TeV lock scale, and the famous 120 orders become $8K$ e-folds of one derived descent — e^{-8K} with $K = 35$ is $10^{-121.6}$, dressed by 1.4 orders of geometry (§24.3). Nothing cancels; nothing is tuned; and the same arithmetic, run one step further, hands over the expansion rate itself. The reader with a calculator can check the whole chain in a minute — which is not a sentence the 10^{120} literature has often been able to write.

O-Hubble-64 (why now?). “Matter and dark energy happen to be comparable right now, in our era. Isn’t that coincidence a little too convenient?”

Response. It would be, if “now” were a dart thrown at cosmic history at random — but observers are not darts. On the derived history the window in which matter and vacuum are within a factor of a few of each other spans roughly two percent of the universe’s history on the logarithmic clock that physics actually runs on — narrow, but not absurd. And observers arrive when their chemistry does: the same inventory that times the vacuum’s takeover times the stellar generations that forge carbon and the main-sequence lifetimes that incubate biospheres, so the two clocks that “coincide” were never independent (§100; the full account is the parent’s, Arisaka, 2026, A10-Cosmos). What remains after that accounting is a factor-of-a-few timing question, not a fine-tuning — roughly the surprise of meeting a person whose age is within a factor of two of your own at a dinner party for adults.

O-Hubble-65 (extra dimensions, really?). “Six-dimensional geometry, a ‘bulk,’ a ‘throat’ — extra dimensions sound like science fiction. Why should I believe any of this is real?”

Response. You should not believe it; you should ask what it predicts, and whether it can be killed. Extra dimensions have a respectable century behind them (Kaluza and Klein unified gravity

and electromagnetism with one in the 1920s), and a disreputable recent habit: frameworks with so much freedom that no observation could ever embarrass them. This program’s defense is not that its geometry is plausible but that it is *cornered*: the same six-dimensional structure must simultaneously deliver the Standard Model’s twenty-eight parameters (Arisaka, 2026, A9-SM), the cosmic inventory, this paper’s expansion rate, a dark-matter frequency at 1.95 GHz, and a collider cluster at 3.85 TeV — with shared integers, so it cannot fit one and dodge another. Whether the sixth dimension is “real” in the picture-book sense is a question physics answers operationally: if one line rings in the window with no partner, a microwave receiver will have read a compact dimension’s light-angle count. Until then, the honest word is the one the codebook uses: conditional — and falsifiable on a printed calendar.

O-Hubble-66 (fine-tuned for life?). “The universe looks incredibly fine-tuned for life. If that’s not luck, and not a multiverse, what is it?”

Response. The framework’s answer is the third road: the constants are not tuned, because there was never a knob. A tuning problem presupposes a dial that could have been set otherwise; if the inventory, the coupling strengths, and the expansion history are theorems of one geometry — with the conditional steps graded and printed — then the life-permitting values are forced, and the appearance of fine-tuning is what forcing looks like from the inside. This dissolves the problem rather than winning the lottery: no luck (nothing was drawn), no multiverse (nothing needed selecting from). The position is falsifiable exactly where the derivations are: every constant this program computes is a place the dissolution can fail in public (Part VI). The deeper question — why there is a geometry, and why one that writes readers — is faced, not dodged, in the capstone (Arisaka, 2026, A11-Origin) and, for this paper’s corner of it, in §100.

O-Hubble-67 (could it all collapse?). “You keep flagging results as ‘conditional’ or ‘on a named conjecture.’ Doesn’t that mean the whole edifice could collapse?”

Response. Parts of it could — and the flags exist so that a failure takes down exactly what depends on it and nothing more. The grading system (§A.4) is a load-bearing map, not a disclaimer: the Friedmann machinery is theorem-grade; the π -locked fractions are Derived-conditional on one named measure conjecture; the vacuum coefficient is selected, with the closing computation pre-registered as a falsifier; and the 10% local excess is not claimed at all, only located. Compare the alternative culture, in which a framework’s dependencies are discovered only at its funeral. Here a failed conjecture demotes specific rows of specific tables by rules printed in advance — and the paper has deliberately arranged *local* falsifiers in front of structural ones, so it can die in its own right without taking the cohort’s honesty down with it (§81). An edifice that documents its own demolition points is not fragile; it is auditable.

O-Hubble-68 (exact, yet error bars?). “You call $H_0 = 66.8 \pm 0.2$ ‘derived exactly.’ If it’s exact, why the error bar?”

Response. Because the derivation and the comparison live in different currencies. The geometric objects are exact: $1/\pi$ is not 0.3183 but $1/\pi$; the closed form of §24.3 has no scatter in it. The ± 0.2 enters where exactness meets measurement: Route II divides the sky’s measured densities

(which carry errors) by the exact fractions, and Route I's laboratory pin (the electroweak vev) carries $\pm 0.18\%$ through the lock scale. So read the number as: an exact structure, priced against imperfect inputs — the error bar belongs to our instruments, not to the theorem — with one honest exception, stated in the triage: Route I's ± 0.2 is *structural* (it carries the conditional status of the selection lemma), a theory-side allowance, not an instrumental one. This is the same distinction as between π itself and any measurement of a circle: the theorem is exact; the comparison never is. The falsifier respects the distinction: what would kill the claim is not scatter but a *converged* value away from the structure — Ω_m settling off $1/\pi$, a siren settling at 73.

O-Hubble-69 (why still telescopes?). “If the expansion rate can be computed from five constants on a calculator, why do we still need billion-dollar telescopes?”

Response. For the same reason a court needs witnesses even when the argument is airtight: physics is not mathematics, and a derivation earns belief only by surviving instruments. The telescopes do three things no derivation can. They *test*: every value the geometry prints — the rate, the fractions, the ages, the drift — is an appointment some instrument must keep or break (Part VI). They *measure what the theory declares it cannot reach*: the depth of our local neighborhood, the dust laws of the ladder, the residual's anatomy — the whole empirical layer of Part IV. And they *discover*: nothing in this framework predicted gravitational-wave astronomy into existence; instruments opened that window, and the framework now stakes its life on it. If the reading is right, telescopes stop being the source of cosmology's numbers and become what they are in every mature science: the jury.

O-Hubble-70 (why do sirens decide?). “Out of all the instruments, why do gravitational waves get the final word?”

Response. Because they are the one channel with no shared ancestry with either disputant. Every optical rung of the ladder shares calibrations with every other; every early-universe inference shares the sound horizon. A binary neutron star shares nothing: general relativity itself certifies the source's absolute loudness, so the waveform hands over the distance with no ladder, and the host galaxy hands over the redshift with no plasma physics (Course J walks every step from the quadrupole formula up). In the language of Course L, the siren breaks the deadlock because a tension between two chains can only be adjudicated by a witness correlated with neither. Both camps accepted this referee in advance — which is precisely why the paper stakes its central claim there (F-Hubble-2), and why Part VII spends a whole Part on the referee's own error budget rather than assuming its innocence.

O-Hubble-71 (what if it lands in the middle?). “What if the sirens converge on 70 — right between the camps?”

Response. Then the most interesting world wins, and this paper has already priced its own discomfort in it. A converged 70 would kill the pure 73.5 reading and strain the derived 66.8 at the $1.4\text{--}2.4\sigma$ level (§70); an earlier version of this answer had a second, lifted reading of this paper's sitting comfortably in that world, and it is withdrawn with the lift. The framework

would be pressured but not dead, the corrected-ladder cluster vindicated, and the open question would sharpen into whether any share of the excess is environmental after all — a question the host-distance shells and the ladder shells can answer directly, and one the paper predicts they answer in the negative. The honest summary, printed in advance: the middle world is pressure, reported as exactly that, and uncomfortable for everyone — this reading included.

O-Hubble-72 (how does it end?). “On this theory, how does the universe end?”

Response. Quietly, and on a schedule already visible in the arithmetic. With $w = -1$ exactly, the vacuum neither dilutes nor decays: the expansion rate settles onto the de Sitter floor $H_\infty = 55.2$ and holds it forever, doubling the universe’s size every 12.3 billion years (Course C). Galaxies not bound to ours slide one by one across the horizon; star formation exhausts its gas over trillions of years; the sky slowly empties. No rip (that would need $w < -1$), no recollapse (no turnaround exists on the derived curve), no cycle — and the ending is not a speculation but a graded consequence, guarded by a falsifier: any converged measurement of $w \neq -1$ retires it (F-Hubble-5). It is, as endings go, a dignified one: the universe does not die of catastrophe; it retires on a fixed pension of 55 kilometers per second per megaparsec.

O-Hubble-73 (is this accepted?). “Is GG-Theory accepted by the cosmology community? And why does this enormous claim come from essentially one author?”

Response. No, it is not accepted — it is a proposal under test, and the paper never pretends otherwise. Its response to the sociology is structural rather than rhetorical: every derivation is shown at graduate level (Part X), every number regenerates from released scripts an opponent can re-run, every falsifier is pre-registered with a date, and the community’s own instruments hold every verdict — there is nothing to take on trust and nowhere for authority, anyone’s, to matter. As for provenance: the history of physics adjudicates claims by exposure, not by headcount or by origin, and it keeps a long memory of consensus positions corrected by outsiders — and of outsiders who were simply wrong. Both precedents are real; both are irrelevant once the tests run. The first appointments are months away (§58); skepticism is not merely permitted in the meantime, it is the intended reading mode, and Part X was written so that a skeptical graduate student can check the claims personally rather than choose whom to believe.

O-Hubble-74 (what would convince you?). “Suppose you are wrong. What would it actually take to convince you — or is this, like so many grand theories, effectively unfalsifiable?”

Response. Name a date and a number, and this paper has already named eleven of each. A gravitational-wave catalog converged at 73: the derivation dies, publicly, by its own printed rule. A converged inventory off $1/\pi$; a detected early-dark-energy lobe; a measured $w \neq -1$; a tensor amplitude inconsistent with 0.0038 in either direction; a silent Kaxion window; a single proton decay: each retires a named layer, several retire everything (§56). These are not gestures at distant posterity — the calendar of §58 runs from this December to the mid-2030s, and the commitments ledger collects, appointment by appointment, the numbers that cannot be revised afterwards. The author’s own answer is printed in Part IX’s personal note: if the

sirens side with the ladder, this reading will have done the one useful thing a wrong theory can do — die cleanly, quickly, and in public, leaving the audit trail that shows which link broke. Unfalsifiability is the one accusation this paper has arranged to make impossible.

P On the derived rate

O-Hubble-1 (circularity). “Dividing a predicted fraction by a measured density to ‘derive’ H_0 sounds circular. Isn’t the CMB analysis that gives ω_b, ω_m already assuming a cosmology?”

Response. The two physical densities $\omega_b = \Omega_b h^2$ and $\omega_m = \Omega_m h^2$ are measured from the *shape* of the acoustic peaks — the baryon-loading ratio of the second-to-first peak, the peak-height envelope, the equality redshift — none of which requires knowing h . Indeed ω_b is confirmed independently by the BBN deuterium abundance, with no CMB at all. So the input is model-independent in h ; the GG content is the *prediction* of the pure fractions Ω_i , and the division returns h with no circularity. Standard Λ CDM does the reverse: it fits Ω_i and reads off h . The non-triviality is that GG’s *a priori* fractions, divided into the measured densities, give a rate that also fits the entire spectrum. And since this objection was first drafted, the framework has removed its premise: the rate is now derived *before* the densities are consulted (Route I, §24), so the division is a confirmation, not the derivation (O-Hubble-49).

O-Hubble-49 (the dimensional anchor). “You still needed the sky’s densities — or BBN’s — to get a number with units. A theory that borrows any measured density has not derived H_0 .”

Response. No measured density enters Route I at all (§24.1). The chain is: GU-5 conversion factors (declared unit bookkeeping, §20.1) $\rightarrow$ the theorem-grade lock scale $\frac{1}{2}E_{\text{Pl}}e^{-35}$ (laboratory-pinned through the electroweak vev, not through any telescope) $\rightarrow$ the seesaw $\rightarrow$ the codebook-selected $\pi/175 \rightarrow 66.8 \pm 0.2$. What a theory of geometry owes is the pure number $H_0 t_{\text{Pl}} \approx 1.2 \times 10^{-61}$ (Eq. (11)), and Route I delivers it. The densities appear only in Route II, the *confirmation* — and even there the BBN variant (Eq. (27)) does without the microwave sky. If “derive” means anything operational, it is this: strike every astronomical measurement from the paper, and the rate 66.8 ± 0.2 survives, with the selection intact: no admitted monomial returns 73.5. The whole route now stands as one printed line — the closed form of §24.3, three laboratory constants and two ledger integers — which the reader can check on a calculator in under a minute.

O-Hubble-50 (the convenient monomial). “Your $\zeta_\Lambda = \pi/175$ was chosen because it matches; the ‘cluster’ just proves the codebook can manufacture whatever number you need.”

Response. Three answers, in ascending strength. *First*, the selection uses GG-A9-SM’s admissibility machinery — stratum availability and structural naming — frozen years earlier for twenty-eight Standard-Model parameters, and under it exactly one cluster member survives; the rules reject the numerically adequate orphan $1/J^2$ even though it is *cheaper*, precisely as

they rejected the cheaper $1/\pi$ for θ_{12} . A machinery that selects against cost is a strange tool for manufacturing convenience. *Second*, the admitted cluster cannot manufacture “whatever number”: it has no member at 73.5 at cost ≤ 5 , and the nearest admitted alternative to the selected 66.8 is 74.8, dearer by one unit — had the true rate been the ladder’s, the codebook’s own rules would not have supplied it. (An earlier version answered this objection with a four-row cluster “66.8–71.2” that “cannot reach 73.5”; that table was drawn from an unrestricted alphabet on which e/L^3 lands on 73.5, and the answer is restated here on the stratum table that supports it.) *Third*, the claim is falsifiable at the coefficient level: F-Hubble-11 pre-registers $\pi/175$ as the required outcome of the outstanding Hosotani computation, two-sided, printed before the computation exists. What remains conditional is one named lemma — the stratum assignment — carried at Derived (cond.), not theorem, and the one interpretive step of the selection (the enumeration of the stratum’s priced objects from its three cohort precedents) is flagged as such in §24.2 rather than smoothed over.

O-Hubble-19 (the neutrino convention). “You quote $\omega_m = 0.1430$ including the neutrino-mass contribution; the baryon-plus-CDM sum is 0.1424. Your matter anchor is convention-dependent at the half-percent level.”

Response. Correct, declared in a footnote where the density is introduced, and priced: the two conventions bracket the matter anchor between 66.9 and 67.0, a spread absorbed in the two-road 66.6 ± 0.2 . The baryon route (66.5) carries no such convention, and the two-route consistency — not either route alone — is what the derivation leans on.

O-Hubble-20 (the radiation share). “Your ‘zero-parameter’ curve quietly uses the measured radiation density. That is a parameter.”

Response. It is a measured *input*, declared as such in the honesty triage and in the curve’s own caption: the photon share follows exactly from the CMB temperature (Fixsen, 2009), the neutrino share from the standard N_{eff} . The inventory claims of the framework cover the three dark and matter shares; the radiation share is today a part in 10^4 of the budget, and its only structural role is to fix landmarks (equality, BBN) that the chart then passes rather than fits.

O-Hubble-21 (double dipping). “You take ω_b, ω_m from Planck, then claim the Planck spectrum ‘confirms’ the result — the same data used twice.”

Response. Two different functions of the same sky, and the distinction is the content. The densities come from *shape* statistics (peak-height ratios, the envelope) that do not constrain h ; the confirmation is the *entire spectrum* — positions, ratios, damping — reproduced at the derived rate with zero background freedom ($\chi^2 = 187$, Figure 15). A wrong inventory would pass the first step and fail the second: dividing correct densities by wrong fractions gives a rate at which the peaks do *not* fit. And the sharpest check, the h -free loading ratio $1/2\pi$, uses no rate at all.

Q On the local ladder and the residual

O-Hubble-2 (the ten-percent systematic). “You attribute the whole $\approx 10\%$ local excess to ‘distance-ladder systematics.’ Isn’t that just a cop-out for the part you can’t explain?”

Response. It is the opposite of a cop-out: it is the framework declining to invent physics it has not derived. GG derives none of the excess and says so: it computes the one environmental channel it could have credited, finds the required depth unobserved, and refuses to invent early-universe physics to close the gap — in the abstract, the benchmark sheet, and here. What it does claim is the location: on one background every rung-free probe reads the same rate, so the displaced instrument must be the one with rungs. Naming the excess as a measurement-side systematic the framework predicts, rather than dressing it as a derived result, is the honest bookkeeping the program demands — and the systematic is no longer a nameless gap: its candidate anatomy (dust, priors, subsample) is published, external, contested, kill-testable on dates, and, as Part IV prints against interest, not yet large enough (§34.3). It is also falsifiable: the systematic is pinned to the ladder, and the siren test bypasses the ladder entirely.

O-Hubble-3 (siren at 73). “What if the standard sirens land at 73, not 67?”

Response. Then this reading is wrong, cleanly and without wriggle-room. The siren reads the background expansion rate using neither a sound horizon nor a distance ladder; if it converges near 73 with adequate statistics, the true background rate is high, GG’s fractions cannot be divided into the CMB densities to give it, and the derivation of Section 24 fails. We have staked the framework on ≈ 67 in writing (kill switch F-Hubble-2). This is the single sharpest test in the paper, and it will be settled this decade.

O-Hubble-22 (the kinship charge). “Part V admits your residual is a Class-1 claim. Strip the geometry and your Hubble ‘resolution’ is ‘the ladder is wrong’ — the oldest dodge in the field.”

Response. For the ten percent: yes, and the paper says so in exactly those words. The difference between a dodge and a prediction is what surrounds the claim: a derived background the microwave sky over-determines, an environmental channel computed and found short rather than invoked, a systematic priced channel-by-channel against the ladder’s error budget rather than gestured at, and a named, dated, ladder-free channel that settles it. A dodge hides; this reading has printed where it dies.

O-Hubble-23 (ten — or five, or nine — percent?). “Your headline systematic is $\approx 10\%$, but your own scenario ledger spans $+4.6\%$ to $+10.0\%$. Which is it?”

Response. The requirement is a function of the void depth, $R = 73.50/[66.8(1 + \epsilon_{\text{Hole}})] - 1$ (Eq. (41)), and the paper carries it at $\epsilon_{\text{Hole}} = 0$: $R = +10.0\%$ on the derived rate. The quoted spread is the spread of void depths the surveys have not supplied, printed so no reader mistakes a recorded scenario for a carried one. The parent GG-A10-Cosmos prints “ $\sim 9\%$ ” against ≈ 67 and this paper “ $\approx 10\%$ ” against 66.8; the two are one number. Every quantitative statement in Part IV uses the exact $R(0)$.

O-Hubble-24 (the Kenworthy contradiction). “You cite the supernova-smoothness analyses to cap void explanations, then adopt a +2.9% void lift yourself. You cannot have both.”

Response. Conceded, and resolved the way the objection implies. An earlier version of this paper held two branches at equal standing — a counts branch worth +2.7–2.9% and a flows branch worth at most +1.7% — and let the shell test choose. The paper now holds neither: the *response law* is derived; the *depth* is a conditional input on which the counts literature (deep) and the smoothness-plus-velocity literature (shallow) disagree; the counts branch is an $\approx 11\sigma$ fluctuation on our own power spectrum, charged in print; and the paper carries no lift, records both branches as scenarios, and lets the shell test (H_0 versus $z_{\min}$, F-Hubble-4 as re-registered) read the profile — a flat profile confirms. You cannot have both, and the paper no longer tries to.

O-Hubble-25 (special pleading). “Your ‘one honest route’ to 73.5 needs a contested void *and* seven aligned systematics. That is not a scenario; that is special pleading.”

Response. It would be special pleading if the paper claimed it, and an earlier version invited the charge by printing such a route as a recorded scenario. This version prints none. The scenario ledger of §33 carries no lift; its deepest row, the disfavored KBC depth, still leaves a +4.6% requirement, and the channel table of §34 stacks every published channel at its central value, fully aligned, and reaches 9.3% (8.3% post-rebuttal) against a +10.0% requirement. No combination the paper prints reaches 73.5, and the verdict says so against interest. What the ledger buys is symmetry: the reader can see exactly how expensive each resolution is, including this paper’s — and the siren charges the same price to everyone.

R On the geometry and the π -fractions

O-Hubble-4 (numerology). “ $\Omega_m = 1/\pi$ is numerology. With enough integers and a few transcendentals you can hit any $\mathcal{O}(1)$ number.”

Response. The charge is answered at length in GG-A10-Cosmos; in brief: one closed-packet topology fixes *three* coupled fractions simultaneously ($1/\pi$, $1 - 1/\pi$, $1/2\pi^2$), not one number, and the same integer ledger (1, 3, 5; 7, 35) that produces them also fixes $\sin^2 \theta_W = 25/108$, $\hat{\alpha}^{-1}(M_Z) = 2^7$, the 28 Standard-Model parameters, and the Kaxion spectrum. A per-observable alphabet search grants itself a fresh formula each time; GG reuses one closed ledger across the whole cohort, and that coherent reuse is the non-numerology. If any single fraction were tuned, the acoustic spectrum would break; it does not.

O-Hubble-5 (D^* conditionality). “The fractions are only D^* — conditional on an unproved lemma (Conjecture VIII). So H_0 is not really derived.”

Response. Correct, and stated plainly: the rate inherits the D^* grade of the fractions. The single remaining conditionality is, however, sharp and localized — Conjecture VIII closes as an *epoch identity* (the inventory is exact at the crossing epoch, 75 Myr before the present), with

the residual reduced to one uniformity lemma exact to one part in 10^{15} plus the declared choice of observing epoch. The honest statement is therefore “derived up to one named topological lemma,” not “derived” unqualified — the same grade the inventory carries in GG-A10-Cosmos, and a far stronger position than a fitted parameter.

O-Hubble-6 (the conditional seesaw). “The seesaw coefficient ζ_Λ is Derived only *conditionally* — one lemma short of a theorem. Doesn’t that undercut the whole dark-energy route?”

Response. The condition is named, singular, and cashable. The selection of $\zeta_\Lambda = \pi/175$ rests on GG-A9-SM’s published admissibility and naming rules plus one lemma — the stratum assignment (§24.2) — and its closure, the Hosotani Wilson-line computation (GG-A5-GDOS O6), carries a pre-registered, two-sided target: $\pi/175$ to $\pm 0.6\%$, falsifier F-Hubble-11. If the computation returns anything else, the geometric route demotes itself from a point back to the admitted stratum cluster, whose nearest members to the observed decade are 66.8 and 74.8 — and the precision rate falls back on the density inversion, which does not use ζ_Λ at all. The reading degrades gracefully and publicly; nothing hides in the conditional.

O-Hubble-26 (why the electroweak sphere?). “Even granting the Hopf arithmetic: why should the *cosmic matter share* care about the volume of the *electroweak group*?”

Response. Because in this framework the cosmological stratum and the particle stratum are projections of one bulk, and the OGN-5 integers $N = 3$, $L = 5$ fix the boundary holonomy as the electroweak $SU(2) \cong S^3$ for *both*. The same identification is not spent once for the inventory: it is the identification that pays out the twenty-eight Standard-Model parameters in GG-A9-SM. One may doubt the identification — that doubt is exactly the D^* conditional, named and tracked — but one cannot call it ad hoc: it is the single most heavily reused move in the cohort.

O-Hubble-27 (the drift below $1/\pi$). “Some current CMB+BAO combinations already pull Ω_m toward 0.30–0.31 — below your exact $1/\pi = 0.3183$. Isn’t the chart already strained?”

Response. Fitted Ω_m values are model- and dataset-dependent, and their spread across combinations currently brackets $1/\pi$ from both sides; the chart’s evidence is the *coherent one-sided pattern* of Figure 16, not any single fit. The break condition is printed where the offset is defined: a firm, converged determination of Ω_m below $1/\pi$ — not a preference in one likelihood combination — breaks every entry of the chart at once. If the drift consolidates, this paper loses; that is what the sentence is for.

O-Hubble-28 (the dimensionful scale). “ π can fix ratios. H_0 has units. Where does the scale come from, if not from data through the back door?”

Response. Through the front door, twice, and labeled both times. On the sharp route the scale enters as the two measured physical densities — the only observational inputs in the triage — divided by dimensionless derived fractions. On the seesaw route it enters as M_{lock} and the Planck mass, two scales fixed far from cosmology. The geometry’s claim was never

that it conjures dimensions from π ; it is that once one dimensionful anchor exists (G_N , as the derived-history caption states), every *ratio* in the expansion history is forced.

O-Hubble-48 (which Ω is which?). “You insert $\Omega_m = 1/\pi$ into $H(t)$, but the fractions evolve — at recombination matter was three-quarters of the budget, not $1/\pi$ of it, and your own theorem says the lock holds only at a crossing epoch 75 Myr ago. So how can the same $1/\pi$ appear in your CMB-based H_0 , which concerns a sky released 380,000 years after the Big Bang?”

Response. The objection conflates two different Ω ’s, and the aside of §23 is printed to keep them apart. The symbols in the Friedmann form (17) are present-epoch constants, $\Omega_i \equiv \rho_i(t_0)/\rho_c(t_0)$; the time dependence lives in the dilution factors $a^{-4}, a^{-3}, 1$, and the *running* fractions $\Omega_i(t) = \rho_i(t)/\rho_c(t)$ evolve exactly as the objection says (Figure 10 draws them) — the equation itself puts the budget at the corridor switch at $\Omega_m(z_*) \approx 0.76$, $\Omega_r(z_*) \approx 0.24$, $\Omega_\Lambda(z_*) \approx 2 \times 10^{-9}$. The π -locked values are a boundary condition at the declared present epoch — inserted once, propagated by the law — and the theorem claims the lock only there: the observing epoch is the $\Omega_m(t) = 1/\pi$ crossing, and Part IV prices the why-now coincidence that declaration carries. As for the CMB: its physics at $z \approx 1100$ never consults a present-epoch fraction — baryon loading, equality, and the sound horizon depend on the physical densities ω_b, ω_m , and the Λ term is one part in 10^9 there. The fractions enter only in the present-day bookkeeping identity $\omega_i = \Omega_i h^2$, i.e. in reading *today’s* rate out of densities the early sky fixed — which is what “the CMB-based H_0 ” means. The one fraction applied directly to the early sky is the ratio $\Omega_b/\Omega_m = 1/2\pi$, and that is legitimate precisely because it is epoch-independent: both sectors dilute as a^{-3} , so the ratio at recombination equals the ratio today.

S On the comparison with other solutions

O-Hubble-7 (everyone has a solution). “Early dark energy, extra radiation, modified gravity, local voids — everyone has a Hubble solution. Why is yours different?”

Response. Every entry in the standard list *fits* the background rate and then adds a mechanism to split it; GG *derives* the background rate from geometry fixed by non-cosmological data, and adds no early-universe mechanism at all. It is the only entry whose CMB endpoint is a derivation rather than a fit, and the only one that predicted the siren outcome *ex ante* rather than accommodating it. It is also the only one that would be refuted, not merely disfavored, by a single clean observation (a siren at 73, or a proton decay).

T On falsifiability and scope

O-Hubble-15 (unfalsifiable framework). “A theory that derives everything can explain any outcome. Isn’t it unfalsifiable?”

Response. The opposite: because the numbers are fixed in advance, the reading cannot absorb a surprise. It has printed, before the measurement exists, that a converged standard siren

at 73 kills it (F-Hubble-2) and that a single proton-decay event retires the whole program (F-Hubble-1). A framework that could explain any outcome would name no such switch; this one names eleven, each with a numerical rule.

O-Hubble-16 (the systematic is doing the work). “You derive 66.8, then let a named 10% ‘systematic’ carry you to 73.5. Isn’t the underived systematic the whole game?”

Response. The systematic carries the paper nowhere it claims to reach. The claim is the background rate and the prediction that the ladder is the displaced instrument; the size of the systematic is explicitly *not* claimed — it is labeled a ladder systematic, bounded where it can be bounded (backreaction $< 10^{-3}$, $U(1)_B$ closed, the void channel priced and not carried), and handed to the siren. A theory that dressed the systematic as derived would be the dishonest one; naming it is the discipline, not the evasion.

O-Hubble-17 (why not just fit H_0). “If the derived rate equals the CMB value, what has the geometry added over simply reading H_0 off the CMB?”

Response. Two things the CMB fit cannot give. First, the CMB *measures* H_0 within Λ CDM using six fitted parameters; the geometry *derives* the same number from $\Omega_m = 1/\pi$ with none, so the agreement is a prediction met, not a fit reproduced. Second, the derivation comes with a whole function $H(t)$ and a de Sitter floor $H_\infty = 55.2$, and it couples H_0 to falsifiers in particle physics (proton decay, the Kaxion) that no CMB fit can reach.

O-Hubble-18 (scope creep). “The paper leans on **GG-A10-Cosmos**, **GG-A5-GDOS**, **GG-A4-Euclid**, **GG-A9-SM**, **GG-C1-Kaxion**. Hasn’t the ‘Hubble’ claim quietly become a claim about the entire framework?”

Response. Deliberately, and it is stated as such. This report derives nothing new; it reads the H0DN measurement through imported results, each carried at its home grade (App. N.3). That is a feature: the Hubble reading is exposed to falsifiers in unrelated arenas, so it cannot be rescued by adjusting a cosmological knob. If the shared ledger is wrong, the Hubble reading falls with it — which is exactly the cross-exposure a single-arena fit lacks.

U On the statistics of the tension

O-Hubble-29 (the universal solvent). “Your ‘a common-mode term collapses 7.1σ to 1.4σ ’ argument could dissolve *any* tension in any field. It is unfalsifiable acid.”

Response. The collapse formula is an identity, not an argument: *if* a common-mode term of size s exists, the significance is what it is. The argument is elsewhere, and it is falsifiable twice over: the candidate channels are enumerated with published allowances (so the acid has a measured concentration), and the hypothesis is handed to a channel immune to it. Acid that names its own neutralizer is called a hypothesis.

O-Hubble-30 (pre-registration theater). “H0DN pre-registered its baseline before computing. You ‘pre-registered’ fractions you already knew matched the CMB.”

Response. Two different registrations, and the paper relies on the honest one. The fractions were

fixed by their use across the non-cosmological ledger — the same integers pay the Standard-Model parameters — not by this comparison; but the reader need not accept that provenance to check the paper, because the operative pre-registration is forward-looking: the siren rule, the drift numbers, the $w = -1$ bet, all printed before the data that will test them exist. Retrodictions are graded as consistency; the wagers are graded as wagers.

O-Hubble-31 (curated landmarks). “Thirteen sub-percent agreements — but you chose which seventeen to display.”

Response. The display set is the standard acoustic and background summary set — the observables Planck’s own parameter tables headline — plus the history landmarks; nothing in it is bespoke. The two most precisely measured numbers in cosmology (θ_* and the loading ratio) are both in the set, which is where curation would have hidden a miss. And the strongest reply is structural: on a one-lever chart (§26.4) a hidden miss is impossible — every observable rides the same displacement, so one bad entry would drag the pattern visibly off its sign structure.

O-Hubble-32 (percent is not sigma). “You quote agreements in percent. In measurement sigmas, θ_* is off by 16σ . Any fitted model failing that badly would be dead.”

Response. Correct, and this is the most important statistical sentence in the paper, so it is answered plainly. A *fitted* model is judged in measurement sigmas because its parameters exist to absorb them; this chart has no parameters, claims percent-level derivation, and is displaced from the data by one coherent, named systematic — the π -offset, whose candidate origin (the unfinished measure normalization) is stated in the theorem’s own honesty clause. The correct test of such an object is not θ_* in sigmas but the *coherence* of the displacement across seventeen observables — and, decisively, the printed break condition: if the offset is real physics rather than the lemma’s residue, a sub-percent inventory determination will refuse to converge on $1/\pi$, and the chart dies. The paper does not grade itself on a curve; it names the exam.

V On the siren and the forecast

O-Hubble-33 (leaning on noise). “Current siren values (~ 68 – 70 , errors of several units) ‘lean toward’ you only if one squints. Cherry-picking.”

Response. The paper’s word is “lean,” used once and priced: central values near 68 – 70 with $\sigma \sim 5$ – 8 are statistically compatible with both camps, which is why the decision rule triggers only at converged $\sigma \lesssim 1$. Nothing in the argument is spent on the lean; the forecast Part exists precisely so that the verdict arrives by threshold, not by squinting.

O-Hubble-34 (the referee’s own bias). “Your budget admits the detectors’ amplitude calibration is a 1–2% common mode. Your 5σ margin is 2%. The referee can fake the verdict.”

Response. A 2% calibration bias displaces the converged value by $\sim 1.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$ — one threshold unit, not the 6.7 that separates the camps. To fake a verdict the calibration would

have to be wrong by ten percent, five times its current bound and two orders beyond its design goal; and the bright, dark, and spectral pipelines carry different calibration exposure and must agree. The budget names calibration as the floor to engineer down for the 5σ era — that is a schedule item, not a loophole.

O-Hubble-35 (dark-siren circularity). “Dark sirens need source population priors; population inference needs cosmology. Your volume channel is circular.”

Response. Jointly inferred, not assumed — the standard treatment fits population and cosmology simultaneously, and the prior sensitivity is the channel’s declared $\sim 1\%$ floor in the budget table. The bright channel carries none of it, which is the point of keeping three pipelines whose floors do not overlap: circularity in one cannot survive agreement across all three.

O-Hubble-36 (the slow decade). “Counterpart events may arrive at one a year or slower. Your ‘decade’ could be twenty years.”

Response. The rate uncertainty is honest — a factor of a few — and the timetable is therefore stated in data eras, not dates. Two things bound the pessimism: the dark-siren channel accumulates hundreds of events per year regardless of kilonova luck and is forecast to reach the two-percent class on its own; and the verdict is threshold-triggered, so a slow decade delays the answer without degrading it. The reading’s exposure does not expire.

W On the program and its manner

O-Hubble-37 (too many miracles). “The cohort claims the Standard Model, cosmology, and biology from five integers. Extraordinary claims require more than internal consistency.”

Response. Agreed — they require exposure, which is what the cohort’s breadth buys. Every sector the ledger touches adds falsifiers rather than immunity: a proton decay, a second Kaxion line in the window, a collider null at the lock scale, a siren at 73 — any one of them retires the same five integers everywhere at once. And this paper is deliberately checkable in isolation: its arithmetic runs on a calculator, its scripts are released, and nothing in its kill tests requires believing any sibling.

O-Hubble-38 (why engage with preprints?). “Twenty-two unrefereed preprints on a personal homepage. Why should a working cosmologist spend an afternoon on this?”

Response. Because the afternoon is cheap and bounded: the claims are numerical, the inputs are published, the scripts are released, and the kill criteria are pre-declared — the reader who finds the broken step performs a public service at low cost, and the reader who does not has learned that a zero-parameter chart survives the standard checks. This Part exists to make adversarial reading efficient; hostile review is not an obstacle the paper endures but the mechanism it is built for.

O-Hubble-39 (style over substance). “The paper is written warm — ‘one clock,’ codas, a personal note. Rhetoric is doing work that derivations should do.”

Response. The register is deliberate and the test is simple: delete every warm sentence and check what remains. What remains is the theorem and its conditional, seventeen tabulated confrontations, three derivation chains, two priced budgets, eleven falsifiers with numerical rules, and a forecast with released code — the paper’s full load-bearing structure, none of it prose. The warmth is a reading aid for a long report, and the codas summarize; they do not argue.

O-Hubble-40 (reverse engineering). “All of this was built knowing the answers. The only honest test of a ‘derivation’ constructed after the data is the data it cannot have known.”

Response. The paper agrees so completely that it has organized itself around the principle: the retrodictions are graded as consistency (and defended against per-observable formula freedom by the one-ledger discipline), while the paper’s weight rests on what no one knows yet — the siren verdict, the drift map, the exact w , the tensor ratio, the Ω_m convergence. Every one is printed with a rule. Reverse engineering cannot write post-dated cheques; this paper’s account is full of them, and they clear or bounce in public.

X On the 2026 landscape and the residual anatomy

O-Hubble-51 (the undersized anatomy). “Your three channels, honestly combined, give 4.3% in quadrature — but you need 10%. Your own arithmetic says the anatomy cannot carry the systematic. Why present it as if it could?”

Response. Conceded, and the paper no longer presents it as if it could. Both combination rules are printed in the verdict (§34.5): quadrature gives 4.3%/3.8%, the linear stack 9.3%/8.3% — and the requirement, +10.0%, sits above both. An earlier version, crediting a +2.8% void lift, found a +7.3% requirement between the two rules and read that as the signature of correlated common-mode channels; with the lift withdrawn the anatomy cannot yet carry the systematic however the channels are correlated, and the verdict section says so in those words. What the paper claims is not that the published channels have found all of it, but that the systematic lives in the ladder — predicted from the one-background reading, tested by the probes that use no rung — and that the channels found so far are real, dated, and two-sidedly testable (F-Hubble-8/9/10). If they are confirmed, the landing point is ≈ 70.6 per reanalysis (69.5–70.2 only if fully chained, which no joint analysis yet licenses), still 2.7–3.8 above the derived rate; roughly another four to five percent of common mode is required and not yet published. The scenario spread is the honesty, not the weakness, and the shortfall is printed as a finding against the paper.

O-Hubble-52 (the consensus and its critics). “You keep noting that H0DN excludes the CCHP group — insinuating that the consensus is manufactured by its membership. That is an ugly charge dressed as a footnote. The combination has published criteria,

and CCHP’s values have published counter-analyses.”

Response. No charge of manufacture is made, and the counter-analyses are printed beside CCHP’s values wherever they appear — the SH0ES selection reading of the same JWST data is quoted in the same breath (§34). The point is narrower and, we believe, undeniable: a covariance-weighted combination certifies what its members *share*; it is silent on what its non-members dispute. Both facts are therefore carried at full weight — 73.50 ± 0.81 as the subject measurement of this paper, and 70.39 ± 1.94 as the standing demonstration that calibrator-level choices worth $\sim 3 \text{ km s}^{-1} \text{ Mpc}^{-1}$ are live inside the ladder community itself (§43.2). And no insinuation is required, because the registry names the experiment that dissolves the boundary: the complete ~ 37 -host JWST sample under a convention agreed in advance (F-Hubble-10). This paper is committed to its outcome in either direction.

O-Hubble-53 (the geometric high probes). “Megamasers at 73.9, TRGB–SBF at 73.8, SNe II at 74.9 — one-step geometric probes with no Cepheids and no shared ladder calibration read high. Your common-mode story cannot touch them. Doesn’t the geometric class alone refute the derived rate?”

Response. The class is carried at full strength (§43.4) as the strongest opposing evidence in the landscape — the paper says so in those words. Three observations keep it short of decisive. *Dispersion:* the same one-step family contains the blinded TDCOSMO lensing value at $71.6^{+3.9}_{-3.3}$ and the γ -ray attenuation route at $61.9^{+2.9}_{-2.4}$, while its ladder-free cousins — the fast radio bursts — read 68.8–69.4; a family whose members span fifteen $\text{km s}^{-1} \text{ Mpc}^{-1}$ is reporting its own unmodeled per-method systematics, not a converged 74. *Precision:* no member is percent-class, and the high trio sits 2.4 – 4.3σ from the derived rate at its own quoted errors — with those errors themselves qualified in the record (the maser value’s peculiar-velocity treatment spans 71.8–76.9; the SNe II error is statistical only; the SBF route hangs on a single anchor galaxy). *The endpoint:* the instrument that settles the class’s own spread is the standard siren — geometric in exactly the same one-step sense, and currently reading $69.9^{+4.1}_{-4.0}$. If the geometric class converges high *and* the sirens do too, the derived rate dies (F-Hubble-2); if the class stays dispersed while the sirens read 67–68, the dispersion was the systematics. That is not a dodge; it is the discriminator doing its work.

O-Hubble-54 (laundering the contested corrections). “You disclaim the dust, prior, and subsample corrections as ‘external’ and ‘never GG-Theory results’ — and then use their central values to walk the ladder to 69.5 and score it against your rate. You cannot both disown a number and spend it.”

Response. The grade does exactly the work it declares. Asserted-external means: quoted with the proposer’s caveats and the published rebuttals attached — the priors channel is carried as *proposed, executed, and rebutted* — *undecided*, with its low edge at zero — never averaged into a GG-Theory value, and never used to revise 66.8 ± 0.2 , which has not moved. The chained 69.5–70.2 is a conditional sentence and is printed as one: *if* the published corrections stand, the ladder lands 2.7–3.4 above the derived rate, and the paper scores that against itself. The framework’s own numbers spend nothing from those channels; its exposure runs the other way

— the registry commits in advance (F-Hubble-8/9/10) to shrinking the anatomy, channel by channel, if the corrections fall. Citing an allowance is not claiming a result, and the two-scenario pricing plus the symmetric kill tests exist precisely to keep that distinction visible on the page.

Y The seven-persona hostile review

The reading was put to a panel of seven hostile readers, each asked for the single question they would most want answered before accepting it. The panel read this paper against its parent GG-A10-Cosmos and the dark-matter capstone GG-C1-Kaxion. Their questions, and the answers, follow.

O-Hubble-8 (distance-ladder cosmologist). “H0DN measured 73.50 ± 0.81 with thirty-eight experts and twenty-two stable variants. Your 66.8 is roughly eight sigma below the best local measurement. Why should the ladder yield to a geometric prediction?”

Response. It should not, and the reading does not ask it to. H0DN’s 73.50 is accepted as the local measurement; the claim is that the local value is a biased reading of a *background* rate 66.8 — read high by a ladder-calibration systematic the framework predicts and does not derive, with no share of it environmental. The background, not the local value, is what the geometry predicts, and it lands on the independent CMB determination. The arbiter that is neither CMB nor ladder is the standard siren (F-Hubble-2).

O-Hubble-9 (standard-siren cosmologist). “Present siren constraints already scatter across 67–73 with large errors. When is ‘decisive,’ and what if the catalog converges at 70?”

Response. Decisive means a converged $\sigma_{H_0} \lesssim 1$, reached plausibly by LVK O5 in the late 2020s (a few tens of well-localized bright sirens, or a large dark-siren catalog). The pre-declared rule is sharp at the extremes (67 ± 1 confirms, 73 ± 1 refutes); an intermediate convergence near 70 would be genuinely ambiguous and is exactly the outcome the paper does *not* claim to have pre-decided — it would leave the reading neither confirmed nor refuted, pending tighter data. Current central values sitting nearer 68–70 than 73 lean toward the derived reading without settling it.

O-Hubble-10 (CMB theorist). “ Λ CDM already fits Planck at 67.4. ‘Your fractions fit the peaks at 67’ is not new — you have relabeled the CMB value as ‘derived.’”

Response. The novelty is provenance, not the fit. Λ CDM fits Ω_m, Ω_b as free parameters; GG *predicts* them a priori from geometry and non-cosmological data, and they land on the peaks. The -0.9% π -offset is a specific falsifiable statement (the exact π -fractions sit 1–3% above Planck’s best-fit values); a refit would have hidden it. That the CMB endpoint is a derivation rather than a displaced projection is precisely the content.

O-Hubble-11 (tension statistician). “A seven-sigma tension is a statement about two datasets. Relabeling 67 ‘derived’ does not lower it — the local data still say 73.”

Response. Agreed, and the paper does not claim to lower it by fiat. It reframes the tension: 67 is no longer one of two discordant measurements but a derived prediction, so the open question

becomes whether the *local* measurement is biased (environment plus calibration), not which of two rival rates is right. The statistical resolution comes only from the ladder-free siren, on which the reading is explicitly staked.

O-Hubble-12 (dark-energy phenomenologist). “The 2024 DESI $w_0 w_a$ hint favors evolving dark energy. Your $w = -1$ exactly is already mildly disfavored.”

Response. Acknowledged and carried as a named falsifier (F-Hubble-5). The DESI hint is parametrization- and sample-dependent and not yet decisive; GG predicts $w = -1$ exactly from the throat-cap geometry — a sharp two-sided bet. Euclid and Rubin settle it this decade; a converged $w \neq -1$ at 3σ refutes the throat-residue interpretation of ρ_Λ , and the paper says so in advance.

O-Hubble-13 (philosopher of science). “Deriving H_0 from π and integers — is this explanation, or redescription dressed as derivation? How is it not merely saving the phenomena?”

Response. The mark that separates the two is unforced multiplicity under falsification. One closed ledger (1, 3, 5; 7, 35) fixes three coupled cosmic fractions, the twenty-eight Standard-Model parameters, and the Kaxion spectrum at once, and the paper commits to observations — a siren at 73, a single proton decay — that would refute it. A redescription cannot be refuted by one measurement; this can. Whether that rises to “explanation” the paper leaves to the reader, but it is not merely saving the phenomena.

O-Hubble-14 (internal cohort auditor). “GG-A10-Cosmos anchors on the 2022 Riess 73.0; this paper uses H0DN 73.50. Is the cohort internally inconsistent?”

Response. They are the same local measurement at two vintages, and both papers now score the H0DN 73.50 ± 0.81 and quote Riess’s 73.0 ± 1.0 beside it. GG-A10-Cosmos — the parent, and the source of record for the derived cosmology — carries the structural 66.8 ± 0.2 and this paper carries the same rate; the one-background reading, the ladder as the predicted outlier, and the siren prediction are identical in both, and the 0.25% between GG-A10-Cosmos’s structural value and the companions’ two-road 66.6 is inside every comparison either makes. Only the vintage of the local anchor once differed, and it is documented in each.

Z The second exchange

A single round flatters the respondent: every panellist gets the last word taken away. So the panel was invited back, this time against the answers — and against the new machinery of Parts IV and VII.

O-Hubble-41 (the distance-ladder cosmologist, again). “Your updated channel table still quadratures to only 4.3% — short of your 10% requirement — and the extra mileage comes from other people’s contested papers. You have strengthened *my* case, not yours.”

Response. Conceded, and printed: on the published arithmetic the requirement now exceeds even the aligned stack — an earlier version of this answer said the opposite, on the strength

of a void lift the parent has since withdrawn — and both contested channels are quoted with their opposing evidence, one of them already rebutted by authors of your own team. What has moved is that the channels exist, are dated, and are tested by your instruments, not our theory; what has not moved is the derived rate, which does not move to meet the corrections, and the prediction that the corrections drain all the way to it. If the siren sides with your ladder, this paper has already written its own obituary; the table you cite will be quoted in it.

O-Hubble-42 (the standard-siren cosmologist, again). “Your forecast ignores electromagnetic selection: we follow up the nearest, brightest, most face-on events. Your per-event errors are optimistic.”

Response. Acknowledged, and partially priced: viewing-angle selection is folded with Malmquist in the budget row, and the conservative scenario (20% per event, $N_{3\sigma} = 39$) exists for exactly this reason. The forecast is deliberately elementary and released; a waveform-level treatment with selection functions would sharpen the timetable without moving the destination, and the panel is warmly invited to run one.

O-Hubble-43 (the CMB theorist, again). “A $\Delta\chi^2 = -22$ improvement from *removing* parameter freedom is unusual enough to warrant suspicion. Show the binning, the covariance, the protocol.”

Response. Conceded, and the number is withdrawn. The protocol is the parent’s and is published with it: eighty-three binned TT points, diagonal covariance, background held at the derived chart, only the tilt swapped — a deliberately blunt instrument. Rerun at the exact plateau tilt and the carried rate, χ^2 on the π -locked background is monotone in the tilt (196/187/180 at 0.9649/0.9634/0.9620), so the binned spectrum shows no preference for the computed value and none is claimed. The remedy the objection asks for stands as the first observational priority: a full-likelihood run with the background fixed at the π -locked inventory, where blunt instruments are retired.

O-Hubble-44 (the tension statistician, again). “Your collapse table treats the common mode as one Gaussian knob. Real calibration systematics have structure, correlations, and signs.”

Response. Granted in full. The collapse table is a scale-of-sensitivity illustration — what *one* coherent term of a given size would do — and the paper’s inferential machinery never uses it: no likelihood is built on it, no verdict read from it. The structured version of the question is the channel table, and the decision rule lives entirely on the siren channel, where the knob does not exist.

O-Hubble-45 (the dark-energy phenomenologist, again). “The evolving- w hints are strengthening. If $w(z)$ is confirmed, your exact -1 dies — and your whole derived $H(t)$, and the age, and the floor, die with it.”

Response. Correct on every count, and the paper would have it no other way: $w = -1$ is carried as a two-sided bet (F-Hubble-5) precisely because the derived history cannot follow a drifting equation of state the way a $w_0 w_a$ fit can. If the hints consolidate at 3σ , the throat-residue

reading of ρ_Λ is finished and the de Sitter terminus with it. The exposure is total, is printed, and is the difference between this reading and a parametrization.

O-Hubble-46 (the philosopher of science, again). “You borrow the language of pre-registration from clinical science. Clinical science has registries with timestamps and third parties. Where is yours?”

Response. The honest answer: the cohort’s registry is self-maintained — dated preprints, a cross-paper falsifier ledger, released scripts — and therefore weaker than a clinical registry, though stronger than the field’s default of nothing. Its best feature is structural: the same integers are spent across twenty-two interlocking papers, so a retroactive edit in one breaks checkable arithmetic in the others. Third-party timestamping of the falsifier ledger is a fair demand, and a good idea this exchange has now put on the record.

O-Hubble-47 (the internal cohort auditor, again). “The residual shorthand and its exact value must agree across the cohort — this paper and GG-A10-Cosmos must print the same number, or the books are stale.”

Response. They agree: both papers carry the ladder excess whole as a predicted systematic — GG-A10-Cosmos as “ $\sim 9\%$ ” against ≈ 67 , this paper as $+10.0\%$ against the derived 66.8 (Eq. (41) and the scenario ledger here, the status ledger in the parent) — and no computation anywhere in the cohort consumes a lifted reading. The audit stands, and it stands on the printed text: the shorthand and the exact requirement appear together at every site, which is what makes a discrepancy of this kind checkable by a reader rather than a matter of trust.

The panel’s second round changes no verdicts, which is itself the finding: the answers hold where they claimed to hold, concede where the first round conceded, and the entire disagreement has been compressed onto measurements with dates. That is what this Part is for.

Z.1 The third exchange: the audit of the amendments

Two rounds are a dialogue; an amendment is a claim, and claims get audited. Two rounds are a dialogue; an answer is a claim, and claims get audited. Taking the concessions of the first two exchanges as given — the combination bookkeeping opened, the lift withdrawn, the siren rule unified, the exclusion owned at full strength — a five-reviewer panel (formal theory, Λ CDM cosmology, phenomenology, observation, and a graduate reader) is asked one question each: *what still stands?* Their five sharpest, with our answers, entered the registry as O-Hubble-75–79.

O-Hubble-75 (law or jurisprudence). “Your single-occupancy and no-empirical-filter clauses were codified only after the audit produced $1/(2L^2)$ and $1/56$ — the codebook has, historically, grown a clause each time a cheaper candidate appeared. Freeze the grammar, or admit it is jurisprudence.”

Response. Conceded and done: the freeze is printed in §24.2 — the rulebook as codified is frozen, any future cohort selection must use it unchanged, and “a clause added after the next inconvenient candidate would itself be a falsification of the method.” The ledger-indexed-tower

clause (2^J is one named depth- J mechanism; an unnamed exponent is not) is part of the frozen text, and under it every standing GG-A9-SM selection stands while $2^{-3}/J$ falls. What remains owed, and is registered as a cohort milestone: a standalone grammar registry, and a re-run of all twenty-eight GG-A9-SM selections under the frozen rulebook, published casualties included.

O-Hubble-76 (the one-lever stake). “You say the seventeen-entry pattern rides one lever. Profile the full Planck likelihood along that lever with everything else locked to the chart. My prediction: the -6 pull in $\Omega_m h^3$ and the θ_* residue are transverse, and you will strand a $\Delta\chi^2$ of order 50–100.”

Response. Accepted as *the* decisive computation, and the objection’s prediction is printed here beside the commitment so the outcome grades both parties. OQ-Hubble-4 is sharpened accordingly: the profiled χ^2 along the degeneracy lever, full covariance, minimum and lever amplitude published with the script. Our own printed transverse pull forbids us from predicting a fall to the fitted 84; what the derived chart cannot survive is a profile that rejects the lever direction itself — and if the stranded residual is as large as predicted, the one-displacement reading dies at its own registered test. Both stakes are now on the same page.

O-Hubble-77 (a pointer is not a rule). “You moved the lifted 68.5 onto the shallow bright subset — good. Now write that rule with the same teeth as F-Hubble-2: thresholds, precision requirement, and the undecidable gap named in advance.”

Response. Withdrawn with the lift. The rule the objection asked for was written, in Part VII, and it is now recorded there as an earlier version’s: with no lift carried there is no shallow reading to defend and no branch to discriminate. What survives is the instrument: a shallow bright subset at $z \lesssim 0.05$ remains the one siren sample that would read a local monopole if one existed, and the framework predicts it reads 66.8 like the deep catalog; the ladder-shell test (F-Hubble-4, as re-registered) carries the confirmation, a step consistent with zero at $\pm 0.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

O-Hubble-78 (the operation behind “conditional”). “You price your adopted contrast at 11σ , disfavor the deep void at 20σ for the same crime, and tell the Kenworthy objection you hold neither branch — yet 68.5 ± 0.6 is the only number your landscape arithmetic ever uses. What operation distinguishes ‘adopted conditionally’ from ‘adopted’?”

Response. The operation is now “not carried,” and it is the one the objection was pressing toward: 66.8 is the only number the landscape arithmetic uses. The scenario ledger’s rows are recorded depths, none adopted; the requirement is scored at the whole excess; and the discriminating measurement is pre-registered with a number (the F-Hubble-4 shell step, zero at $\pm 0.3 \text{ km s}^{-1} \text{ Mpc}^{-1}$, which confirms). The asymmetry the earlier answer conceded — headline arithmetic on one branch, labeled at every use — is gone with the branch, and the concession is recorded here so the record shows the objection did its work.

O-Hubble-79 (which sentence is canonical?). “Your appendix says a formal weighting ‘would itself be a modeling choice’ and calls 66.6 an adopted center — but if the baryon

routes are the model-cleanest, as you argue, the clean-route value is 66.4–66.5, not 66.6. Which sentence is canonical?”

Response. The adopted-center sentence is canonical, and it now appears wherever the boxed value does — Part III included, so no reader can re-derive the old objection from the main text alone. The clean-route fact is printed with it: the two baryon routes read 66.4–66.5, the matter route 67.0, and 66.6 ± 0.2 is the center of that structural span, which the companions carry beside the structural 66.8 that GG-A10-Cosmos and this paper carry, with every conclusion surviving any choice in the span. The residual weighting question is not adjudicated by prose at all: it dissolves at OQ-Hubble-4, where a single profiled likelihood replaces every argument about which route deserves the pen.

The third exchange changes no number and retires no test; what it records is that the paper’s sharpest remaining exposures carry a stated counter-prediction printed beside the paper’s own, so that the data grade both. And there will be data. That is what this Part is for.

Part XII

References

The curated bibliography, split into the twenty-two-paper GG-Theory cohort (Part A) and all other works (Part B), each entry carrying a Relevance note and its first in-text citation

Unlike a conventional reference list, every entry below carries a *Relevance* note: a sentence or several saying what the cited work contributes and where in the present paper it bears, closing with the section of first citation. 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, audit, or landscape row each one underwrites.

The notes are written for two readers at once. For the newcomer, they trace the intellectual lineage of the Hubble problem: from Leavitt’s periods and Hubble’s first straight line, through the factor-of-two wars and the precision era, to the 7σ split, the 2026 landscape of forty-five determinations, and the gravitational-wave referee — and then to the twenty-two-paper GG-Theory cohort, in which the present paper reads the contested constant as one point on a derived expansion history, one story above the Standard Model of GG-A9-SM. For the specialist, they make the paper’s debts explicit and checkable: each observational entry names the measurement it supplies (with its arXiv identifier where the landscape table quotes one), so that no citation is merely decorative and every row of Table 13 can be traced to its source.

The cohort papers are cited throughout under the canonical keys and titles fixed by the program’s cross-paper coherence map; this paper’s place in the cohort — the Hubble member of the Consequence branch, the dedicated companion of GG-A10-Cosmos — is set out in Part I’s cohort section. Within each of the two groups 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. Constitutional source of the axioms, postulates, admissibility rules, and OGGEF semantics from which every structural input used in this report descends; the inheritance statement of Part II points here. First cited in §3.

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](https://doi.org/10.5281/zenodo.22064298).

Relevance. Source of the six-dimensional stage $M_4 \times I_r \times S_\theta^1$, BI-6, OGN-5, and the compact-angle holonomy whose cosmological projections the derived history reuses. First cited in §3.

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](https://doi.org/10.5281/zenodo.22064300).

Relevance. Proves the BI-6 legality theorem behind Eq. (9); the 2/5 coefficients of the Arisaka Ratio are fixed here, and so is the boundary-condition termination of the gauged baryon symmetry, under which no $U(1)_B$ vector survives below the lock scale — the structural reason this paper carries no baryonic lift. First cited in §3.

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](https://doi.org/10.5281/zenodo.22064302).

Relevance. Promotes the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ of Eq. (10) to theorem status and states the zero-free-parameters discipline against which this paper's honest provenance labels are drawn. First cited in §6.1.

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](https://doi.org/10.5281/zenodo.22064304).

Relevance. Exact boundary reduction from which four-dimensional observables and open-system dynamics are inherited; the operational-dictionary reading of the Hubble tension used here is a GDOS statement. First cited in §6.1.

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](https://doi.org/10.5281/zenodo.22064306).

Relevance. Source of the four Gi pillars TGG, CCI, GFF, GRIP invoked in the recombination-switch proposition to order the DGI plasma before the frozen attractor. First cited in §6.1.

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](https://doi.org/10.5281/zenodo.22064308).

Relevance. Supplies the Qi runtime and the Arisaka Equation current-locking $J_E/Q_B = J_e = J_B = J_{Q_i}$ cited in the cross-ledger coherence table as the sibling structure of the $U(1)_{Q_i}$ radiation channel. First cited in §6.1.

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 consequence-trio member that recovers textbook quantum mechanics as theorems of the boundary open-system channel. Cited here for the cohort map: it is the measurement-theory sibling of the same projection that makes every Hubble reading a *registered observable* of one background rather than a contradiction. First cited in §3.

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. Fixes the 28-parameter Standard-Model ledger, $\sin^2 \theta_W = 25/108$, $\hat{\alpha}^{-1}(M_Z) = 2^J$, and the $U(1)_B \times U(1)_{Q_i}$ structure; the cross-ledger coherence argument of Part II tests whether the same ledger reaches the Hubble rate. First cited in §3.

K. Arisaka. *The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry*, Arisaka-GG-A10-Cosmos preprint.

Preprint, available at <https://preprints.arisaka-gg.org/a10-cosmos/>.

doi:10.5281/zenodo.22064314.

Relevance. The companion cosmology recast, and the parent of this paper: the source of record for the derived cosmology — the two-route background rate (66.8 from geometry, 66.6 from the densities), the π -locked inventory, the one-background reading in which the calibrated ladder is the predicted outlier, the standard-siren prediction, the computed clustering amplitude $S_8 = 0.825$ and the tension it shares, and the derived radiation null — to which this paper is the dedicated Hubble companion. First cited in §3.

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. The capstone synthesis of the cohort; cited for the cross-domain reach (DNA alphabet, conscious-episode duration) invoked in the cross-ledger coherence table, and as the unified summary referenced in the comparison sections. First cited in §6.1.

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 Predictions-Branch sibling carrying the TeV lock cluster at $M_{\text{lock}} \approx 3.85$ TeV — the very scale whose fourth power sets the vacuum-energy seesaw coefficient ζ_Λ of the Hubble rate derived here. A collider result at the lock scale is collateral evidence for the cosmological reading. First cited in §6.1.

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 Predictions-Branch sibling that spends the same OGN-5 integer ledger (1, 3, 5; 7, 35) on the generation count; cited in the cohort map as an independent falsifier of the same integers that fix the π -locked inventory here. First cited in §6.1.

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 Predictions-Branch sibling: the electron as the lightest charged permanent resident of the same boundary whose cosmological inventory this paper reads. Cited for the branch context of the cohort map. First cited in §6.1.

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 of the Foundations Branch: the two real sources of the boundary’s complex structure — the compact phase circle and the causal clock — from which the single imaginary unit of boundary physics is derived. Cited at roster level in the cohort map; the chronology read here presupposes that complex structure and imports nothing further from it. First cited in §6.1.

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 mirrors-and-matter member 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, with the matter–antimatter asymmetry at the end of the chain. Its four-endowment definition of time is kin, at roster level, to the chronology of the derived $H(t)$ read here. First cited in §6.1.

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 home of the switch this paper hands its whole program to. F-Hubble-1 says that one unambiguous proton-decay event retires GG-Theory in its entirety, and the reason a

cosmology paper can afford to print a rule that severe is that the proton’s absolute stability is a consequence of the exact local $U(1)_B$ rather than an assumption bolted on for the occasion; the paper that works that consequence out, question by question, is this one. It supplies no number to the Hubble chain — the gauge factor’s home is GG-A2-GG6 and GG-A9-SM — and it is cited here for the exposure it creates rather than for anything it contributes. First cited in §56.

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, and the background to a habit this paper uses without arguing for it. The vacuum residue that sets the derived rate is small because a Wilson-line phase admits no local counterterm, not because a symmetry forbids it, and a reader who finds that substitution unfamiliar will find the century of argument behind it there rather than here. It supplies no number and no theorem the Hubble chain uses. First cited in §74.

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 classification of protection mechanisms that this paper’s seesaw coefficient occupies one rung of. The compact phase circle whose holonomy fixes ζ_A , and whose non-locality is the reason the residue survives at the meV scale rather than being renormalized away, is one entry in that classification; a reader who wants to know why that rung and not another should read the classification rather than this paper’s use of it. Cited at roster level in the cohort map, and named here because the phrase the map’s caption uses of it — protection by topology rather than by symmetry — is the mechanism behind the number in §24.3. First cited in §6.1.

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; source of the count — one line, alone, in the 1.59–2.13 GHz laboratory window — behind kill switch F-Hubble-7; the earlier twin-peak reading is withdrawn there. First cited in §6.1.

K. Arisaka. *Black-Hole Entropy under GG-Theory: The GDOS Reading of the Horizon as an OGGE 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 Cosmology-Branch sibling at the strong-gravity end, recovering the exact $A/4$ horizon-entropy coefficient from the same GDOS boundary reading; cited in the cohort map as the horizon-thermodynamics member of this paper’s own branch. First cited in §6.1.

B. Other references

Abbott, B. P., Abbott, R., Abbott, T. D., et al. (LIGO/Virgo and partner collaborations) (2017). A gravitational-wave standard siren measurement of the Hubble constant. *Nature*, 551, 85. arXiv:1710.05835.

Relevance. The first standard-siren H_0 : GW170817 with its kilonova counterpart, 70^{+12}_{-8} — the worked example of the siren Part, and the per-event precision its central forecast scenario adopts. First cited in §64.

Abdul Karim, M., Aguilar, S., Ahlen, S., et al. (DESI Collaboration) (2025). DESI DR2 results II: Measurements of baryon acoustic oscillations and cosmological constraints. *Physical Review D*, 112, 083515. arXiv:2503.14738.

Relevance. Source of the BBN+BAO comparison value $68.51 \pm 0.58 \text{ km s}^{-1} \text{ Mpc}^{-1}$ quoted by H0DN; the second, CMB-independent early-universe channel against which the 5.0σ tension is stated. First cited in §15.

Alfradique, V., Bom, C. R., et al. (2026). An improved measurement of the Hubble constant with LIGO/Virgo/KAGRA O4a dark standard sirens. Preprint (arXiv:2603.20195).

Relevance. The sharpest siren-family value to date: 17 dark sirens with DESI Legacy photometric redshifts plus GW170817 give $69.9^{+4.1}_{-4.0}$ — band 2, ladder- and ruler-free; catalog-completeness systematics dominate. First cited in §43.1.

Alpher, R. A., & Herman, R. (1948). Evolution of the Universe. *Nature*, 162, 774.

Relevance. The prediction that the hot beginning should survive as a few-kelvin relic glow — the forecast the 1965 discovery vindicated; cited in the century narrative of Part I. First cited in §1.

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 $S_8 = 0.759^{+0.024}_{-0.021}$, the shear comparator of the S_8 section and its appendix course; the computed $S_8^{\text{GG}} = 0.825$ sits 2.5σ above it, and the tension is shared. First cited in §61.1.

Aubourg, É., Bailey, S., Bautista, J. E., et al. (2015). Cosmological implications of baryon acoustic oscillation measurements. *Physical Review D*, 92, 123516.

Relevance. The reference inverse-distance-ladder analysis: the early-route logic run downward, landing with the CMB — the mirror image of the ladder described in the routes-anatomy section. First cited in §10.

Baade, W. (1956). The period-luminosity relation of the Cepheids. *Publications of the Astronomical Society of the Pacific*, 68, 5.

Relevance. The first great halving: the two Cepheid populations separated, and the distance scale doubled at a stroke — the century chart’s first correction. First cited in §2.

Bernal, J. L., Verde, L., & Riess, A. G. (2016). The trouble with H_0 . *Journal of Cosmology and Astroparticle Physics*, 2016(10), 019.

Relevance. The canonical statement of the sound-horizon lever arm — how r_s and the inferred rate trade against each other — recapped at the recombination switch. First cited in §31.

Bhardwaj, A., et al. (2025). A globular-cluster Mira anchor for the distance ladder. *Astrophysical Journal* (arXiv:2507.10658).

Relevance. The cluster-Mira rung: 73.06 ± 2.67 from 41 globular-cluster Miras with no metallicity dependence over 1.6 dex — an independent stellar-variable rung siding with the high ladder. First cited in §43.1.

Birrer, S., 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. Independent strong-lensing time-delay determination of H_0 , read as the standing counterweight to the derived rate: it is ladder-free and sound-horizon-free, so a converged time-delay value at 73 refutes the one-background reading exactly as a siren would (F-Hubble-4 as re-registered), and its present $71.6^{+3.9}_{-3.3}$ sits 1.5σ above 66.8. First cited in §15.

Blas, D., Lesgourgues, J., & Tram, T. (2011). The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes. *Journal of Cosmology and Astroparticle Physics*, 2011(07), 034.

Relevance. The Boltzmann solver behind the zero-parameter temperature-spectrum test imported from the parent GG-A10-Cosmos. The instrument to which the derived chart hands its inputs: the zero-freedom TT confrontation of §26 is run with it, and the tier ladder of Part VIII names it explicitly as the machinery the geometry frames but does not replace. First cited in §10.

Boubel, P., et al. (2024). An improved Tully–Fisher estimate of H_0 . *Monthly Notices of the Royal Astronomical Society*, 533 (arXiv:2408.03660).

Relevance. The CosmicFlows-4 Tully–Fisher route: $73.3 \pm 2.1 \pm 3.5$, its systematic term driven by discordant zero-point calibrators — one of the high geometric probes carried at full strength in the counterweight. First cited in §43.4.

Breuval, L., Riess, A. G., Kervella, P., Anderson, R. I., & Romaniello, M. (2022). An improved calibration of the wavelength dependence of metallicity on the Cepheid Leavitt law. *Astrophysical Journal*, 939, 89.

Relevance. Sources the metallicity channel’s status: the term is significant and wavelength-independent from 0.5 to 4.5 μm — the chromatic test a photometric artifact would fail — which is what caps the channel at the half-percent class. First cited in §34.

- Breuval, L., Riess, A. G., Casertano, S., et al. (2024).** Small Magellanic Cloud Cepheids observed with the Hubble Space Telescope provide a new anchor for the SH0ES distance ladder. *Astrophysical Journal*, 973, 30.
Relevance. The SMC Cepheid anchor entering the H0DN variant catalog (V07); part of the anchor-redundancy argument of Part I. First cited in §12.
- Brout, D., Scolnic, D., Popovic, B., et al. (2022).** The Pantheon+ analysis: Cosmological constraints. *Astrophysical Journal*, 938, 110.
Relevance. The Pantheon+ cosmological analysis; part of the supernova-standardization context whose calibration choices enter the residual channel table. First cited in §9.
- Buchert, T., & Räsänen, S. (2012).** Backreaction in late-time cosmology. *Annual Review of Nuclear and Particle Science*, 62, 57.
Relevance. The review of the averaging problem; frames the rigorous bound of the backreaction section and its appendix course. First cited in §35.
- Buchert, T. (2000).** On average properties of inhomogeneous fluids in general relativity: dust cosmologies. *Gen. Rel. Grav.* **32**, 105.
Relevance. The spatial-averaging formalism defining the kinematic backreaction $\mathcal{Q}_D$; the basis of the bound in §35 and Appendix H.4. First cited in §35.
- Calabrese, E., et al. (ACT Collaboration) (2025).** The Atacama Cosmology Telescope DR6: Constraints on extended cosmological models. *JCAP* (arXiv:2503.14454).
Relevance. The ACT team’s own extended-model search: no evidence for early dark energy, varying constants, or modified recombination — one side of the two-sided EDE status of Class 2, opposite the profile-likelihood reading of Poulin et al. (2025). First cited in §46.
- Camphuis, E., Quan, W., Balkenhol, L., et al. (2025).** Joint analysis of SPT-3G, *Planck*, and ACT CMB temperature and polarization data. arXiv:2506.20707.
Relevance. Source of the combined Planck+SPT+ACT early-universe value 67.24 ± 0.35 km s⁻¹ Mpc⁻¹ used as this paper’s reference anchor H_0^{ref} (the companion GG-A10-Cosmos uses the 67.4 convention; the conversion is stated in Part II and Appendix D), and of the SPT-3G-alone value 66.66 ± 0.60 quoted in §30 — the tightest single-instrument CMB determination, landing directly on the derived rate. First cited in §2.
- Carrick, J., Turnbull, S. J., Lavaux, G., & Hudson, M. J. (2015).** Cosmological parameters from the comparison of peculiar velocities with predictions from the 2M++ density field. *Monthly Notices of the Royal Astronomical Society*, 450, 317.
Relevance. The 2M++ peculiar-velocity model used in the H0DN flow corrections; variant V19 tests its removal and shifts the central value by only -0.44 km s⁻¹ Mpc⁻¹. First cited in §17.
- 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 $\sim 1\%$ precision. *Astronomy & Astrophysics*, 708, A166. arXiv:2510.23823.

- Relevance.** The subject measurement of this report: the covariance-weighted Local Distance Network consensus $H_0^{\text{HODN}} = 73.50 \pm 0.81 \text{ km s}^{-1} \text{ Mpc}^{-1}$, its pre-registered baseline, its thirty-plus structured variants, and its orthogonal-path consistency test. First cited in §12.
- Chen, H.-Y., Fishbach, M., & Holz, D. E. (2018).** A two per cent Hubble constant measurement from standard sirens within five years of design sensitivity. *Nature*, 562, 545.
- Relevance.** The reference forecast for the siren channel’s precision clock; the basis of the dark-siren acceleration adopted in the events-to-verdict timetable. First cited in §65.
- Chen, H., Wang, K., et al. (2026).** The local distance ladder: a 2026 review. Preprint (arXiv:2606.26831).
- Relevance.** The covariance-network review: seven ladder routes combine to 73.30 ± 0.92 (5.6σ), and the residual Freedman–Riess disagreement is localized in SN-calibrator subsampling and tip/JAGB standardization conventions — not anchors, not crowding. First cited in §43.2.
- Cruz Reyes, M., & Anderson, R. I. (2022).** A 0.9% calibration of the Galactic Cepheid luminosity scale from Gaia DR3 cluster Cepheids. *Astronomy & Astrophysics*, 672, A85 (arXiv:2208.09403).
- Relevance.** The faint- G bypass of the bright-star zero-point regime: a 0.9% Milky-Way Leavitt-law calibration agreeing with the ladder’s at 0.3σ — the reason the Gaia channel is priced at its low edge in §34. First cited in §34.
- de Bernardis, P., Ade, P. A. R., Bock, J. J., et al. (2000).** A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404, 955.
- Relevance.** The Boomerang first-acoustic-peak measurement that pinned the universe flat — the first CMB pillar of the concordance the dark-energy course narrates. First cited in §9.
- Desmond, H., Stiskalek, R., Nájera, J. A., & Banik, I. (2026).** The subtle statistics of the distance ladder: On the distance prior and selection effects. *Monthly Notices of the Royal Astronomical Society*, in press (arXiv:2511.03394).
- Relevance.** The proposal that the ladder’s implicit uniform-in-modulus distance priors bias every Bayesian rung low in distance and high in H_0 — an Eddington-class common mode. One half of the statistics channel of the residual anatomy, carried at 0–3% with its rebuttal beside it. First cited in §34.3.
- Di Valentino, E., Mena, O., Pan, S., et al. (2021).** In the realm of the Hubble tension — a review of solutions. *Classical and Quantum Gravity*, 38, 153001.
- Relevance.** The encyclopaedic review of proposed resolutions; with the CosmoVerse white paper, the map from which Part V’s six classes are drawn. First cited in §44.
- Di Valentino, E., Said, J. L., Riess, A. G., et al. (2025).** The CosmoVerse White Paper: Addressing observational tensions in cosmology with systematics and fundamental physics. *Physics of the Dark Universe*, 49, 101965.
- Relevance.** The community survey of tension-resolution proposals; frames the six comparison classes examined in Part III. The census against which Part V’s six rival classes are organized,

and the standing measure of how crowded the solution space had become before the landscape’s three-band reading. First cited in §11.

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. (1965). Cosmic black-body radiation. *Astrophysical Journal*, 142, 414.

Relevance. The companion paper that identified the excess antenna temperature as the relic radiation of the hot beginning. First cited in §1.

Ding, C., et al. (2025). The Gaia DR3 parallax zero-point at bright magnitudes from orbital parallaxes. *Astronomical Journal*, 169, 211 (arXiv:2502.08068).

Relevance. The bright-star ($G < 8$) zero-point check: $-38.9 \pm 10.3 \mu\text{as}$ with formal errors underestimated by $\sim 2\times$, and no established sign for the induced H_0 shift — the verified state of the Gaia channel, decided by DR4 (December 2026). First cited in §34.

Dodelson, S. and Schmidt, F. (2020). *Modern Cosmology*, 2nd ed. Academic Press.

Relevance. Standard graduate reference for the growth of structure and the S_8 amplitude used in the derivation course of App. I.5. First cited in §21.1.

Domínguez, A., et al. (2023). A new derivation of the Hubble constant from γ -ray attenuation using improved optical depths. *Monthly Notices of the Royal Astronomical Society*, 527, 4632 (arXiv:2306.09878).

Relevance. The anomalously low outlier of the landscape: $61.9^{+2.9}_{-2.4}$ at fixed $\Omega_m = 0.32$ (free fit $65.6^{+5.6}_{-5.0}$), EBL-model dependent — retained as the reminder that not every anomaly points the same way. First cited in §43.1.

Efstathiou, G. (2020). A lockdown perspective on the Hubble tension (with comments from the SH0ES team). arXiv:2007.10716.

Relevance. The sharpest ladder-critical audit in the literature; several allowances in the residual channel table trace to the issues it raises. First cited in §34.

Einstein, A. (1917). Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, 142.

Relevance. The cosmological constant introduced to hold the universe still — the term whose two later careers (regret, then dark energy) frame the century narrative. First cited in §1.

Eisenstein, D. J., & Hu, W. (1998). Baryonic features in the matter transfer function. *Astrophysical Journal*, 496, 605.

Relevance. The analytic transfer function used by the released backreaction and variance scripts; keeps the paper’s bounds free of external Boltzmann-code dependencies. First cited in §H.4.

Eisenstein, D. J., Zehavi, I., Hogg, D. W., et al. (2005). Detection of the baryon acoustic peak in the large-scale correlation function of SDSS luminous red galaxies. *Astrophysical Journal*, 633, 560.

Relevance. The discovery of the baryon-acoustic ruler in galaxy clustering — the ancestor of the DESI channel that now checks the CMB value without the CMB. First cited in §9.

Etherington, I. M. H. (1933). On the definition of distance in general relativity. *Philosophical Magazine*, 15, 761.

Relevance. The reciprocity (duality) theorem $d_L = (1 + z)^2 d_A$, valid in any metric theory with photon-number conservation — the bedrock identity of the distance toolkit. First cited in §D.1.

Fixsen, D. J. (2009). The temperature of the cosmic microwave background. *Astrophysical Journal*, 707, 916.

Relevance. The definitive CMB temperature: the measurement from which the photon share of the radiation density follows exactly, as invoked in the objections on declared inputs. First cited in §23.

Freedman, W. L., & Madore, B. F. (2010). The Hubble constant. *Annual Review of Astronomy and Astrophysics*, 48, 673.

Relevance. The standard review of the measurement’s first eight decades; companion to the century course and its figure. Also the TRGB physics behind Course D’s second rung — the core-mass argument for the tip as a standard candle — whose calibration conventions the subsample channel of Part IV prices. First cited in §2.

Freedman, W. L., Madore, B. F., Gibson, B. K., et al. (2001). Final results from the Hubble Space Telescope Key Project to measure the Hubble constant. *Astrophysical Journal*, 553, 47.

Relevance. The measurement that ended the factor-two war at 72 ± 8 ; the ten-percent waypoint of the century-long fall traced in the measurement-history course and its figure. First cited in §2.

Freedman, W. L., Madore, B. F., Hoyt, T. J., et al. (2025). Status report on the Chicago-Carnegie Hubble Program (CCHP): Measurement of the Hubble constant using the near-infrared tip of the red giant branch. *Astrophysical Journal*, 985, 203. arXiv:2408.06153.

Relevance. The independent TRGB pipeline: the published JWST value is 70.39 ± 1.94 (with JAGB and Cepheid variants spanning 67.8–72.1) — the low side of the subsample dispute priced in the residual anatomy, and part of the JWST cross-check of Priority 4. First cited in §34.

Freedman, W. L. (2021). Measurements of the Hubble constant: Tensions in perspective. *Astrophysical Journal*, 919, 16.

Relevance. The TRGB-side perspective ($H_0 \approx 69.8$) whose spread against the Cepheid route informs the primary-indicator-family allowance in the residual channel table. First cited in §34.

Friedmann, A. (1922). Über die Krümmung des Raumes. *Zeitschrift für Physik*, 10, 377.

Relevance. The honest solution of Einstein’s equations: a matter-filled universe must expand or contract — the dynamical law behind everything this paper derives. First cited in §1.

Gall, C., Izzo, L., Wojtak, R., & Hjorth, J. (2026). The Hubble constant from blue Type Ia supernovae. *Astrophysical Journal*, 1001, 233. arXiv:2411.05642.

Relevance. The dust-free cross-check of the dust channel: a ladder built only from intrinsically blue, low-reddening supernovae returns $H_0 = 70.0 \pm 2.1$ — no extinction model at all, landing where the two-population reanalysis lands. First cited in §34.3.

- Gao, H., et al. (2024).** Measuring the Hubble constant with localized and nonlocalized fast radio bursts. *Astronomy & Astrophysics*, 698, A215 (arXiv:2410.03994).
- Relevance.** The FRB channel: 108 localized bursts give $69.40^{+2.14}_{-1.97}$ (527 nonlocalized CHIME bursts: 68.81 ± 0.68 , dispersion-measure-model dependent) — a camp-free probe sitting squarely in band 2. First cited in §43.1.
- García Escudero, H., et al. (2025).** A sound-horizon-agnostic measurement of H_0 from BAO, lensing, and galaxy clustering. *Astrophysical Journal Letters*, 998, L41 (arXiv:2509.16202).
- Relevance.** The high edge of the ruler-free span: 70.0 ± 1.7 , prior-sensitive — the reason this paper prints the honest r_s -free range 66.7–70.2 rather than the convenient half of it. First cited in §43.3.
- Giovanetti, C., et al. (2026).** A generic baryon-density tension for early-time resolutions of the Hubble tension. Preprint (arXiv:2604.05095).
- Relevance.** The deuterium audit as a generic wall: early-time fixes that raise H_0 raise the inferred ω_b against the LUNA-pinned BBN value — wielded against Classes 2–4. First cited in §46.
- Hallgren, A., et al. (2025).** Diverse host-galaxy reddening distributions and their effect on Type Ia supernova distances. *Astronomy & Astrophysics*, 706, A129.
- Relevance.** The simulation-level face of the dust channel: imposing a universal extinction prior on intrinsically diverse host populations provably biases the recovered relative distances — the mechanism’s existence proof, independent of any H_0 claim. First cited in §34.3.
- 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 any early-lobe S_8 story, and part of the no-early-modification case of Class 2. First cited in §46.
- Högås, M., & Mörtzell, E. (2026).** Physically-motivated priors in the local distance ladder significantly reduce the Hubble tension. *Monthly Notices of the Royal Astronomical Society* (arXiv:2601.22215).
- Relevance.** The executed recalibration of the statistics channel: physically motivated distance priors shift the ladder from 73.0 ± 1.0 to 70.6 ± 1.0 and the quoted tension from $\sim 5\sigma$ to $\sim 2\sigma$. Read against its rebuttal (Stiskalek et al., 2026); the pair is why the channel is printed “undecided.” First cited in §34.3.
- Holz, D. E., & Hughes, S. A. (2005).** Using gravitational-wave standard sirens. *Astrophysical Journal*, 629, 15.
- Relevance.** Developed the standard-siren program and its name; source of the amplitude-distance reading of Eq. (50). First cited in §64.
- Huang, C. D., et al. (2023).** The Mira distance to M101 and a Mira-based Hubble constant. *Astrophysical Journal* (arXiv:2312.08423).

- Relevance.** The Mira rung: 211 Miras in M101 give 72.37 ± 2.97 — an oxygen-rich variable route independent of Cepheids and TRGB, siding with the high ladder. First cited in §43.1.
- Hubble, E., & Humason, M. L. (1931).** The velocity-distance relation among extra-galactic nebulae. *Astrophysical Journal*, 74, 43.
- Relevance.** The velocity-distance law confirmed to twenty times the original depth; the second point on the century chart’s opening plunge. First cited in §1.
- Hubble, E. (1925).** Cepheids in spiral nebulae. *The Observatory*, 48, 139.
- Relevance.** Cepheids resolved in Andromeda: the island-universe debate settled and the universe of galaxies opened. The moment the ‘nebulae’ became galaxies and the universe acquired an outside: the first rung of Part II’s century course, and the birth of the distance ladder as an instrument. First cited in §1.
- Hubble, E. (1929).** A relation between distance and radial velocity among extra-galactic nebulae. *Proc. Natl. Acad. Sci.* 15, 168.
- Relevance.** The founding measurement of the expansion, $v = H_0 d$; cited for the historical opening of the Introduction. The founding one-number cosmology — whose graduation from measured constant to derived quantity is this paper’s closing argument (Part IX). First cited in §1.
- Hughes, N., Wilensky, M., & Bull, P. (2025).** A Bayesian jackknife of the Hubble tension. Preprint (arXiv:2511.19341).
- Relevance.** The model-agnostic bias audit over 16 key probes: “no bias anywhere” is strongly disfavored, many bias assignments fit comparably, and the agnostic posterior reads $H_0 \approx 68^{+4.7}_{-1.3}$ (95%) — the landscape’s warning against convicting any single culprit. First cited in §43.1.
- Jedamzik, K., Pogosian, L., et al. (2025).** Hints of primordial magnetic fields at recombination. *Nature Astronomy*, 10, 317 (arXiv:2503.09599).
- Relevance.** The baryon-clumping mechanism of Class 3: 5–10 pG fields hastening recombination, currently a 1.8σ , dataset-dependent hint — carried at exactly that strength. First cited in §47.
- Jensen, J. B., Blakeslee, J. P., Cantiello, M., et al. (2025).** The TRGB-calibrated SBF distance scale with JWST and HST. *Astrophysical Journal*, 987, 87. arXiv:2502.15935.
- Relevance.** The surface-brightness-fluctuation rung of the TRGB-SBF path; backbone of the statistically independent orthogonal path O2 in the H0DN consistency test. First cited in §43.4.
- Jia, X. D., Hu, J. P., Yi, S. X., & Wang, F. Y. (2025).** Uncorrelated estimations of H_0 redshift evolution from DESI baryon acoustic oscillation observations. *Astrophysical Journal Letters*, 979, L34.
- Relevance.** The published $H_0(z)$ descending trend (6.4σ , uncorrelated DESI bins) — the redshift-shell signature this paper names as the survey-only test of whether any share of the ladder excess is environmental, carried with its provenance cautions in §33.4. First cited in §33.4.
- Jimenez, R., & Loeb, A. (2002).** Constraining cosmological parameters based on relative galaxy ages. *Astrophysical Journal*, 573, 37.

Relevance. The original cosmic-chronometer proposal — galaxies as clocks — behind the differential-age channel of the over-determination program. First cited in §15.

Kamionkowski, M., & Riess, A. G. (2023). The Hubble tension and early dark energy. *Annual Review of Nuclear and Particle Science*, 73, 153.

Relevance. Review of the tension and the sound-horizon lever arm; the standard linear-response logic recapped in Part II descends from this literature. First cited in §11.

Karwal, T., & Kamionkowski, M. (2016). Dark energy at early times, the Hubble parameter, and the string axiverse. *Physical Review D*, 94, 103523.

Relevance. The original early-dark-energy proposal; ancestor of the phenomenological EDE class compared against the derived reading in Class 2. First cited in §46.

Keenan, R. C., Barger, A. J., & Cowie, L. L. (2013). Evidence for a ~ 300 megaparsec scale under-density in the local galaxy distribution. *Astrophysical Journal*, 775, 62.

Relevance. The deep reading of the Local Hole (the “KBC void”): a $\sim 30\%$ near-infrared count deficit to ~ 300 Mpc; the contested upper scenario of the void-channel ledger, recorded and not carried. First cited in §33.1.

Kenworthy, W. D., Scolnic, D., & Riess, A. G. (2019). The local perspective on the Hubble tension: local structure does not impact measurement of the Hubble constant. *Astrophysical Journal*, 875, 145.

Relevance. The dynamical brake on the deep reading: the smooth supernova Hubble diagram leaves room for at most a sub-percent void-induced shift of the SN-anchored rate; one of the two reasons the deep void is disfavored and the void channel is priced and not carried. First cited in §33.1.

Kolb, E. W., & Turner, M. S. (1990). *The Early Universe*. Addison-Wesley (Frontiers in Physics 69).

Relevance. The classic graduate text on thermal history, entropy conservation, and freeze-out — the standard companion to the thermal annex of Course K. First cited in §K.5.

Kreisch, C. D., Park, M., Calabrese, E., et al. (2024). Atacama Cosmology Telescope: constraints on early dark energy from DR6 and DR4+Planck. *Journal of Cosmology and Astroparticle Physics*, 2024, 055.

Relevance. Current CMB squeeze on f_{EDE} ; the observational pressure quoted in the EDE comparison and in kill switch F-Hubble-3. First cited in §46.

Kumar, S., et al. (2026). The Hubble constant from thirteen bright standard sirens in AGN discs. Preprint (arXiv:2607.03672).

Relevance. Thirteen binary-black-hole mergers with claimed AGN-flare counterparts give $70.50^{+3.37}_{-2.89} \pm 1.56$ — carried as a candidate only, the flare associations being unconfirmed; if validated, a second bright-siren channel a decade early. First cited in §43.1.

Leavitt, H. S., & Pickering, E. C. (1912). Periods of 25 variable stars in the Small Magellanic Cloud. *Harvard College Observatory Circular*, 173, 1.

- Relevance.** The period-luminosity law: the key that made cosmic distance measurable, and the first rung of every ladder since. First cited in §1.
- Lee, A. J., et al. (2024).** The JAGB method: a blind carbon-star distance ladder in seven supernova hosts. CCHP series (arXiv:2408.03474).
- Relevance.** The blind carbon-star route of the CCHP program: $67.80 \pm 2.17 \pm 1.64$ on the N4258 zero point — the lowest of the CCHP family, and one pole of the subsample dispute priced in §34. First cited in §43.2.
- Lee, N., et al. (2026).** Modified recombination II: varying the electron mass against DESI DR2. Preprint (arXiv:2606.06495).
- Relevance.** The key negative result of Class 3: the $m_e(z)$ variant that resolves the CMB-only fit fails once DESI DR2 is included — one of the four audits that cap modified recombination below the ladder. First cited in §47.
- 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.
- Relevance.** Derived the expansion law from general relativity and extracted the constant two years before Hubble’s plot; opens the measurement-history course. First cited in §1.
- Lindgren, L., Bastian, U., Biermann, M., et al. (2021).** Gaia Early Data Release 3: Parallax bias versus magnitude, color, and position. *Astronomy & Astrophysics*, 649, A4.
- Relevance.** The map of the Gaia parallax zero-point — the anchor-floor systematic of the residual-anatomy channel table; its residual position and color dependence sets the allowance adopted there. First cited in §10.
- Liske, J., Grazian, A., Vanzella, E., et al. (2008).** Cosmic dynamics in the era of extremely large telescopes. *Monthly Notices of the Royal Astronomical Society*, 386, 1192.
- Relevance.** The design study for measuring the redshift drift with an extremely large telescope; the instrument context for the drift predictions of the derived curve. First cited in §59.
- Loeb, A. (1998).** Direct measurement of cosmological parameters from the cosmic deceleration of extragalactic objects. *Astrophysical Journal Letters*, 499, L111.
- Relevance.** Revived the redshift-drift measurement for the Lyman- α forest era; with Sandage (1962), the basis of the drift predictions of Table 18. First cited in §59.
- Louis, T., La Posta, A., Atkins, Z., et al. (2025).** The Atacama Cosmology Telescope: DR6 power spectra, likelihoods and Λ CDM parameters. arXiv:2503.14452.
- Relevance.** The ACT DR6 release: the second, fully independent CMB instrument whose agreement with Planck underwrites the early-universe network argument. First cited in §15.
- Lu, W., et al. (2026).** A sound-horizon-free measurement of H_0 from the EFTofLSS power spectrum and bispectrum. *Physical Review D*, 114, 023518 (arXiv:2602.21733).
- Relevance.** The first $P+B$ ruler-free determination: $70.2^{+2.2}_{-2.4}$ baseline (68.6 ± 1.8 with Pantheon+) — the top of the honest r_s -free span. First cited in §43.3.

Abac, A., et al. (LIGO–Virgo–KAGRA Collaboration) (2025). GWTC-4.0: Constraints on the cosmic expansion rate and modified gravitational-wave propagation. Preprint (arXiv:2509.04348).

Relevance. The 142-source dark-siren catalog value $76.6^{+13.0}_{-9.5}$ — the wide, mass-model-dependent start of the siren trajectory the discriminator section tracks. First cited in §51.

LIGO–Virgo–KAGRA Collaboration (2026). GWTC-5.0: Constraints on the cosmic expansion rate and modified gravitational-wave propagation. Preprint (arXiv:2605.27227).

Relevance. 236 sources with the DES-Y6 catalog and GW170817: $71.0^{+9.0}_{-7.1}$, a 25.7% tightening over GWTC-4.0 with no gravitational-wave-propagation anomalies — the siren trajectory drifting toward band 2 as catalogs deepen. First cited in §51.

Lynch, G. P., Knox, L., & Chluba, J. (2024). Reconstructing the recombination history: non-parametric constraints and the Hubble tension. *Physical Review D* (arXiv:2404.05715).

Relevance. The ceiling of Class 3, printed by its own machinery: with full spline freedom in the ionization history, CMB+BAO allow at most $H_0 < 70.7$ at 95% — the citable cap this paper quotes in place of any point value. First cited in §43.3.

Lynch, G. P., Knox, L., & Chluba, J. (2024). DESI and the Hubble tension in light of modified recombination. *Physical Review D* (arXiv:2406.10202).

Relevance. The DESI audit of Class 3: the spline histories that raise H_0 drag Ω_m against the DR2 measurement, leaving SH0ES-level rates at $\sim 2\sigma$ residual tension. First cited in §47.

Madhavacheril, M. S., et al. (ACT Collaboration) (2023). The Atacama Cosmology Telescope DR6 gravitational lensing map and cosmological parameters. *Astrophysical Journal*, 962, 113 (arXiv:2304.05203).

Relevance. The lensing route with no primary-anisotropy fit: 68.3 ± 1.1 with BAO, and Λ CDM-consistent lensing at the percent level — the resolution of the old A_L anomaly. First cited in §43.3.

Maggiore, M. (2008). *Gravitational Waves, Volume 1: Theory and Experiments*. Oxford University Press.

Relevance. The standard reference for linearized gravity, the TT gauge, and the quadrupole formula — the full derivations compressed in Course J. First cited in §J.1.

Marra, V., Amendola, L., Sawicki, I., and Valkenburg, W. (2013). Cosmic variance and the measurement of the local Hubble parameter. *Phys. Rev. Lett.* **110**, 241305.

Relevance. The reference computation of the cosmic variance of the locally measured H_0 ; corroborates the $\lesssim 0.4\%$ scatter used in §35. First cited in §35.

Mazurenko, S., Banik, I., & Kroupa, P. (2025). The redshift dependence of the inferred H_0 in a local void solution to the Hubble tension. *Monthly Notices of the Royal Astronomical Society*, 536, 3232.

Relevance. Shows that a local void predicts a declining $H_0(z)$ — the redshift-shell signature this paper names as the sharpest survey-only test of an environmental share — with amplitude tracking the void depth, and predicted flat here. First cited in §33.4.

Mészáros, P. (1974). The behavior of point masses in an expanding cosmological substratum. *Astronomy & Astrophysics*, 37, 225.

Relevance. The radiation-era stall of matter perturbations (the Mészáros effect), which imprints the equality scale on the power spectrum. First cited in §I.2.

Mirpoorian, S. H., Jedamzik, K., & Pogosian, L. (2024). Modified recombination and the Hubble tension. *Physical Review D*, 111, 083519 (arXiv:2411.16678).

Relevance. The strongest positive case inside Class 3: spline ionization histories can cut the tension below 2σ in some data combinations — quoted with its dataset dependence, against the four audits that cap the class. First cited in §47.

Moresco, M., Amati, L., Amendola, L., et al. (2022). Unveiling the Universe with emerging cosmological probes. *Living Reviews in Relativity*, 25, 6.

Relevance. The reference review of cosmic chronometers and the other ladder-free, horizon-free probes of $H(z)$; source of the differential-age method the over-determination section leans on. First cited in §15.

Mörtsell, E., Goobar, A., Johansson, J., & Dhawan, S. (2022). Sensitivity of the Hubble constant determination to Cepheid calibration. *Astrophysical Journal*, 933, 212.

Relevance. The Cepheid-calibration sensitivity study behind the metallicity-term allowance in the residual channel table. First cited in §34.

Murakami, Y. S., Riess, A. G., Stahl, B. E., et al. (2023). Leveraging SN Ia spectroscopic similarity to improve the measurement of H_0 . *Journal of Cosmology and Astroparticle Physics*, 2023, 046.

Relevance. Part of the SN Ia calibration-refinement literature whose debates the H0DN variant V13 (no SNe Ia) renders moot for the central value. First cited in §14.

Pantos, V., & Perivolaropoulos, L. (2026). Dissecting the Hubble tension: 88 measurements, two populations. Preprint (arXiv:2601.00650).

Relevance. The meta-split, quantified: 30 distance-ladder measurements average 72.73 ± 0.39 while 58 ladder-free measurements average 69.37 ± 0.34 — internally discrepant at up to 6.5σ ($\geq 3.9\sigma$ with correlations); the statement of the comparison problem in two numbers. First cited in §43.1.

Penzias, A. A., & Wilson, R. W. (1965). A measurement of excess antenna temperature at 4080 Mc/s. *Astrophysical Journal*, 142, 419.

Relevance. The accidental discovery of the cosmic microwave background — the moment the Big Bang became measured history; cited in the century narrative that opens Part I. First cited in §1.

Perlmutter, S. et al. (1999). Measurements of Ω and Λ from 42 high-redshift supernovae. *Astrophys. J.* **517**, 565.

Relevance. The companion discovery of cosmic acceleration; cited with [Riess et al. \(1998\)](#) for the dark-energy surprise. With the High- z team’s discovery, fixes the acceleration whose derived timing ($t_{\text{acc}} = 7.8$ Gyr, the exact $\text{arcsinh}(1/\sqrt{2})$ epoch) Course C computes with nothing tuned. First cited in §1.

Pesce, D. W., et al. (2020). The Megamaser Cosmology Project. XIII. Combined Hubble-constant constraints. *Astrophysical Journal Letters*, **891**, L1 (arXiv:2001.09213).

Relevance. The purest one-step geometric probe: six maser-galaxy distances give 73.9 ± 3.0 , sliding across 71.8–76.9 with the peculiar-velocity treatment of six objects — the leading member of the high geometric counterweight. First cited in §43.4.

Peskin, M. E., & Schroeder, D. V. (1995). *An Introduction to Quantum Field Theory*. Reading, MA: Addison-Wesley.

Relevance. Standard reference for the gauge-theory reading of an ultra-weakly coupled vector, cited in the record of the earlier graviphoton estimate; that estimate is closed by the boundary-condition termination of GG-A3-B16, and the cohort’s radiation term is a derived null rather than a decoupling-temperature estimate. First cited in §G.5.

Pietrzyński, G., Graczyk, D., Gellenne, A., et al. (2019). A distance to the Large Magellanic Cloud that is precise to one per cent. *Nature*, **567**, 200.

Relevance. The one-per-cent DEB distance to the LMC; one of the three geometric anchors of the H0DN baseline network. The anchor whose one-per-cent geometry stiffens the ladder’s floor; its independence from Gaia’s zero point is part of why the parallax channel is carried at only 0.5–1% in the anatomy. First cited in §10.

Planck Collaboration (2020). Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics*, **641**, A6. arXiv:1807.06209.

Relevance. The baseline-dictionary parameter source: the N_{eff} preference $\Delta N_{\text{eff}} = 0.04 \pm 0.18$ quoted in the exposure caveat, and the CMB-inferred $S_8^{\text{ref}} \approx 0.83$. First cited in §9.

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. The flagship early-dark-energy resolution of the tension, and the source of the sound-horizon lever-arm logic recapped in the recombination-switch section; the standard against which the derived reading’s no-early-modification stance is contrasted in Class 2. First cited in §31.

Poulin, V., Smith, T. L., & Karwal, T. (2023). The ups and downs of early dark energy solutions to the Hubble tension: A review of models, hints and constraints circa 2023. *Physics of the Dark Universe*, **42**, 101348.

Relevance. Status review of the EDE band and its S_8 cost; the comparison matrix and the EDE comparison draw their phenomenological state of play from here. First cited in §17.

Poulin, V., et al. (2025). Early dark energy after ACT DR6 and DESI DR2: a profile-likelihood view. *Physical Review D*, 113, 063519 (arXiv:2505.08051).

Relevance. The other side of the EDE status: a prior-free profile likelihood on ACT DR6 + DESI DR2 gives $f_{\text{EDE}} = 0.09 \pm 0.03$ and $H_0 = 71.0 \pm 1.1$ with no SH0ES input — a live statistical-method dispute with the ACT team’s own null, and this paper’s F-Hubble-3 kill test in motion. First cited in §46.

Reid, M. J., Pesce, D. W., & Riess, A. G. (2019). An improved distance to NGC 4258 and its implications for the Hubble constant. *Astrophysical Journal Letters*, 886, L27. arXiv:1908.05625.

Relevance. The megamaser distance to NGC 4258; the geometric anchor shared by the otherwise disjoint orthogonal paths of the H0DN consistency test. First cited in §10.

Riess, A. G. et al. (1998). Observational evidence from supernovae for an accelerating universe and a cosmological constant. *Astron. J.* **116**, 1009.

Relevance. One of the two discovery papers of cosmic acceleration; cited for the dark-energy surprise that makes the expansion rate genuinely time-dependent. First cited in §1.

Riess, A. G., Casertano, S., Yuan, W., et al. (2019). Large Magellanic Cloud Cepheid standards provide a 1% foundation for the determination of the Hubble constant. *Astrophysical Journal*, 876, 85.

Relevance. The 74.03 ± 1.42 waypoint of the precision-era track drawn in the century figure. The waypoint at which the tension crossed 4σ and ‘crisis’ entered the field’s vocabulary — a fixed point of the century course, and the moment the precision-amplifies-bias clock of Course L became audible. First cited in §8.

Riess, A. G., Yuan, W., Macri, L. M., et al. (2022). A comprehensive measurement of the local value of the Hubble constant with $1 \text{ km s}^{-1} \text{ Mpc}^{-1}$ uncertainty from the Hubble Space Telescope and the SH0ES team. *Astrophysical Journal Letters*, 934, L7.

Relevance. The SH0ES measurement 73.04 ± 1.04 , quoted as 73.0 ± 1.0 beside the H0DN 73.50 that GG-A10-Cosmos and this paper score; also the source of the deceleration and jerk priors $q_0 = -0.55$, $j_0 = +1$ held fixed by the H0DN analysis. First cited in §2.

Riess, A. G., Scolnic, D., Anand, G. S., et al. (2024). JWST observations reject unrecognized crowding of Cepheid photometry as an explanation for the Hubble tension at 8σ confidence. *Astrophysical Journal*, 977, 120.

Relevance. The JWST instrument cross-check on HST Cepheid photometry; part of the Priority-4 instrument-swap test that the derived reading treats as confirming $\epsilon_{\text{local}}^{(B)} > 0$ is not an HST artifact. First cited in §34.

Riess, A. G., Scolnic, D., Anand, G. S., Breuval, L., Casertano, S., et al. (2024). JWST validates HST distance measurements: Selection of supernova subsample explains differences in JWST estimates of local H_0 . *Astrophysical Journal*, in press (arXiv:2408.11770).

Relevance. The SH0ES side of the subsample dispute: on the same JWST fields the combined three-method result is 72.6 ± 2.0 , with the full-sample prediction for those galaxies 72.8 — the CCHP–SH0ES split localized to sample selection, not measurement. First cited in §34.

Riess, A. G., Li, S., Anand, G. S., Yuan, W., Breuval, L., et al. (2025). The Perfect Host: JWST Cepheid observations in a background-free SN Ia host confirm no bias in Hubble-constant measurements. *Astrophysical Journal Letters*, 992, L34. arXiv:2509.01667.

Relevance. The sharpest ladder-internal robustness test: a background-free Cepheid field returning 73.49 ± 0.93 — evidence that whatever moves the ladder is not photometric, quoted in the counterweight clause where it is first cited. First cited in §34.5.

Said, K., Howlett, C., Davis, T., et al. (2025). DESI peculiar velocity survey — Fundamental plane. *Monthly Notices of the Royal Astronomical Society*, 539, 3627.

Relevance. The DESI fundamental-plane calibration entering the H0DN variant catalog (V03) and the O1 orthogonal path. One of the orthogonal-path legs of the H0DN consistency architecture, whose scope — internal covariance, not external truth — the certification-boundary analysis of Courses D and L delimits. First cited in §12.

Sandage, A., & Tammann, G. A. (1975). Steps toward the Hubble constant. V. *Astrophysical Journal*, 196, 313.

Relevance. The low camp of the factor-two war, primary source; the historical stalemate the century course narrates. The factor-two war is this paper’s standing reminder that confident error bars have excluded the next era’s value before — the cautionary-precedents lesson of Part VIII, one century early. First cited in §2.

Sandage, A. (1958). Current problems in the extragalactic distance scale. *Astrophysical Journal*, 127, 513.

Relevance. The second great halving: brightest “stars” recognized as HII regions, and the distance scale doubled again — the correction that brought H_0 near its modern range. First cited in §2.

Sandage, A. (1962). The change of redshift and apparent luminosity of galaxies due to the deceleration of selected expanding universes. *Astrophysical Journal*, 136, 319.

Relevance. The original redshift-drift proposal — the expansion watched, not inferred; the purest channel in the over-determination program. First cited in §59.

Sandage, A. (1970). Cosmology: A search for two numbers. *Physics Today*, 23(2), 34.

Relevance. The famous framing of pre-1998 cosmology as the measurement of H_0 and q_0 — the expectation the supernova teams were built to satisfy, and overturned. First cited in §9.

Schmidt, B. P., Suntzeff, N. B., Phillips, M. M., et al. (1998). The High-Z Supernova Search: Measuring cosmic deceleration and global curvature of the Universe using Type Ia supernovae. *Astrophysical Journal*, 507, 46.

Relevance. The High- z team’s method paper — the search designed to measure deceleration that found acceleration. The method paper behind the discovery pair: the pipeline ancestor of every supernova standardization whose dust and population choices the residual anatomy now prices. First cited in §9.

Schöneberg, N. (2024). The 2024 BBN baryon abundance update. *JCAP* (arXiv:2401.15054).

Relevance. The deuterium accountant: $\omega_b = 0.02218 \pm 0.00055$ marginalized over reaction rates, with no microwave sky anywhere — the audit that both anchors the CMB-free ruler and walls off the early-side fixes. First cited in §43.3.

Schutz, B. F. (1986). Determining the Hubble constant from gravitational wave observations. *Nature*, 323, 310.

Relevance. The founding proposal: compact-binary waveforms as self-calibrating distance indicators — the forty-year-old idea on which this paper stakes its verdict. First cited in §64.

Scolnic, D., Brout, D., Carr, A., et al. (2022). The Pantheon+ analysis: The full data set and light-curve release. *Astrophysical Journal*, 938, 113.

Relevance. The Pantheon+ SN Ia compilation; the deep Hubble-flow rung of the network and part of the calibration-debate context retired by variant V13. First cited in §10.

Scolnic, D., Riess, A. G., Wu, J., Li, S., Anand, G. S., et al. (2023). CATS: The Hubble constant from standardized TRGB and Type Ia supernova measurements. *Astrophysical Journal Letters*, 954, L31. arXiv:2304.06693.

Relevance. The standardized-TRGB grid: 73.22 ± 2.06 baseline with none of its 108 analysis variants below 71.6 — the second ladder-internal robustness result in the counterweight clause. First cited in §34.5.

Scolnic, D., et al. (2024). The Hubble tension in our own backyard: DESI’s fundamental plane and the Coma cluster. *Astrophysical Journal Letters*, 979, L9 (arXiv:2409.14546).

Relevance. The highest entry of the landscape: 76.5 ± 2.2 from the DESI fundamental plane anchored on Coma — whose distance rests on twelve Type-Ia supernovae, so the route is not independent of the standardization question it is sometimes cited against. First cited in §43.4.

Slipher, V. M. (1917). Nebulae. *Proceedings of the American Philosophical Society*, 56, 403.

Relevance. The systematic redshifts of the spirals — the first observational hint, years before the law, that the universe is in motion. First cited in §1.

Spergel, D. N., Verde, L., Peiris, H. V., et al. (2003). First-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters. *Astrophysical Journal Supplement*, 148, 175.

Relevance. The WMAP concordance fit that made Λ CDM the standard model of cosmology — the six-parameter framework whose H_0 inference the tension now interrogates. First cited in §9.

Stiskalek, R., Desmond, H., & Banik, I. (2025). Testing the local supervoid solution to the Hubble tension with direct distance tracers. *Monthly Notices of the Royal Astronomical Society*, 543, 1556.

Relevance. Field-level forward modeling of the CosmicFlows-4 velocities: the data prefer voids smaller than ~ 70 Mpc, and the deep profiles that would help H_0 most contradict the measured bulk flow — the result that downgrades the KBC-deep scenario from contested to disfavored in the hole ledger. First cited in §33.1.

Stiskalek, R., Riess, A. G., Desmond, H., Lavaux, G., & Scolnic, D. (2026). Forward-modeling Milky Way Cepheids: selection effects and physical priors in the Gaia–HST calibration. *Monthly Notices of the Royal Astronomical Society*, in press (arXiv:2603.09880).

Relevance. The rebuttal inside the statistics channel: a uniform-in-volume prior *without* a selection model manufactures a ~ 0.05 mag zero-point counter-bias, and the consistent treatment reinforces the published H_0 — the reason the channel is carried with a zero-edge range. First cited in §34.3.

Verde, L., Schöneberg, N., & Gil-Marín, H. (2024). A tale of many H_0 . *Annual Review of Astronomy and Astrophysics*, 62, 287.

Relevance. The review whose many-measurements framing the H0DN network operationalizes; quoted in Part I for the decade-long scrutiny of the tension. First cited in §11.

Vincenzi, M., Brout, D., Armstrong, P., Popovic, B., Taylor, G., et al. (DES Collaboration) (2024). The Dark Energy Survey Supernova Program: Cosmological analysis and systematic uncertainties in the DES-SN5YR data sample. *Astrophysical Journal*, 975, 86.

Relevance. The dust channel’s principal counterweight: the five-year systematics analysis argues in its error budget that part of the host-mass step is intrinsic rather than dust-driven — cited with a body-level pointer, since the abstract is silent on the step. First cited in §34.3.

Vogl, C., Taubenberger, S., Csörnyei, G., et al. (2025). No rung left behind: A cosmological distance ladder free of Type Ia supernovae. *Astronomy & Astrophysics*, 702, A41. arXiv:2411.04968.

Relevance. The SN II expanding-photosphere ladder free of SNe Ia; the EPM variant class (V05) of the H0DN catalog. Supplies the landscape’s SNe II entry (74.9 ± 1.9 , statistical only): one of the geometric counterweight’s members, carried at full strength and with its stated qualification, in Part V and O-Hubble-53. First cited in §12.

Wang, B., López-Corredoira, M., & Wei, J.-J. (2024). The Hubble Tension Survey: 216 measurements, one field. *Monthly Notices of the Royal Astronomical Society*, 527, 7692 (arXiv:2311.18443).

Relevance. The scatter audit: the literature’s spread exceeds its quoted errors, and the headline five-sigma is nearer three in equivalent-frequency terms — the reason the landscape is read in bands, not rows. First cited in §43.1.

Watkins, R., Allen, T., Bradford, C. J., Ramon, A., Walker, A., et al. (2023). Analyzing the large-scale bulk flow using CosmicFlows-4: increasing tension with the standard cosmological model. *Monthly Notices of the Royal Astronomical Society*, 524, 1885.

Relevance. The contested companion evidence in the hole ledger: a bulk flow of $419 \pm 36 \text{ km s}^{-1}$ at $200 h^{-1}$ Mpc, far above the Λ CDM expectation — carried as contested, with selection and estimator critiques in play. First cited in §33.1.

Weinberg, S. (1989). The cosmological constant problem. *Reviews of Modern Physics*, 61, 1–23.

Relevance. The definitive statement of the vacuum-energy problem that Course F’s seesaw addresses; the classic accounting of the 10^{120} . First cited in §F.1.

Weinberg, S. (2008). *Cosmology*. Oxford University Press.

Relevance. Graduate reference for the Friedmann background and the acoustic-scale physics behind the two physical densities ω_b, ω_m (App. E.7). First cited in §21.1.

Whitbourn, J. R., & Shanks, T. (2014). The local hole revealed by galaxy counts and redshifts. *Monthly Notices of the Royal Astronomical Society*, 437, 2146.

Relevance. Regional count-and-redshift depths $\delta \approx -0.15$ to -0.24 over 150–250 Mpc; the counts-branch evidence in the void-channel ledger, recorded and not carried. First cited in §33.1.

Whitford, A. M., et al. (2023). Evaluating bulk-flow estimators for CosmicFlows-4. *Monthly Notices of the Royal Astronomical Society* (arXiv:2306.11269).

Relevance. The estimator audit of the CF4 bulk-flow excess: $428 \pm 108 \text{ km s}^{-1}$ under the minimum-variance estimator with uncertainties treated fully — the tension with Λ CDM softer than the headline, which is how the hole ledger carries it. First cited in §50.

Wiltshire, D. L. (2007). Cosmic clocks, cosmic variance and cosmic averages. *New J. Phys.* **9**, 377.

Relevance. The timescape proposal — the one non-linear, two-phase reading that claims percent-level backreaction; discussed as the honest loophole (and its cost) in §35. First cited in §35.

Wojtak, R., & Hjorth, J. (2024). Consistent extinction model for type Ia supernovae in Cepheid-based calibration galaxies and its impact on H_0 . *Monthly Notices of the Royal Astronomical Society*, 533, 2319. arXiv:2403.10388.

Relevance. The dust channel’s opening move: imposing a Milky-Way-like extinction model on the calibrator galaxies fits the data better ($\Delta\text{BIC} = -11$) and lowers the ladder to 70.5 ± 1.0 — a coherent $\approx -2.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$ shift, about half the gap. First cited in §34.3.

Wojtak, R., & Hjorth, J. (2025). Stretch to stretch, dust to dust: lower-value local H_0 measurement from two-population modeling of type Ia supernovae. *Astronomy & Astrophysics*, 702, A176. arXiv:2506.22150.

Relevance. The dust channel in full: a two-population Bayesian hierarchical remodeling of Pantheon+, driven by the observed stretch bimodality, returns $H_0 = 70.59 \pm 1.15$ — the central value the two-scenario arithmetic of §34.5 adopts for the channel. First cited in §34.3.

Wu, J., Scolnic, D., Riess, A. G., Anand, G. S., Beaton, R., et al. (2023). Comparative analysis of TRGBs (CATs) from unsupervised, multi-halo-field measurements: Contrast is key. *Astrophysical Journal*, 954, 87.

Relevance. The tip-contrast relation ($\sim 5\sigma$): the TRGB’s measured tip varies field to field at the 0.05–0.1 mag level unless standardized — the convention-level lever inside the subsample channel. First cited in §34.

Zaborowski, E., et al. (DESI Collaboration) (2024). A sound-horizon-free determination of the Hubble constant from DESI. *JCAP*, 06, 020 (arXiv:2411.16677).

Relevance. The equality-scale route: $66.7^{+1.7}_{-1.9}$ (DESY5), $67.9^{+1.9}_{-2.1}$ (Pantheon+), $67.8^{+2.0}_{-2.2}$ (Union3) with no sound horizon anywhere — the low half of the ruler-free span that closes “shorten the ruler” as a road to 73. First cited in §43.3.