

The Kaxion

Cold Dark Matter as the Non-Spatial Compact-Angle Pseudoscalar of GG-6

Katsushi Arisaka

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

Contact: arisaka@physics.ucla.edu

Abstract

Nearly a century ago, astronomers found that galaxies move as though they hold far more matter than we can see. Stars at the edges of spinning galaxies orbit too fast; the afterglow of the Big Bang bears the imprint of a substance outweighing ordinary atoms five to one. We call it dark matter, and no experiment has yet identified what it is.

This paper offers an answer not invented for the purpose. In earlier papers of this series we built a theory from a few founding rules and showed that, with no adjustable numbers, it reproduces the known particles and forces. That same construction is forced to contain one more field: an exceedingly light, electrically neutral relic of the geometry we call the *Kaxion*. It is automatically dark, stable, and cold, and the early universe makes it by the standard mechanism; requiring it to make just the observed amount — about 27 percent of the cosmos — fixes the one scale the geometry has not yet supplied. It also helps relax the strong nuclear force’s vacuum angle to zero — a job the construction was not built to do.

With that one scale set and nothing else free to tune, the theory is committed to three things: the mass, $8.1 \mu\text{eV}$, and with it the microwave frequency at which to listen, near 1.95 GHz, inside a narrow band; the faintness of its glow into light; and a count — one note, standing alone inside that band, with no partner tone at any spacing. Two programs aimed at this band — one running, one designed — would together reach this frequency and faintness, and would then either find the Kaxion or close the window. If they do, the substance that first gathered the galaxies will have been sounding one faint note since the universe’s opening microsecond — a pitch set by one calibration and a handful of whole numbers.

Executive Summary

The dark-matter problem is usually posed as a menu: measure the missing mass, then audition candidate particles invented for the occasion, each carrying free parameters tuned afterwards to fit. This series inverts that order. It begins from a short constitution — four axioms and seven postulates (Arisaka, 2026, A1-Constitution) — and derives a six-dimensional stage whose two extra dimensions are not places one could travel to but a *scale* and an *angle*: pure numbers rather than directions (Arisaka, 2026, A2-GG6). Consistency alone then fixes the particle content and a ledger of five whole numbers, with not one continuous dial left to turn (Arisaka, 2026, A3-BI6); the known particles and forces, and the standard account of cosmology, follow as shadows this stage casts on its four-dimensional boundary (Arisaka, 2026, A9-SM; Arisaka, 2026, A10-Cosmos). The consequence is unusual: by the time one asks what the dark matter is, the theory *already* contains a forced candidate and *already* fixes how much of it there should be. This paper asks whether the one accounts for the other. It does — and the candidate, the *Kaxion*, arrives with a single frequency to listen at and two fingerprints to confirm it by.

What the Kaxion is. The extra “angle” dimension is a circle that is not a place but a phase — a dial with no preferred setting. By an elementary fact of field theory, a field living on such a dial is an extremely light particle with essentially no mass of its own. Quantum consistency forces exactly such a field into the theory, and the process that selects the theory’s ground state settles it into the deepest, most stable valley of the energy landscape. That settled field is the Kaxion. The decisive point is that it is not added to explain dark matter: it is already present, and obligatory, before the dark-matter question is ever posed. The constitution’s guiding rule — *maximum symmetry, minimum freedom* — forbids inventing a new sector when one the theory already owns will serve.

Why it is dark, and why it is cold. The Kaxion carries no electric charge and no color, and every one of its interactions is divided by a single enormous energy scale, so it is very nearly inert and effectively immortal. It is never cooked up in a hot thermal bath. Instead, in the young universe the field sits pinned by the friction of cosmic expansion, held away from the bottom of its potential; once the expansion slows enough, it is released and begins to oscillate, and from that instant its density thins out exactly as ordinary matter does. A field oscillating in step with itself is a calm, cold condensate, not a hot gas: its internal motions are negligible. The Kaxion is therefore *cold* dark matter of precisely the kind known to build galaxies correctly — not an ultralight “fuzzy” field that would smear their fine structure away.

Why the amount comes out right. In the cosmology paper of the series (Arisaka, 2026, A10-Cosmos) the present contents of the universe are not fitted but locked to the number π : the total matter fraction is $1/\pi$, the ordinary-atom fraction $1/(2\pi^2)$, and their difference,

$$\Omega_{\text{DM}} = \frac{1}{\pi} - \frac{1}{2\pi^2} \simeq 0.268,$$

is the dark-matter share — matching what is measured in the Big Bang’s afterglow. That is the target the Kaxion must hit. The observed abundance calibrates the one scale the geometry has not yet supplied, at an utterly ordinary value of the field’s starting tilt, which is then read back rather than chosen. The dark-matter density is therefore not a prediction of this paper but its one calibration, declared as such; what the geometry fixes is the ladder of pitches around that scale and the count.

The one frequency, and the faintness of its glow. The constitution also hands the Kaxion a second job — to relax the strong nuclear force’s vacuum angle to zero, the relaxation half of the mechanism by which the series explains why that force, too, respects mirror symmetry so exactly. Because it does this job, the Kaxion inherits a precise tie to the strong interaction, and a single well-measured nuclear quantity then fixes its mass: $8.1\ \mu\text{eV}$, a microwave tone at 1.95 GHz. The same bookkeeping fixes how feebly it converts into light, placing it exactly on the *DFSZ line* — among the faintest and hardest targets the axion-search community has set. At the calibrated scale, every uncertainty in these numbers traces to laboratory nuclear measurement; the placement of the scale itself is the calibration.

Two fingerprints that cannot be tuned away. Whatever its exact mass, the geometry stamps the Kaxion with two signatures that carry no adjustable number beyond the one calibrated scale. First, its note must fall inside one narrow band of frequencies — a short ladder of whole-number steps around the calibrated pitch, the ladder read from the geometry’s integers and stated at the grade it earns. Second, the note stands *alone*: the geometry supplies exactly one light field of this kind, so a second note inside the band would refute the count as this paper states it — and tell us which of the parents’ registered modes is light — while an empty band would refute the placement (Part VI states this precisely, with the conditional the count carries). Either answer is the geometry read directly, in a laboratory.

Why no rival survives — inside this theory. The same constitution that produces the Kaxion forecloses every alternative, each by a rule adopted for independent reasons: no supersymmetry, hence no WIMP; strong CP already solved, hence no room for a second, separate axion; the particle content and cosmic history fixed, hence no sterile neutrino, dark photon, or primordial black holes to spare. The laboratory has been reaching the same verdict on its own — direct searches find no WIMP and are entering the glare of solar neutrinos, colliders find no superpartners, and astronomical data have cornered the sterile neutrino and the fuzzy field. The candidates the data are quietly retiring are exactly the ones this theory forbids. The Kaxion is what remains: required by consistency, and sufficient for the budget.

What it accounts for in the sky. Two recent astronomical surprises point the same way. The *James Webb* Space Telescope finds galaxies assembling earlier and more massive than expected — a demand for *more* early structure, which counts against any candidate that would suppress it, while the cold Kaxion clears those limits with room to spare and inherits, at most, the modest

collapse-threshold head start its cosmology companion carries as a working value. And the long-standing puzzle of dwarf-galaxy centers is met honestly: the Kaxion’s one faint lever is far too weak to reshape them, so it predicts ordinary galaxy cores — a closed prediction, not a free knob. Astronomy sets the Kaxion’s stage; only the laboratory can identify it.

How, and when, it can be tested. This is not a prediction for some distant generation of instruments. The Kaxion’s note near 1.95 GHz falls inside the band that two well-advanced programs — IBS-CAPP in Korea and ADMX’s extended-range search at Fermilab — are being built to sweep, at the very faintness the Kaxion requires, over the next several years to a decade. A receiver covering the band completely, checking as it goes that any note it finds stands alone, would turn a bare detection into an identification — the fence, the address on the standard model line, and the count together. The test cuts both ways: the Kaxion can be *found*, at a stated frequency with stated fingerprints, or *killed* by stated measurements — with every criterion written down in advance.

What is proved, and what is still assumed. We keep an explicit ledger of proved versus assumed throughout, because the honesty is structural: it tells a reader which measurement tests a theorem and which tests a working assumption. Established in the earlier papers of the series and inherited here: the field’s existence and its origin in the tensor sector, its coupling to the strong force, and the coldness of the condensate. Resting on named, clearly flagged assumptions: the calibration that fixes the bulk normalization through the π -locked inventory, the ladder that pins the laboratory frequency within its discrete window, the inflation-era pinning of the angle that makes a single misalignment angle meaningful, and the count — exactly one light field of this kind, so the note stands alone — which is conditional on five groups of modes the parent papers register with open masses. Each is stated as what it is wherever it is used, and each is testable.

A closing word on why the answer matters. Dark matter is not a curiosity at the margin of the cosmic account; it is five-sixths of all the matter there is, and it was its patient, invisible gravity that first gathered gas into galaxies early enough for stars to ignite, forge the carbon and oxygen we are made of, and seed the worlds on which questions like this one came to be asked. To identify it is to learn something about the origin of our own existence. If the Kaxion is real, the substance that built us has been sounding a single faint note since the first microsecond of the universe — a note whose ladder and whose solitude are fixed by five whole numbers and the geometry of a dimension that is not a place but an angle, and whose pitch is set once by the amount of dark matter there is. We now know the pitch to listen for, and we will soon know whether the universe is playing it. Either way, the question will at last be answered the way physics answers questions: not by preference, but by measurement.

Keywords: GG-Theory; GG-6; BI-6; OGN-5 integer ledger (1, 3, 5; 7, 35); GRIP; non-spatial compact angle; Green-Schwarz mechanism; Kaxion; dark matter; strong CP; $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$; misalignment; haloscope; single line; fenced window [1.59, 2.13] GHz; ADMX; HAYSTAC.

One Predicted Line Against Every Published Search — the Kaxion at $8.1 \mu\text{eV}$, on the DFSZ Line




Figure 1: The Kaxion against the haloscope landscape: the published exclusion envelopes (red) of the ADMX, CAPP, RBF/UF, HAYSTAC, TASEH/CAST-CAPP, QUAX and ORGAN campaigns, digitized from the community AxionLimits compilation (O’Hare, 2020); the KSVZ and DFSZ model lines; and the one prediction — the Kaxion (red star) at $8.1 \mu\text{eV}$ on the DFSZ line ($E/N = 8/3$), inside the closed window of the seven loadings $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$ (orange), every other frequency a mandated silence. The black star marks the 1.036 GHz candidate excluded in early 2026 — F-Kaxion-1d was faced and the prediction survived. Three published searches reach inside the window and are labeled; not one is within a factor of ten of the DFSZ line. Two announced programs are drawn dashed and only as far as each states: ADMX-EFR on the DFSZ line it targets across $2-4 \text{ GHz}$, and ADMX’s four-cavity array — a stated mass range and no stated sensitivity — as a coverage bracket above the plane. Full version and sources: Figure 8 in Part VII.

From Five Integers to Four Falsifiable Numbers — the Derivation Flow of the Kaxion




Figure 2: Derivation flow of the Kaxion: from the constitution to four falsifiable numbers. Each downward step adds no free parameter beyond the one calibrated scale: the constitution forces the 6D bulk and the BI-6 anomaly, which carry the five-integer OGN-5 ledger (1, 3, 5; 7, 35); the GDOS boundary law and the GRIP runtime select and stabilize the deepest vacuum — the Ground Geometric Attractor — and with it the Kaxion. The four key quantities then follow from named upstream ingredients: the **mass** from the π -locked abundance and the gluonic portal $m_K f_K = \sqrt{\chi_{\text{QCD}}}$; the **photon coupling** from $E/N = 8/3$ on the DFSZ line; the **window** from the ladder ansatz around the calibrated benchmark, anchored at the throat depth $K = 35$ through the discrete cap ladder; and the **line count** — single — from the mode analysis of the anomaly companion. The dashed link records that the same decay constant f_K fixes both the mass and the photon coupling. Full version: Figure 7 in Part V.

Nine Numbers, and the Measurement That Can Kill Each One

The paper, stripped of narrative. Its content is nine numbers and the machinery that stands behind them — each with its value, its status, the Part in which it is proved, and the measurement that can end it. The status column is the honest one: *calibrated* where one scale is set by the observed abundance and declared as such; *derived* where a theorem is combined with a Standard-Model input; *derived-conditional* where the derivation rests on a named open computation; *ladder* for the discrete window read from the geometry’s integers by an ansatz; *fitted* for the angle read back from the abundance. Nothing in the last column waits on future theory; all of it waits on data, and two programs — IBS-CAPP in Korea, running, and ADMX-EFR at Fermilab, designed and awaiting a funding decision — would take it between now and the early 2030s. In full: Table 35, p. 195.

Quantity	Value	Status	Proved	Killed by
1 mass m_K	$8.1 \mu\text{eV} (\pm 1.3\%)$	calibrated	(via Part V	F-Kaxion-1c/1d
2 frequency ν_K	$1.95 \text{ GHz} (\pm 25 \text{ MHz})$	f_K calibrated	(via Part V	F-Kaxion-1c/1d
3 window	$[1.59, 2.13] \text{ GHz}$	discrete ladder	Part V	F-Kaxion-1c
4 line count	single (no partner tone)	derived-cond. (T-Kaxion-5)	Part VI	(count test, Part VI)
5 coupling $g_{K\gamma\gamma}$	$(1.2 \pm 0.07) \times 10^{-15} \text{ GeV}^{-1}$	derived (DFSZ; via f_K)	Part V	F-Kaxion-1c
6 decay constant f_K	$7.1 \times 10^{11} \text{ GeV}$	calibrated; bench. $\Xi_{\text{cap}}=43$	Part V	(via 1–3)
7 misalignment angle θ_i	$1.18 \pm 0.04 \text{ rad}$	the convention’s output	Part V	(via 8)
8 abundance Ω_{DM}	$1/\pi - 1/(2\pi^2) = 0.2676$	derived (GG- A10-Cosmos, cond.)	Part III	CMB
9 lifetime τ_K	$\simeq 4 \times 10^{32} t_U$	derived	Part V	(stability)

Overview of Contents

Part I. Introduction	20
Part II. What is the Kaxion?	30
Part III. Cosmology of the Kaxion	62
Part IV. Context	79
Part V. Derivation of the basic quantities	97
Part VI. Unique fingerprints	120
Part VII. Detection	124
Part VIII. Future detection	144
Part IX. Simulating the Discovery	148
Part X. The KAXIO Proposal	161
Part XI. Discussion	179
Part XII. Conclusions	194
Part XIII. Technical Appendices	199
Part XIV. Common Objections and Responses	258
Part XV. References	280

List of Figures

1	The Kaxion against the haloscope landscape (showcase; full version in Part VII).	5
2	Derivation flow of the Kaxion (showcase; full version in Part V).	6
3	The open-system funnel shared across the GG-Theory cohort.	23
4	The twenty-two-paper cohort map.	26
5	Roadmap of the fifteen Parts of GG-C1-Kaxion.	28
6	The Kaxion on the two fibers: identity on the circle, occupancy on the interval.	58
7	Derivation flow: from the constitution to the four falsifiable numbers.	98
8	The Kaxion against the published haloscope landscape, with window and DFSZ front.	132
9	SNR versus integration time for a camped search, by instrument class.	139
10	Kill-test timeline: the survived 2026 test, the scan fronts, and the kill date.	147
11	Flow chart of the discovery simulation: inputs, modules, outputs, budget.	149
12	Synthetic camped spectra with expected signal and systematic envelope.	154
13	Annual and daily modulation of the line with statistical and systematic errors.	155
14	Single-mode acceptance and the count test, with the Q_L systematic band.	156
15	Conceptual design of KAXIO: two independently tuned sub-arrays.	163
16	The KAXIO RF and readout chain, common to both data streams.	165
17	Daily modulation with KAXIO: one month and one year, with statistical errors.	168
18	One year of KAXIO: the annual modulation as a precision measurement.	169
19	The verdict clock: what KAXIO knows at a week, a month, three months, a year.	170
20	The KAXIO exclusion reach across the window on the null tape.	172
21	Hierarchy of the program-defined names, from the framework to the Kaxion.	203

22	The archived two-lock simulation: four peaks and the booster path.	257
----	--	-----

List of Tables

1	The constitution: four axioms and seven postulates, and the role of each here.	31
2	What the same five integers have already bought.	34
3	The five-link existence chain.	45
4	The three non-spatial mode families and their division of labor.	53
5	The triptych: three Ground Geometric Attractors, one anatomy.	60
6	The Kaxion’s cosmic timeline (onset numbers from Section 13).	64
7	The π -locked inventory versus Planck 2018.	71
8	The dial’s swing amplitude through history (local-halo value in the last row).	76
9	Structural scorecard.	83
10	The constitutional verdict: what each rival requires, and what forbids it.	86
11	The discrete cap window: each admissible loading fixes one decay constant.	106
12	Error budget for the four headline predictions.	115
13	Theorem-versus-conjecture ledger for every load-bearing claim of Parts II–V.	117
14	Why no adjustable competitor can reproduce the signature.	130
15	The experimental landscape against the Kaxion window.	131
16	The routes into the Kaxion window, with realistic entry dates.	136
17	Camping on the benchmark: power, photon rate, and time to 5σ by class.	138
18	The search program.	145
19	Simulation inputs I: the frozen theory sheet.	149
20	Simulation inputs II: the standard-halo-model block with assigned systematics.	150
21	Simulation inputs III: the instrument model with engineering justifications.	151
22	The worked signal-power budget, factor by factor with uncertainties.	151
23	Time to a five-sigma detection per instrument class, with systematic range.	152
24	Modulation observables: periods, amplitudes, statistical reach, systematics.	155
25	Single-mode suppression of an in-window second line versus loaded quality factor.	156
26	The master systematic-error budget of the discovery simulation.	157
27	KAXIO design requirements, each traced to its origin and falsifier.	162
28	KAXIO design specification by subsystem.	165
29	KAXIO sensitivity and noise budget at the benchmark.	166
30	KAXIO-specific systematic errors, extending the Part-IX budget.	167
31	The staged per-rung campaign across the discrete window.	171
32	The KAXIO four-step run plan with deliverables on both tapes.	173
34	KAXIO cost classes: reused, modified, and new elements.	175
33	The KAXIO risk register with likelihood, impact, and mitigations.	175
35	The paper in nine numbers: value, status, proof location, and falsifier for each.	195
36	Summary of abbreviations.	200
37	Glossary of frequently used symbols.	203

38	Every headline quantity: symbol, value, status, and the course that proves it.	206
39	End-to-end audit.	235
40	What each Kaxion number inherits, and from where.	237
41	The claims of the paper, graded.	238
42	The open items, the single computation that promotes each, and the payoff.	239
43	The baryonic port: two Larmor reference rows and three non-Larmor concepts.	244
44	The harmonic-booster coupling and noise budget, with the staged readout path.	247
45	The notation of this paper at a glance, in two columns of symbol and meaning.	256

Contents

Executive Summary	2
List of Figures	8
List of Tables	9
Preface	19
Part I. Introduction	20
1 The dark-matter problem	20
1.1 A brief history of the dark-matter problem	20
1.2 The astronomical evidence	21
1.3 The detection challenge	21
2 The central claim: the Kaxion dark-matter mass as an open-system observable . . .	22
3 Three reading paths	22
4 The cohort inheritance and the roadmap of the fifteen Parts	23
4.1 The cohort inheritance	24
4.2 Roadmap of the fifteen Parts	27
5 Part I in plain English: what the dark-matter problem is, and what would count as solving it	29
Part II. What is the Kaxion?	30
6 The GG-Theory framework, reviewed for newcomers	30
6.1 The constitutional method	30
6.2 GG-6: the six-dimensional stage	32
6.2.1 The manifold and the metric	32
6.2.2 The gauge group, and why baryon number is gauged	33
6.2.3 The OGN-5 integer cascade	33
6.2.4 The track record	34
6.2.5 The energy ladder, numerically	35
6.3 BI-6 and the forced tensor sector	36
6.3.1 A two-paragraph anomaly primer	36
6.3.2 The rank-one criterion and the locked coefficients	36
6.3.3 How the inflow cancels the anomaly: descent in three lines . . .	37
6.4 GDOS: from the bulk to the boundary observer	38
6.4.1 Why a projection is constitutional	38
6.4.2 A one-paragraph open-systems primer	38
6.4.3 What GDOS proves, and what the Kaxion takes from it	38
6.4.4 The boundary law in its three dress codes, displayed	39
6.5 GRIP: the runtime that selects vacua	40
6.5.1 Why a Lagrangian is not enough	40
6.5.2 A Morse-theory primer in one figure's worth of words	40
6.5.3 The landscape: a free energy, not an energy	41
6.5.4 The cascade and the uniqueness theorem	41
6.5.5 From GA to GGA: the ground-state refinement	42
6.5.6 The two regimes: GRIP-In and GRIP-Out	42
6.6 The tour in one paragraph	43
7 The five-link existence chain	43
7.1 Link 1: the stage provides the circle	43
7.2 Link 2: the ledger provides the gears	43
7.3 Link 3: anomaly freedom provides the tensor	44

7.4	Link 4: the boundary provides its pseudoscalars — and exactly one stays light	44
7.5	Link 5: the runtime provides occupancy	44
8	The Kaxion existence theorem	45
9	The clock and the lifted dial: constructing the light mode	47
9.1	The tensor remnant: dualizing B_2	47
9.2	The holonomy echo: the Wilson line as a field	48
9.3	The kinetic matrix and the light eigenstate	49
9.4	The registered exposures of the parents	50
9.5	Why the Kaxion is light: protection, quantified	51
9.6	The zero-mode among three towers	52
9.7	The KK in the name	53
10	Born, selected, stabilized: the GRIP cascade on the compact-angle stratum	53
10.1	The stratum and its landscape	53
10.2	The cascade, run explicitly	54
10.3	The landscape, quantified	54
10.4	The Kaxion as Ground Geometric Attractor	54
10.5	The Kaxion on the two fibers: matter of the circle, occupancy of the interval	56
10.6	Birth, selection, stabilization	58
10.7	Why the Big Bang cannot fail to make Kaxions	59
10.8	What Part II has established	60
11	Part II in plain English: why the geometry has no choice but to produce a light pseudoscalar	61
 Part III. Cosmology of the Kaxion		62
12	The cosmic event chain	62
12.1	One ledger, three epochs	62
12.2	The equation of motion and its two regimes	62
12.3	The timeline at a glance	64
13	Misalignment production and the initial angle	65
13.1	When does the pendulum let go? The onset temperature, twice	65
13.2	The exact solution: Bessel functions and the envelope law	66
13.3	The misalignment angle	67
13.4	How violently is the angle tested?	67
13.5	The isocurvature question, quantified	68
14	The decay constant and the misalignment angle	69
15	Abundance and the π -locked inventory	70
15.1	The target is inherited, derived-conditionally	70
15.2	The misalignment abundance, derived	71
15.3	The closure: from the π -chart to the misalignment angle	73
15.4	What the condensate does to the rest of cosmology	73
16	Why the Kaxion is cold	74
16.1	Pressure: the oscillating condensate is dust	74
16.2	Velocity: the coldest matter in the universe	74
16.3	Wavelength: cold waves, not fuzzy waves	75
16.4	The condensate's diary: the amplitude through cosmic history	75
16.5	Wave pressure, quantified: the Jeans scale	76
16.6	Verdict	77
17	Interactions with matter	77
17.1	The portal sheet, previewed	77
17.2	Why every recoil experiment is blind	77
17.3	Why resonant conversion works anyway	77

17.4	What Part III has established	78
18	Part III in plain English: where the dark matter came from, in one chain of events	78
Part IV. Context		79
19	The candidate landscape: empirical status and structural comparison	79
19.1	The empirical status as of 2026	79
19.2	The structural comparison	81
20	The GG-Theory verdict: within the framework, the Kaxion is required and sufficient	84
20.1	Required: the Kaxion must exist, and it must suffice	84
20.2	Excluded: what each rival needs that the constitution forbids	84
20.3	The shape of the claim, stated honestly	86
20.4	The message to the community	87
21	Astrophysical significance I: early, massive structure in the JWST era	87
21.1	The puzzle, in brief	87
21.2	The JWST surprise, quantified	88
21.3	Status in 2026: attenuated, not dissolved	88
21.4	Direction matters: the surprise excludes the suppressing candidates	89
21.5	The Kaxion clears the necessary condition by sixteen orders of magnitude	89
21.6	The head start on offer: what the cosmology companion still carries	89
21.7	Honest fences on the early-structure claim	90
21.8	Where the Kaxion stands on early structure	90
21.9	The not-fuzzy fence, and the one early-structure contingency	90
21.10	Lever one: minicluster seeds (post-inflationary contingency)	90
21.11	Closing: where the Kaxion stands on early structure	91
22	Astrophysical significance II: the cusp–core problem and the diversity of rotation curves	91
22.1	The problem, quantified	92
22.2	The standard solutions and their limits	92
22.3	The Kaxion is cold, and is not fuzzy	93
22.4	The Kaxion is not self-interacting dark matter	93
22.5	The one candidate lever, and why the termination closes it	93
22.6	The magnitude, now computed: the portal does not core	93
22.7	What the Kaxion predicts for galaxy centers	94
22.8	Where the Kaxion stands on cusp–core	94
22.9	What Part IV has established	95
23	Part IV in plain English: how the candidate compares, and what it would take to tell them apart	96
Part V. Derivation of the basic quantities		97
24	Strong-CP alignment and the legitimacy of the QCD susceptibility	97
24.1	The constitutional assignment	97
24.2	The gluonic portal theorem	97
24.3	Relaxation, residual, and the quality problem	99
24.4	Consequence: the susceptibility owns the mass	100
24.5	Where the susceptibility number comes from	100
25	Calibrating the bulk normalization through the π -lock	101
25.1	Ground rules for this derivation	101
25.2	The pipeline, assembled end to end	101
25.3	The pipeline in closed form	102
25.4	Cross-validation against the precision literature	103
25.5	The normalization, cross-checked against the series	104

25.6	The closure	104
26	The mass and the laboratory line	105
26.1	The mass from the single decay constant	105
26.2	The decay constant and its residual window	105
26.3	The compact-angle mass-uniqueness conjecture, stated and derived	106
26.4	The cap index: from benchmark toward theorem	110
27	The photon coupling, factor by factor	111
27.1	The exact matching formula	111
27.2	$E/N = 8/3$: the ledger fixes the ratio	111
27.3	Numbers, with errors	112
27.4	The Kaxion sits on the DFSZ line	112
27.5	The baryonic entry, and why it opens no laboratory channel	113
28	Abundance recheck and lifetime	113
28.1	Closing the loop	113
28.2	Lifetime: stability derived, not assumed	113
29	Error budget and the theorem-versus-conjecture ledger	114
29.1	The itemized budget	114
29.2	Why not a few percent? (And when it will be.)	116
29.3	The master ledger	117
29.4	What Part V has established	118
30	Part V in plain English: the four numbers, and where each one comes from	118
Part VI. Unique fingerprints		120
31	The single line and the count	120
31.1	The claim	120
31.2	The count test	120
32	The baryonic entry, stated once	120
33	Finite- Q morphology and falsifiers	121
33.1	What the line looks like in a real receiver	121
33.2	The line shape and the search template	121
33.3	The pre-declared falsifier	121
33.4	What a discovery sequence looks like	122
34	Part VI in plain English: one line, and the cheap test that it is alone	122
Part VII. Detection		124
35	The haloscope principle and the QGI readout	124
35.1	Resonant conversion	124
35.2	The haloscope as a measurement event of the series' own theory	124
35.3	The QGI readout	125
35.4	The response, derived end to end	125
36	Current and near-future experiments	126
36.1	The decisive simplification: the Kaxion is on the DFSZ line	126
36.2	The unique window: closing the two discrete questions	127
36.2.1	The first rejected alternative: the randomized-angle history	127
36.2.2	Why the line lies in the gigahertz window	129
36.2.3	Uniqueness, and the scorecard	129
36.3	The experimental landscape	130
36.4	IBS-CAPP: the best-situated instrument	131
36.5	ADMX-EFR: the backstop at the upper edge	134
36.6	HAYSTAC: the high-mass frontier, and a technology partner	135
36.7	The near-term targeting plan	135

36.8	The bottom line	136
36.9	The discovery calendar: F-Kaxion-1c as dates	136
37	The pre-registered single-frequency campaign: camping on the frequency	137
37.1	The trials factor collapses to the window	137
37.2	Signal-to-noise at the benchmark, instrument by instrument	138
37.3	The campaign protocol	139
37.4	The first kill test has already been faced	141
38	Part VII in plain English: how you would actually look, and how close the field already is	142
Part VIII. Future detection		144
39	The multi-mode count-test receiver	144
40	The spin couplings, and why no spin port is staged	144
41	Search strategy and roadmap	145
41.1	The kill-test timeline	146
42	Part VIII in plain English: what an instrument built for this prediction would look like	146
Part IX. Simulating the Discovery		148
43	Why a discovery can be simulated at all — and what is simulated here	148
44	The simulation pipeline and its inputs	148
45	The signal model: equations and worked numbers	150
46	The noise model and the time to detection	152
47	Results I: the camped spectra	153
48	Results II: the modulation fingerprints	153
49	Results III: single-mode acceptance, and why the count test is a deliberate act . . .	154
50	The systematic-error budget	157
51	Discovery week — and what failure looks like	157
52	Reproducibility and code release	158
53	The bottom line: why this discovery cannot be mistaken	158
54	Part IX in plain English: what discovery week would look like, written down before it happens	159
Part X. The KAXIO Proposal		161
55	Executive summary of the proposal	161
56	Design requirements	161
57	Detector design	162
58	The RF and readout chain	164
59	Sensitivity and noise	166
60	Systematic errors	166
61	Simulated performance	167
62	Staged coverage of the discrete window	170
63	The four-step run plan	172
64	Blinding, data management, and openness	173
65	Risk register	174
66	Cost class, schedule, and organization	175
67	The case for KAXIO	177
68	Part X in plain English: a specific machine, at a specific place, on a specific clock .	178

Part XI. Discussion	179
69 Summary: the Kaxion in theory and in the laboratory	179
69.1 The theory: a field the framework could not omit	179
69.2 The numbers: every observable a gear ratio of one ledger	180
69.3 One physical line, and the discipline that fixes it	180
69.4 The five credentials, derived rather than assumed	180
69.5 The experiment: one line, two fingerprints, three instruments	181
69.6 The honesty ledger	181
70 Significance and uniqueness	181
71 KAXIO: a flagship for the dark sector	184
72 Uniqueness, nobility, and innovation	185
72.1 The classic dark-matter puzzles, resolved or dissolved	185
72.2 The structural innovations	187
72.3 The nobility clause	189
73 Future directions, open items, and recommendations	189
73.1 The experimental road to a verdict	189
73.2 Cross-paper interfaces and the division of labor	190
73.3 The three named open items	190
73.4 Closing open item (iii): the dark-baryon coring magnitude	191
73.5 The path to few-percent	192
73.6 Relation to the sister predictions	192
74 Part XI in plain English: what the paper has settled, and what it has not	193
 Part XII. Conclusions	 194
75 Conclusions	194
75.1 A mystery a century old	194
75.2 The first exact answer	194
75.3 What this paper has established	195
75.4 An answer we can check — soon	196
75.5 KAXIO: a new window on the dark universe	196
75.6 The objective limitations	196
75.7 Why it matters that we find out	197
75.8 A personal note	197
 Part XIII. Technical Appendices	 199
A Reference dictionaries	200
A.1 Summary of abbreviations	200
A.2 Glossary of key symbols	203
A.3 Global naming convention for cross-cohort citation	205
A.4 Comprehensive summary tables	205
B The conceptual on-ramp: primers and foundational questions	206
B.1 Primer: Why is the Kaxion not just another axion?	206
B.2 Primer: Why exactly $8.1 \mu\text{eV}$?	207
B.3 Primer: What happens the day a line appears?	207
B.4 Primer: If the window is excluded, why can the frequency not simply be adjusted?	208
B.5 Six foundational questions about the Kaxion	208
B.5.1 What is the DFSZ line, and why does the Kaxion sit exactly on it?	208
B.5.2 Why does the Kaxion carry a baryonic ledger entry, and what became of the 5:2 readout?	210
B.5.3 What is the cap index Ξ_{cap} , and why is it $K + N + L = 43$?	212

	B.5.4	Why is the Kaxion mass so small — and why can GG-6 predict it when the axion and the WIMP cannot?	214
	B.5.5	Why is the cosmic inventory locked to π , and what becomes of the coincidence problem?	216
	B.5.6	Why is the Kaxion cold, and why are there no galaxies or stars without it?	217
C		From-scratch derivation courses	219
	C.1	Dualization and periodicity	219
	C.2	The GFF Hessian at the compact-angle attractor	221
	C.3	The QCD-coupled axion: the mass relation and the photon coupling, derived	222
	C.4	The baryonic ledger entry, derived — and the withdrawn 5:2 readout	224
	C.5	The cap index $\Xi_{\text{cap}} = K + N + L = 43$: the throat normalization and its integer	225
	C.6	Misalignment production: the relic abundance from the equation of motion	227
	C.7	The π -locked inventory: an inherited result, with the key steps reproduced	229
D		Detection, discovery, and the KAXIO instrument	231
	D.1	Haloscope sensitivity: the Sikivie power and the radiometer scan-rate wall	231
	D.2	Finite- Q transfer functions	232
	D.3	The KAXIO instrument and the one-year projection	233
	D.4	The discovery-week error budget and the verdicts	234
E		Registries, audits, and roadmaps	235
	E.1	Consolidated benchmark audit: the closure pipeline end to end	235
	E.2	Claims and falsifier registry	236
	E.3	Cross-paper inheritance table	236
	E.4	Named-theorem-and-conjecture deck	237
	E.5	Promotion roadmap for the open items	238
F		Retired readings of earlier revisions: the record	240
	F.1	The doublet and the $1/K$ split	240
	F.2	The channel ratio and the spin ports	241
	F.2.1	Non-Larmor nucleon readout: the principle	242
	F.2.2	Concept B': internal-field (hyperfine-enhanced) nuclear magnon resonance	242
	F.2.3	Concept B'': heterodyne (up-converted) nuclear readout	243
	F.2.4	Concept B''': the baryon-number force readout	243
	F.2.5	Feasibility band: signal, integration time, and the chiral coefficient	244
	F.3	The harmonic-booster module of the KAXIO design	245
	F.4	The two-lock simulation: the channel-ratio panel and the four-peak week	247
	F.5	Primer (archived): the withdrawn $1/35$ split — and why it could not renormalize away	248
	F.6	Why did earlier revisions expect twin peaks split by $1/35$? (archived)	248
	F.7	The flux-cell measure: $2\pi/K$ and the single-transit gearing $1/K$	250
	F.8	The archived four-peak template and its matched filter	252
	F.9	The archived two-port ratio estimator and dual-mode remark	253
	F.10	The archived inversion step of the baryonic-ledger course	254
G		Notation	256
Part XIV. Common Objections and Responses			258
H		Methodology and the derivation chain	258
I		The geometry and its protection	261
J		Cosmology and the relic pipeline	264
K		Couplings, locks, and the strong-CP assignment	267
L		Experimental practice	269

M	Instruments and astrophysics	273
N	Foundations and philosophy	275

Part XV. References	280
----------------------------	------------

Preface

This Preface is addressed to a graduate student or postdoc who is in — or is deciding whether to join — a haloscope collaboration, and who is wondering what the next five years of cold nights at the dilution refrigerator are for.

Consider, first, where the field stands. Dark matter was established as a fact of the sky across ninety years of astronomy, and for the last fifty of them physics has searched for its particle by sweeping *spaces*: the WIMP program swept decades of mass and ten orders of magnitude of cross-section, guided by naturalness arguments that migrated as they were excluded; the axion program scans decades of frequency because the QCD axion’s mass is, by construction, unknown. These campaigns produced fine instruments and airtight null results — and not one of them could be pointed at a *single frequency*, because no theory staked one. Every anomaly the field has debated — an annual modulation, a 3.5 keV line, most recently a 5.1σ candidate at 1.036 GHz — died ambiguous for the same reason: one signature, no pre-registered number, nothing to check it against.

This paper offers a target with its numbers already on paper. A frequency — 1.95 GHz, the marked rung of a closed window: a band to sweep, with a bin to camp on first. A strength — exactly on the DFSZ line you already calibrate against. A shape, an annual breathing, a sidereal wobble. And a count no other candidate offers: the line stands alone inside its fenced band, because the geometry supplies exactly one light field of this kind. Parts VI–VIII derive this; Part IX simulates the discovery week it implies, with statistical and systematic errors on every panel; Part X designs the observatory — KAXIO — that covers the whole band and runs the count test live. The simulation script is available from the author on request; the input tables are yours to edit.

What should you do with it on Monday morning? If your collaboration’s scan crosses — or has ever crossed — the band between 1.59 and 2.13 GHz, re-bin the archived spectra for a second line anywhere in the band: it costs no hardware, no beam time, and no permission, and it runs this theory’s count test for free (§37.3). If you operate a cavity that can reach the window, the camping arithmetic of §37 says a decisive exposure of the benchmark costs *minutes* of live time, not years. And if you are choosing where to invest a career: Part X’s run plan has a one-year outcome on the existence question — discovery or window closure — that cannot be inconclusive.

Two warnings, so the invitation is honest. The theory whose numbers you would be testing is not the community’s consensus framework; it is a single deductive structure, published in a twenty-two-paper cohort, and this paper’s claims stand or fall with theorems proved elsewhere in it. The objections Part and the honesty ledger (Table 13) say which links are theorems, which are derived, and which rest on the named conjectures and the one calibration. And the falsifiers cut both ways by design: a swept and silent window kills the calibrated placement and its ladder; a line outside the window does the same, or, at the randomized-angle mass ($\gtrsim 6$ GHz), kills the inflation-era pinning; a second line inside the window kills the count and names which registered mode is light. Which object dies is written down beside each, and so are the dates.

That is the wager, and it is why the experimentalist, not the theorist, is this paper’s real addressee. The dark-matter problem has waited ninety years not for a better detector but for a prediction worth pointing one at. The pages that follow are that prediction, with your name — whoever you are, at whichever magnet — already written into its verdict.

Part I

Introduction

The dark-matter problem, the twenty-two-paper cohort, and the roadmap of the fifteen Parts

1 The dark-matter problem

The case for dark matter has been built over ninety years, in four acts (Bertone & Hooper, 2018; Trimble, 1987) — a cluster anomaly, the galactic rotation curves, the precision cosmic inventory, and a long and so-far-null direct search. We recount them in turn, because the candidate of this paper will be tested against every fact the four acts pin down, and then state the modern challenge that organizes the rest of the work: not detection, but *identification*.

1.1 A brief history of the dark-matter problem

In 1933 Fritz Zwicky applied the virial theorem — the bookkeeping rule that converts the speeds of a cluster’s galaxies into the mass needed to hold them together — to the Coma cluster and found that the velocity dispersion of its galaxies implied a gravitating mass exceeding the luminous mass by more than an order of magnitude (Zwicky, 1933). He named the deficit *dunkle Materie*. The observation was largely set aside for four decades, until Vera Rubin and Kent Ford, measuring optical rotation curves of Andromeda and then of dozens of spirals, established that orbital velocities do not fall off in the Keplerian $v \propto r^{-1/2}$ manner outside the luminous disk but stay *flat* to the last measured point (Rubin & Ford, 1970). A flat rotation curve means $M(< r) \propto r$: every spiral galaxy is embedded in a halo of unseen mass extending far beyond its stars.

The third act made the deficit precise. The acoustic peaks of the cosmic microwave background — the frozen sound waves of the infant universe — respond differently to matter that couples to photons (baryons) and matter that does not; the Planck mission measured the two components separately, $\Omega_b h^2 = 0.0224$ and $\Omega_c h^2 = 0.1200 \pm 0.0012$ (Planck Collaboration, 2020) — the invisible component outweighing the visible five to one, with one-percent precision, at redshift eleven hundred. Whatever dark matter is, it was already in place before a single star existed.

The fourth act has been a search, and so far a null one. The theoretically favored candidate of the 1980s–2010s, the weakly interacting massive particle (WIMP) (Bertone, Hooper, & Silk, 2005; Lee & Weinberg, 1977), motivated a generation of magnificent experiments — tonne-scale xenon time-projection chambers, cryogenic crystals, indirect searches, collider hunts — which have driven the spin-independent cross-section bound below $\sim 10^{-47} \text{ cm}^2$ at the optimal mass without a confirmed signal. The null results do not falsify dark matter (the gravitational evidence is overwhelming and independent); they falsify an ever-growing portion of one candidate’s parameter space, and they have returned the field to its widest question in fifty years: *not where is it, but what*

is it. Through that opening, the axion family — light pseudoscalars with topological couplings, proposed originally to solve the strong-CP problem (Peccei & Quinn, 1977; Weinberg, 1978; Wilczek, 1978) — has moved from alternative to front-runner.

1.2 The astronomical evidence

It is worth itemizing what the evidence actually pins down, because the candidate of this paper will be tested against every item.

Dynamics. Flat rotation curves in spirals; virial masses of groups and clusters; the velocity dispersions of dwarf spheroidals, the most dark-matter-dominated objects known (mass-to-light ratios in the hundreds).

Lensing. Strong and weak gravitational lensing maps the total mass directly, luminous or not. The Bullet Cluster is the canonical exhibit (Clowe et al., 2006): in this pair of colliding clusters the hot X-ray gas (most of the *baryonic* mass) lags at the collision point while the lensing mass tracks the galaxies that passed through — gravity points at something collisionless that is not the gas. Modified-gravity explanations must work hard here; a collisionless particle does not.

The CMB and BAO. The relative heights of the acoustic peaks fix Ω_b and Ω_c separately, as above; the baryon acoustic oscillation scale in galaxy surveys — the same frozen sound waves, measured in today’s maps of galaxies — confirms the same inventory at low redshift.

Structure formation. Baryons cannot begin collapsing until recombination frees them from photon pressure; the observed amplitude of today’s structure requires a component that began collapsing *earlier* — electrically neutral, hence decoupled from the photon fluid. Moreover the observed small-scale power requires that component to be *cold*: relativistic (“hot”) dark matter free-streams away exactly the dwarf-galaxy-scale structure we observe in abundance.

Collecting the items, any successful candidate must be: (i) electrically neutral and colorless; (ii) cold (non-relativistic from early times, with negligible free-streaming); (iii) effectively collisionless at cluster precision; (iv) stable over $t_U = 13.8$ Gyr; and (v) produced with $\Omega_c h^2 \simeq 0.120$. These five requirements — and nothing about them is negotiable — are precisely the five properties that Parts II–V derive, rather than assume, for the Kaxion: neutrality from Corollary 8.3, coldness from Section 16, collisionlessness from the $1/f_K$ portal suppression, stability with thirty-two orders of headroom for the physical line from Section 28, and the abundance from the π -locked closure of Section 25.

1.3 The detection challenge

The modern problem is sharper than detection; it is *identification*. Suppose tomorrow a haloscope records a persistent narrow excess. What, beyond its frequency, identifies it? A generic axion-like particle has two free parameters — mass and coupling — and any point in the open plane can be declared a model after the fact. A detection without a prior, parameter-free prediction confirms “something” but identifies nothing; this is the epistemic weakness of candidates with adjustable dials, and it has haunted the field’s occasional anomalies (the 3.5 keV line, the DAMA modulation, low-threshold excesses) — each consistent with *some* model, none demanded by one.

The inversion this paper offers is the point of the entire exercise. The Kaxion arrives with one calibrated scale, declared as such, and *no* adjustable dial after it: its mass sits on a discrete window whose benchmark row is calibrated through the π -locked abundance (Part V), its photon coupling sits exactly on the DFSZ line with the bracket $0.75(4)$ fixed by group theory and chiral dynamics, and — decisively — it carries a fingerprint that a generic axion cannot fake: a discrete window, a fixed address on the DFSZ line, and a line that must stand alone, because the geometry supplies exactly one light field of this kind. A signal matching the sheet identifies the Kaxion, its cap loading, and with them the series. A null across the window retires the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d; the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinned branch or the relic is not this sector's. A second line inside the window refutes the count and names a registered mode. Neither is a kill of the framework, and this paper does not claim one. A line outside the window retires the same objects as the null, or, at $\gtrsim 6$ GHz, the pinning conjecture (Part VI). The remainder of this paper builds that sheet (Parts II–V), sharpens the fingerprints (Part VI), and maps the laboratory program that can settle the question (Parts VII–VIII).

2 The central claim: the Kaxion dark-matter mass as an open-system observable

Let us set down, 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, 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. In plain words: what we call “the world” is what a finite observer registers of a deeper six-dimensional structure; the rules of that registering are fixed, not chosen; and anything the registering must contain — like the field this paper is about — arrives with its properties already decided. It follows that the Kaxion dark-matter mass, $m_{\text{Kaxion}} \approx 8.1 \mu\text{eV}$, is a *registered observable* of one projection — its ladder read from the geometry's integers by an ansatz, its count fixed conditionally, its absolute placement calibrated once through the relic closure — rather than chosen — which is why this paper can derive it, and why a single clean measurement can refute it. The funnel below is shared, in structure, with the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs.

3 Three reading paths

Before the derivation begins, it is worth saying plainly that this paper serves three audiences whose needs differ, and that it has been arranged so each can take its own road.

The haloscope experimentalist should go straight to the laboratory spine: Part VI for the count and the window, Part VII for the landscape, the camping arithmetic, and the discovery calendar, Part VIII for the dedicated instruments, Part IX for the simulated discovery week with its error budgets, and Part X for the KAXIO proposal. Part V can be consulted afterwards, as the provenance




Figure 3: 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 Kaxion dark-matter mass.

file for every number the spine spends; the objections most likely on this reader’s mind are answered in the “Experimental practice” and “Instruments and astrophysics” sections of Part XIV.

The cosmologist should begin with Part III — the misalignment history, the isocurvature accounting, and the π -locked abundance target — then take Part V for the closure of the decay constant and the mass, and Part IV for the candidate landscape, the JWST early-structure argument, and the cusp–core prediction. The relic pipeline’s honest fences (the misalignment-completeness and mass-uniqueness conjectures) are flagged where they enter and collected in the master ledger (Table 13); the objections written for this reader are Part XIV’s “Cosmology and the relic pipeline” section.

The newcomer — a graduate student meeting GG-Theory for the first time — has a self-contained course: Part II builds the framework from its constitution with no prior exposure assumed, the tutorial appendices of Part XIII — the four primers (Appendices B.1–B.4) and the from-scratch derivations that follow them — supply every technical ingredient, and Parts III–V can then be read in order as one narrative from the early universe to the laboratory line. The abbreviations dictionary and the symbol glossary (Appendices A–B) are meant to be kept open alongside.

Whichever road is taken, all three converge on the same destination: the four falsifiable numbers of Part V, the falsifiers of Part VI, and the dated verdict of Parts VII–X.

4 The cohort inheritance and the roadmap of the fifteen Parts

Before the derivation begins (Part II), it is useful to lay out two reader-orientation maps on the same opening: a panorama of the twenty-two-paper cohort in which this paper stands (§4.1, Figure 4), and a panorama of the fifteen Parts of the present paper (§4.2, Figure 5). 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 fifteen-Part arc.

4.1 The cohort inheritance

The present paper is not free-standing. It is the dark-matter entry of the *Cosmology & Astrophysics Branch* of the GG-Theory program, and it inherits its entire apparatus — the geometry, the anomaly structure, the runtime, and the cosmology — from eleven parent papers, the *GG-A series* GG-A1-Constitution–GG-A11-Origin. A reader meeting the Kaxion for the first time need not have read them: Part II reconstructs, from scratch, the few ingredients the Kaxion actually uses. But it helps to see once, in a single picture, where each ingredient comes from. Figure 4 is that family tree — drawn for the full twenty-two-paper cohort — with the present paper, GG-C1-Kaxion, standing in the Cosmology & Astrophysics frame.

The constitutional root is GG-A1-Constitution (Arisaka, 2026, A1-Constitution), which fixes the four axioms and seven postulates — among them the rule of *Maximum Symmetry, Minimum Freedom* that forbids inventing a dark-matter sector while a forced one stands, and the assignment (Problem P7) of the strong-CP resolution to the Kaxion.

The *GG-6 Trio* builds the bulk. GG-A2-GG6 (Arisaka, 2026, A2-GG6) proves the six-dimensional geometry with its two non-spatial coordinates — a *scale* and a *phase* — and the integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. Those five integers are not chosen but *forced* by the anomaly-free BI-6 closure (GG-A3-BI6), each carrying a definite meaning — $M = 1$ the 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 that also fences the laboratory window. Exactly two transcendental constants then appear in every formula, tied one to each coordinate: e to the scale, whose exponential throat depth e^{-K} carries the hierarchy, and π to the phase, the angle of the compact circle that fixes the π -locked inventory — no other constant is admitted. GG-A3-BI6 (Arisaka, 2026, A3-BI6) proves the anomaly identity BI-6 that forces the Green–Schwarz two-form B_2 whose light remnant is the Kaxion, and fixes the baryonic trace ratio $2/5$. GG-A4-Euclid (Arisaka, 2026, A4-Euclid) supplies the historical lineage and the reconstruction of the universal constants from the geometry.

The boundary law is GG-A5-GDOS (Arisaka, 2026, A5-GDOS): the projection of the bulk to the four-dimensional boundary (GDOS), with the conserved-current package through which the Kaxion is permitted to couple. The *runtime pillars* are GG-A6-GRIP (Arisaka, 2026, A6-GRIP), which defines the GRIP cascade that births, selects, and stabilizes the Kaxion as the Geometric Attractor (the single-patch history in which one misalignment angle is meaningful is carried as the named conditional C-Kaxion-d, Part III), and GG-A7-Qi4 (Arisaka, 2026, A7-Qi4), the quantum-information quartet whose current-locking Arisaka equation transports the ledger’s portal weights to the boundary.

The *Consequence Trio* are the textbook readouts. GG-A8-QM (Arisaka, 2026, A8-QM) promotes quantum mechanics to a boundary channel; GG-A9-SM (Arisaka, 2026, A9-SM) derives the Standard Model, supplying the charge content and the scale chain that fixes the lock scale (the anomaly ratio $E/N = 8/3$ itself descends from the GG-A2-GG6/GG-A3-BI6 embedding and is evaluated as pure group theory in Part V, per the provenance table of Part IV); and GG-A10-Cosmos (Arisaka,

2026, A10-Cosmos) derives the π -locked Λ CDM inventory $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ that the Kaxion must saturate, together with the conjectures this paper develops and fences. The closing capstone GG-A11-Origin (Arisaka, 2026, A11-Origin) defines the Ground Geometric Attractor that selects the Kaxion as the unique deepest vacuum of its stratum.

Beneath the cohort sit the two phenomenology branches. The Predictions Branch carries GG-P1-KK (Arisaka, 2026, P1-KK) (the Kaluza–Klein lock cluster at $M_{\text{lock}} \approx 3.85$ TeV, searchable at the LHC), GG-P2-N3 (Arisaka, 2026, P2-N3) (the three-generation count forced by the ledger), and GG-P3-Electron (Arisaka, 2026, P3-Electron) (the electron as the unique ground state of the boundary). The Cosmology & Astrophysics Branch carries GG-C2-Hubble (Arisaka, 2026, C2-Hubble) (the two Hubble constants as the two ends of one geometric corridor), GG-C3-BH (Arisaka, 2026, C3-BH) (the exact Bekenstein–Hawking coefficient), and the present paper, GG-C1-Kaxion (the Kaxion: a single line in the window fenced by $K = 35$). Each is a sharp, independent falsifier of an upstream cohort theorem.

On the map’s left flank stands the *Foundations Branch* — the two member papers whose subject is the boundary’s complex structure and its mirrors — and the present paper’s stake in it is concrete. GG-P5-CPT (Arisaka, 2026, P5-CPT) shares this paper’s strong-CP resolution: the two-fate architecture in which the weak CP phase is locked by geometry while the one relaxed angle $\bar{\theta}$ is handed to the compact-angle pseudoscalar — its Part VII names the Kaxion as the companion that relaxes what remains — and it carries the three mirrors of matter, C, P, and T as orientation reversals of the three real bulk factors, whose breaking the kaxion sector survives. GG-P4-Real (Arisaka, 2026, P4-Real) is the origin-of- i member: 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 the boundary’s quantum description descends. Both carry P-series designators; their place at the spine’s flank, not the phenomenology rows, marks their foundational role.

This subsection is the panorama only; the parent-by-parent detail of what the Kaxion imports is the business of Part II and of the provenance table of Part IV. The twenty-two-paper map, identical in every paper of the cohort, is Figure 4, which closes this subsection.




Figure 4: 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 GG-6 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.

4.2 Roadmap of the fifteen Parts

The rest of the paper is organized into fourteen further Parts, plotted in Figure 5. The present Part is nearly done — the ninety-year case for dark matter, the identification problem that organizes the search, the central claim, the three reading paths, and the cohort the paper stands in — and the itinerary below says where each piece lives.

Part II answers the question of the title — *What is the Kaxion?* — as a self-contained tour assuming no prior exposure to GG-Theory: the six-dimensional bulk, the boundary law, the five-link existence chain, the tensor remnant and the lifted holonomy, and the GRIP cascade that births, selects, and stabilizes the field, closing at the existence theorem T-Kaxion-1. Part III, the *Cosmology of the Kaxion*, then runs the object through cosmic history as one event chain — inflation exit, Kaxion birth, dark energy: misalignment production, the decay constant and the angle, the π -locked abundance, the coldness, and the whisper-quiet interactions with matter. Part IV, the *Context*, places the candidate among its rivals: the empirical status of the field’s candidate landscape, the framework’s own verdict that the Kaxion is required and sufficient, and the astrophysical reach from the JWST early-structure era to the cusp-core problem.

Part V is the quantitative heart, highlighted in the roadmap because it carries the paper’s central claim. From the series theorems it derives the basic quantities one by one: the strong-CP alignment that legitimizes the QCD susceptibility, the decay constant closed from the π -lock, the mass $m_K \approx 8.1 \mu\text{eV}$ with its laboratory line at 1.95 GHz, the photon coupling on the DFSZ line factor by factor, the abundance recheck and the lifetime, and the error budget with the theorem-versus-conjecture ledger. Every falsifiable number the later Parts spend is registered in this Part.

Parts VI–X are the laboratory arc. Part VI sharpens the *Unique fingerprints* no generic axion can imitate — the fenced window, the DFSZ address, and the count with its cheap reanalysis test. Part VII, *Detection*, maps the haloscope principle and the current experiments onto the pre-registered single-frequency campaign: camping on the frequency rather than scanning past it. Part VIII, *Future detection*, designs the dedicated instrument for the count test — a multi-mode receiver that interrogates the fenced window. Part IX, *Simulating the Discovery*, computes the data of discovery week in advance — the camped spectra, the modulation fingerprints, and the count-test acceptance. Part X assembles everything into *The KAXIO Proposal*: a second-generation Kaxion observatory at Fermilab — two sub-arrays, one fenced window, one count, one year.

The paper then steps back. Part XI, the *Discussion*, weighs the construction: summary, significance, the classic puzzles resolved and the innovations registered, and the road to a verdict. Part XII, the *Conclusions*, says plainly why the Kaxion matters — and why finding it matters to us. Three Parts of back matter close the volume: the *Technical Appendices* of Part XIII (the graduate machine room in seven blocks, from the reference dictionaries to the promotion roadmaps), the *Common Objections* of Part XIV (the strongest first-reading objections, answered with pointers into the paper), and the *References* of Part XV (cited works in two groups, each entry carrying a relevance note).




Figure 5: Roadmap of the fifteen Parts of GG-C1-Kaxion, grouped into eight phases — *Setup* (Part I), *Theory* (Parts II–III), *Context* (Part IV), *Derivation* (Part V), *Fingerprints* (Part VI), *Laboratory* (Parts VII–X), *Verdict* (Parts XI–XII), and the *back matter* (Parts XIII–XV). **Part V** is highlighted because it carries the paper’s central claim: the Kaxion mass $m_K \approx 8.1 \mu\text{eV}$ at its calibrated benchmark — the laboratory line at 1.95 GHz — together with the photon coupling, the abundance, and the lifetime, from which every falsifiable number of the later Parts descends. The fingerprints of Part VI and the laboratory arc of Parts VII–X spend the numbers Part V registers.

Figure 5 renders this Part-by-Part development as a single flowchart, so the big picture is visible at a glance.

But every road begins at the same place — with the object itself. Before the numbers, the falsifiers, or the calendar can carry any weight, one has to see *why* a rigid six-dimensional geometry is compelled to contain a light, neutral, cold pseudoscalar at all. That is the work of Part II, which builds the Kaxion from the constitution upward — the bulk and its integer ledger, the boundary law, the runtime that chooses a ground state — and ends where every later Part begins: at the existence theorem for a single light compact-angle field.

5 Part I in plain English: what the dark-matter problem is, and what would count as solving it

Eighty-five percent of the matter in the universe has never been seen. It bends light, it holds galaxies together, and it left its fingerprint on the microwave background, and in ninety years nobody has caught a single particle of it. That is the problem this paper is about.

The difficulty is not a shortage of candidates. There are many, and most of them work. A WIMP with the right mass and cross-section fits; an axion with the right decay constant fits; a sterile neutrino with the right mixing fits. The trouble is the phrase "with the right," because in each case the right value is supplied by hand, and a candidate whose parameters are free predicts a region rather than a point. A region can be searched for decades without being excluded.

What this paper offers is narrower and, if it is wrong, easier to kill. The claim is that a specific six-dimensional geometry, adopted in the parent papers for reasons that had nothing to do with dark matter, forces a light pseudoscalar into existence and leaves nothing about it adjustable after one declared calibration — not its mass, not its coupling, not its count. If that is right, the candidate sits at one frequency, and an experiment that scans through that frequency at the required depth either finds it or ends it. The rest of the paper is the derivation of the frequency and the account of what it would take to look.

Part II

What is the Kaxion?

A self-contained tour of GG-6, GDOS, and GRIP — and the pseudoscalar they force

6 The GG-Theory framework, reviewed for newcomers

The claim of this Part is unusual for a dark-matter paper, so we state it before any formalism. We will not argue that the Kaxion is a plausible candidate, an attractive candidate, or a well-motivated candidate. We will argue that it is a *forced* one: that a reader who accepts the published theorems of the GG-A series — theorems proved for reasons having nothing to do with dark matter — has already accepted the existence of a light boundary pseudoscalar, whether or not the dark-matter problem is ever mentioned.

That argument is only as persuasive as the reader’s grasp of the theorems it leans on, and we do not assume prior exposure to the series. This section is therefore a self-contained tour of the three structures the Kaxion descends from: GG-6, the six-dimensional stage on which everything stands; GDOS, the projection that turns bulk physics into what a four-dimensional observer actually registers; and GRIP, the dynamical runtime that decides which of the allowed configurations the universe occupies. A reader fluent in the series may skim to Section 7; a graduate student meeting GG-Theory for the first time should find everything needed here, with pointers to the parent papers for proofs.

6.1 The constitutional method

The GG-Theory program does not begin with a Lagrangian; it begins with a constitution (Arisaka, 2026, A1-Constitution) — four axioms and seven postulates from which the Lagrangian is *derived*. Three constitutional clauses do almost all the work in this paper, and the reader needs only these.

The Geometry–Gauge Equivalence Principle (GGEP) demands that gravity and the gauge interactions be two facets of a single composite connection on one manifold — no “gravity plus matter sprinkled on top.” It is what forces extra dimensions: four dimensions cannot carry a connection rich enough to house the Standard Model and curvature in one object.

The Operational Geometry–Gauge Equivalence Principle (OGGEP) demands that physics be expressed as what a *finite boundary observer* registers. We do not live in the bulk; every measurement ever performed is a boundary event. OGGEP is the clause that makes the bulk-to-boundary projection (GDOS, Section 6.4) constitutional rather than optional.

Maximum Symmetry, Minimum Freedom (MSMF) is the anti-tuning rule: no continuous free parameter may appear anywhere in the framework unless absolutely required; discrete, topological data are admissible, continuous dials are not. MSMF is the clause a newcomer should watch most

carefully in this paper, because it acts three times on the Kaxion: it forbids a tunable circle radius (Section 6.2), it forbids a scanned misalignment angle (Part III), and it forbids inventing a new dark sector while a forced one stands available (Remark 8.5).

Finally, the falsification postulate (MF) requires every quantitative claim to carry a pre-declared kill criterion. This paper honors it in Parts VI–XI and in the objection ledger of Part XIV: the Kaxion can be killed by specific measurements, stated in advance.

For reference — since this paper will invoke clauses by name — the complete constitution is eleven lines, collected in Table 1 together with the duty each clause performs here. Read its right-hand column as a promissory index: AX-2 forces the extra dimensions, AX-3 mandates the projection, AX-4 is the anti-tuning rule that acts three times on the Kaxion, and every postulate row is discharged at the section it names:

Table 1: The constitution of GG-Theory (GG-A1-Constitution): four axioms and seven postulates, with the role each plays in this paper. The point of reprinting the constitution here rather than pointing at GG-A1-Constitution is the third column. Each clause is stated with the specific job it does in the Kaxion argument, so a reader who doubts a later step can find the axiom it rests on without leaving the paper. Nothing in the table is derived; it is the list of what the derivation is allowed to use.

	Clause	One-line content; role in C1-Kaxion
AX-1	LC (Locality and Causality)	boundary physics closes under the Lorentz group; guarantees the Kaxion is an honest 4D pseudoscalar.
AX-2	GGEP (Geometry–Gauge Equivalence)	geometry and gauge are one composite connection; forces the extra dimensions of Link 1.
AX-3	OGGEP (Operational GGEP)	physics = what a finite boundary observer registers; mandates the GDOS projection (§6.4).
AX-4	MSMF (Max. Symmetry, Min. Freedom)	no continuous free parameters, discrete data only; acts three times on the Kaxion (§6.1).
POS-1	GI (Gauge Invariance)	no accidental global symmetries; forces gauged $U(1)_B$ — the ledger’s baryonic entry.
POS-2	UEP (Universal Equivalence)	all sectors couple through the one connection; no dark sector may be “decoupled by hand.”
POS-3	AF (Anomaly Freedom)	the quantum theory must be exactly anomaly-free; forces B_2 — the Kaxion’s mother field (Link 3).
POS-4	CP (Compositional)	boundary dynamics is CPTP; the open-system structure behind Eqs. (6.19)–(6.21).
POS-5	RG (Renormalization Group)	scale flow is geometrized as the r coordinate; the scale ladder of §6.2.
POS-6	OGN (Occam Gauge Normalization)	structural integers form the OGN-5 ledger; the gear ratios of every fingerprint (the absolute scale is calibrated).
POS-7	MF (Measurement & Falsifiability)	every claim carries a pre-declared kill criterion; discharged by F-Kaxion-1 (Part VI).

6.2 GG-6: the six-dimensional stage

6.2.1 The manifold and the metric

The bulk paper GG-A2-GG6 (Arisaka, 2026, A2-GG6) proves (T1, T2) that the constitution forces a six-dimensional bulk with the product structure and canonical metric

$$\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S_\theta^1, \quad (6.1)$$

$$ds_6^2 = e^{-2r} \eta_{\mu\nu} dx^\mu dx^\nu + dr^2 + d\theta^2, \quad r \in [0, K], \quad \theta \in [0, 2\pi), \quad (6.2)$$

with $G_{\theta\theta} = 1$ and *no free radius*: no adjustable R_θ , no adjustable scale profile, no compactification modulus. In textbook coordinates the six directions are $(t, x, y, z; r, \theta) = \{\text{time, three space, scale, phase}\}$, and the names are exact, not poetic.

Why exactly six. The dimension is not an aesthetic choice; GG-A2-GG6 T1 forces it by counting structural roles. The constitution requires three irreducible geometric functions: a Lorentzian arena for registered events (four dimensions, by LC and the observed signature); a geometrized renormalization-group flow (one more dimension — POS5 demands that scale be a coordinate, not a parameter); and a carrier of integer-quantized Abelian holonomy (one compact dimension — the minimal smooth geometry whose Wilson loop (6.4) is integer-valued is a single circle). Four plus one plus one, with T1 proving that no dimension can serve two roles and no role can be dropped: $D = 6$, exactly. Contrast the logic with string theory’s $D = 10$ (fixed by worldsheet anomaly cancellation) or Randall–Sundrum’s $D = 5$ (chosen phenomenologically): here each dimension holds a named job, and the Kaxion will turn out to live precisely at the intersection of the two non-spatial jobs — a topological angle (the phase circle) remembered across a scale flow (the throat).

The scale coordinate. The interval coordinate is the renormalization group made literal: $r = \ln(E_{\text{GUT}}/\mu)$, the logarithm of an energy ratio. Moving inward along r is lowering the observation scale μ ; an RG trajectory, which the textbooks draw as a curve on a coupling-constant plot, is in GG-6 a geodesic in a real dimension. The scale profile e^{-2r} in (6.2) then says something quantitative: every unit step in r red-shifts all boundary energies by one factor of e . The interval has length $K = 35$ (the integer derived below), so the total red-shift from the ultraviolet end to the infrared end is $e^{-35} = 6.3 \times 10^{-16}$. This single number is the resolution of the hierarchy problem in GG-6: the famous sixteen orders of magnitude between gravity and the weak scale are not a mystery to be explained by new particles, they are the exponential of an integer,

$$M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K} = \frac{1}{2} \times (1.22 \times 10^{19} \text{ GeV}) \times 6.3 \times 10^{-16} \simeq 3.85 \text{ TeV}, \quad (6.3)$$

the mass of the Kaluza–Klein lock cluster predicted for colliders (Arisaka, 2026, A9-SM; Arisaka, 2026, P1-KK). The unification scale sits at the other natural exponent, $E_{\text{GUT}} = E_{\text{Pl}} e^{-2\pi} \simeq 2.28 \times 10^{16} \text{ GeV}$. *Why is gravity 10^{16} times weaker than the weak interaction? Because thirty-five.*

The phase coordinate. The circle coordinate θ is a pure angle — dimensionless, periodic, carrying no unit of length. GG-A2-GG6 is explicit that it is *non-spatial*: a value of θ is an internal phase, not a location; local motion in θ is a gauge transformation, a redundancy rather than a displacement; no signal propagates along it. The only gauge-invariant content of the circle is its holonomy, the Wilson loop

$$W = \text{tr } P \exp \left(i \oint_{S_\theta^1} A_\theta d\theta \right), \quad (6.4)$$

whose eigenvalues are quantized to integers by single-valuedness of the fields. A reader who knows the Aharonov–Bohm effect already knows this object: the electron circling a solenoid acquires a phase that is physical even though the local field along its path vanishes. The phase circle of GG-6 is an Aharonov–Bohm solenoid built into the fabric of the theory, and Section 9 will show that the Kaxion is the slow wandering of the tensor’s dual angle, while the circle’s own holonomies — six of them, one per Cartan direction of the rank-six group — are lifted.

Physical picture. The phase circle is a clock dial, not a road. Nothing travels along it; what is physical is where the hands point, and after a full circuit the hands must return to a quantized position. A periodic, dimensionless, integer-quantized internal angle is precisely the natural habitat of a periodic field — and this paper is, at bottom, the study of what happens when this dial is displaced and released.

6.2.2 The gauge group, and why baryon number is gauged

GG-A2-GG6 (T3) forces the bulk gauge–gravity group

$$G_{\text{Bulk}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}. \quad (6.5)$$

Each factor is forced, not chosen. $\text{Spin}(1, 5)$ is the Lorentz group of the six-dimensional bulk. $\text{SU}(5)$ is the *smallest* simple group containing the Standard Model — the Georgi–Glashow spine — selected by MSMF from among all candidates. $\text{U}(1)_B$ is gauged baryon number: because the constitution forbids accidental global symmetries, the observed proton stability ($\tau_p > 10^{34}$ yr) must be protected by a *local* symmetry, and gauging B forbids every proton-decay operator outright — in sharp contrast to classic $\text{SU}(5)$, whose X, Y bosons mediate $p \rightarrow e^+ \pi^0$. $\text{U}(1)_{Q_i}$ is the informational current factor required by the four-current identity of GG-A7-Qi4 (Section 6.4). For this paper, the decisive feature of (6.5) is the presence of *two* Abelian factors: they are the ledger entries whose $-2/5$ coefficients lock the Kaxion’s portals to the anomaly bookkeeping.

6.2.3 The OGN-5 integer cascade

With the stage and the group fixed, GG-A2-GG6 (T6) proves that every remaining dimensionless structural quantity collapses to five integers,

$$(M, N, L; J, K) = (1, 3, 5; 7, 35), \quad (6.6)$$

derived in a five-step Diophantine cascade with no input anywhere:

$$\begin{aligned}
M &= 1 && \text{one compact circle (T1: dimensional minimality);} \\
L &= 5 && \text{the minimal simple-unification index (SU(5));} \\
N &= 3 && \text{from the BI-6 trace ratio } (L - N)/L = 2/5 \text{ (Section 6.3);} \\
J &= 2N + 1 = 7 && \text{the minimal odd closure index;} \\
K &= LJ = 35 && \text{the scale depth, closed by the dual-lock identity below.}
\end{aligned} \tag{6.7}$$

The cascade terminates in an overdetermination that is the series' favorite consistency check:

$$K = LJ = M^2 + N^2 + L^2, \quad 35 = 5 \times 7 = 1 + 9 + 25. \tag{6.8}$$

Two independent expressions — a product and a sum of squares — must agree, and do. A newcomer should pause on (6.8): in a framework with even one continuous dial, such a coincidence could be arranged; in a framework with none, it either happens or the framework dies.

6.2.4 The track record

Why should a newcomer take five integers seriously as physics? Because the same ledger has already been spent, repeatedly, on quantities that were measured long before the series existed. Table 2 collects nine examples from the parent papers; none has a tunable parameter behind it.

Table 2: What the same five integers have already bought (parent papers GG-A9-SM, GG-A10-Cosmos). Every entry is parameter-free; deviations are against PDG/CODATA/Planck values. The table is evidence about the method rather than about the Kaxion. None of these numbers was computed for dark-matter reasons and none is adjusted here; they are what the same five integers produced in the parent papers, set against the standard compilations. A reader who thinks the Kaxion result is a fit should account for this column first.

Quantity	OGN-5 formula	Predicted	Measured	Dev.
$\hat{\alpha}^{-1}(M_Z)$	2^J	128	127.951 ± 0.009	+0.04%
$\alpha^{-1}(0)$	$2^J + N^2$	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%
Recombination z_* (atomic chain)	GG-A10-Cosmos	1091.1	1089.9 ± 0.3	+0.1%
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%

Two kinds of entry, and an honest look-elsewhere. The nine rows are not equally strong evidence, and a careful reader should grade them. The *rigid* entries — $\hat{\alpha}^{-1} = 2^J$, $\sin^2 \theta_W = L^2/4N^3$, $v = M_{\text{lock}}(2/L)^N$, $z_* = e^J$ — are integer or exponential relations in the ledger integers whose functional form admits essentially no freedom: once the integers are fixed (by the gauge structure,

independently), the value is forced, and matching a quantity measured to 10^{-3} – 10^{-4} with such a form carries a small look-elsewhere probability. The π/e *fractions* — $\alpha_s = 1/(\pi e)$ and the three cosmological shares — are weaker, and we say so plainly: π and e are flexible symbols, and the number of comparably simple expressions in $\{\pi, e, \text{small integers}\}$ that land within the percent-level error bar of an $\mathcal{O}(1)$ quantity is not small (for α_s , of order several), so $1/(\pi e)$ is suggestive rather than decisive. What guards the table against numerology is not the individual coincidences but their *provenance*: each formula was fixed in the parent papers (GG-A9-SM, GG-A10-Cosmos) by the geometry, independently of the comparison made here — they are pre-, not post-dictions — and all nine use the *same* five integers, not a fresh choice per row. The Kaxion sheet of this paper is held to that standard: its integers are the table’s integers, declared in advance, and its decisive fingerprints — the fenced window, a ladder the integers set by ansatz around one calibrated scale, and the count, conditional on the parents’ registered modes — are integer-shaped statements rather than π/e fractions, each graded where it is stated (Part VI).

A note on the cosmological shares. Three rows (Ω_m , Ω_{DM} , Ω_Λ) are present-epoch density *fractions*, and a fraction of the critical density is time-dependent — $\Omega_m(z)$ runs from near unity in the past toward zero in the future — so “ $\Omega_m = 1/\pi$ ” is a statement about *now*, not a timeless constant. The epoch-independent content, and what the abundance closure of Part III actually uses, is twofold: the comoving combination $\Omega_{\text{DM}} h^2 = 0.1215$, and the constant ratio of the two matter species,

$$\frac{\rho_{\text{DM}}}{\rho_b} = \frac{1/\pi - 1/(2\pi^2)}{1/(2\pi^2)} = 2\pi - 1 \simeq 5.28 \quad (\text{measured } 0.1200/0.02237 = 5.36, -1.5\%), \quad (6.9)$$

which redshifts identically for both species and therefore holds at every epoch. We read the π -locked fractions as the present-epoch evaluation of these timeless quantities, not as a claim that a fundamental constant equals today’s matter fraction.

This table is the newcomer’s reason to read on. The Kaxion prediction of this paper — mass, coupling, abundance, window, count — is one more withdrawal from the same account: the integers $K = 35$, $J = 7$, $L = 5$ will reappear below in new roles, and the reader should recognize each appearance as the *same* five integers rather than as fresh assumptions.

6.2.5 The energy ladder, numerically

Because the scale arithmetic carries the entire hierarchy, we evaluate it once, explicitly, so that every later scale can be read off one ladder. The two anchor theorems are $E_{\text{GUT}} = E_{\text{Pl}} e^{-2\pi}$ and the lock theorem of GG-A9-SM, $M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K}$; read together they place the ultraviolet end of the K -deep throat at $\frac{1}{2} E_{\text{Pl}}$, with E_{GUT} reached $2\pi - \ln 2 \simeq 5.6$ e-folds inside it (the labeling $\mu(r) = E_{\text{GUT}} e^{-r}$ of GG-A2-GG6’s figures is the same ladder shifted by that constant, and is not used for any number here). Numerically, $e^{-2\pi} = 1.867 \times 10^{-3}$ and $e^{-K} = e^{-35} = 6.305 \times 10^{-16}$, so

$$\begin{aligned} E_{\text{GUT}} &= 1.22 \times 10^{19} \times 1.867 \times 10^{-3} = 2.28 \times 10^{16} \text{ GeV}, \\ M_{\text{lock}} &= \frac{1}{2} \times 1.22 \times 10^{19} \times 6.305 \times 10^{-16} = 3.85 \text{ TeV}. \end{aligned} \quad (6.10)$$

A useful consistency check a student can do in one line: the ratio of the two rungs must be a pure exponential of ledger data,

$$\frac{E_{\text{GUT}}}{M_{\text{lock}}} = 2 e^{K-2\pi} = 2 e^{28.72} = 5.9 \times 10^{12}, \quad (6.11)$$

and indeed $2.28 \times 10^{16} / 3.85 \times 10^3 = 5.9 \times 10^{12}$. The electroweak scale then follows from the Wilson–Hosotani vev theorem (GG-A9-SM), $v = M_{\text{lock}}(2/L)^N = 3849 \times (0.4)^3 = 246.3$ GeV — three rungs, zero dials. The Kaxion will add a fourth rung at the bottom of this ladder, $m_K \sim 10^{-14}$ GeV, generated not by the scale profile but by topology (Section 9.6); the ladder is how a single rigid geometry holds objects thirty orders of magnitude apart without a single tuned number.

6.3 BI-6 and the forced tensor sector

6.3.1 A two-paragraph anomaly primer

A gauge theory of *chiral* fermions — fermions whose left- and right-handed components carry different charges, as the Standard Model’s do — faces a quantum subtlety. Currents that are conserved classically can fail to be conserved at one loop: the famous triangle diagram (a fermion loop with three gauge-boson legs) produces $\partial_\mu J^\mu \neq 0$. If the violated current is one that the gauge symmetry itself depends on, the theory loses unitarity and is simply inconsistent — an anomaly is not an inconvenience but a death sentence. Consistency therefore imposes exact arithmetic conditions on the fermion charges; the miraculous-looking anomaly cancellations of the Standard Model are an example.

In six dimensions the bookkeeping is done by an *anomaly polynomial*: an eight-form I_8 , built from traces of curvature two-forms, that encodes every one-loop anomaly at once via the descent procedure. Green and Schwarz discovered (Green & Schwarz, 1984) that even when $I_8 \neq 0$, the theory can be rescued *if* the polynomial factorizes,

$$I_8 = \frac{1}{2} X_4 \wedge X_4, \quad (6.12)$$

because then a single two-form field B_2 , transforming anomalously and coupled through an inflow term, cancels the entire defect. The price of the rescue is the permanent presence of B_2 in the spectrum. That price is the Kaxion’s birth certificate.

6.3.2 The rank-one criterion and the locked coefficients

The anomaly paper GG-A3-BI6 (Arisaka, 2026, A3-BI6) proves the sharp version for GG-6: a single Green–Schwarz tensor can cancel the reducible anomaly of (6.5) *if and only if* the anomaly matrix has rank one — and for the GG-6 fermion content it does, with the unique primitive coefficient vector $(a, b, c, d) = (1, -1, -\frac{2}{5}, -\frac{2}{5})$. The theory is therefore forced to contain B_2 with the Chern–Simons-improved field strength

$$H_3 = dB_2 + \omega_3(\Omega) - \omega_3(A_5) - \frac{2}{5}\omega_3(A_B) - \frac{2}{5}\omega_3(A_{Q_i}), \quad (6.13)$$

obeying the modified Bianchi identity that names the series:

$$dH_3 = X_4 = \text{tr } R^2 - \text{Tr } F_5^2 - \frac{2}{5} F_B^2 - \frac{2}{5} F_{Qi}^2, \quad (6.14)$$

together with the inflow action

$$S_{\text{GS}} = \int_{\mathcal{M}_6} B_2 \wedge X_4. \quad (6.15)$$

The tensor is anti-self-dual at the quantum level (GG-A3-BI6), and nothing about it is adjustable. In particular the fraction $\frac{2}{5}$ is not modeling; it is the trace-norm ratio of the baryon and hypercharge generators, computable in two lines. With

$$Y_0 = \text{diag}\left(-\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}, \frac{1}{2}, \frac{1}{2}\right), \quad B_0 = \text{diag}\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0\right) \quad (6.16)$$

(the hypercharge and baryon-number generators in the fundamental of $\text{SU}(5)$),

$$\frac{\text{Tr}(B_0^2)}{\text{Tr}(Y_0^2)} = \frac{3 \times (1/3)^2}{3 \times (1/3)^2 + 2 \times (1/2)^2} = \frac{1/3}{5/6} = \frac{2}{5} = \frac{L - N}{L}. \quad (6.17)$$

The ratio of the gravitational entry to either Abelian entry of the coefficient vector is thus locked at $|a/c| = |a/d| = \frac{5}{2}$ — the integer pair $(5, 2)$ that fixes the ledger's baryonic entry, which opens no channel below M_{lock} (Part VI).

Physical picture. The tensor sector is the theory's topological bookkeeper. Every gauge factor writes its curvature into the ledger X_4 with a fixed weight, and B_2 is the page on which the entries are reconciled. One does not get to decline the bookkeeper: without it the theory is anomalous, i.e. not a theory. Two facts matter downstream: the bookkeeper *exists*, with exactly one physical degree of freedom to spare (Section 9); and its entries come in the fixed proportion $1 : -1 : -\frac{2}{5} : -\frac{2}{5}$, so anything that later couples through the ledger inherits integer-locked ratios.

6.3.3 How the inflow cancels the anomaly: descent in three lines

For the student who wants the mechanism and not just its name. (1) The one-loop anomaly of the 6D chiral theory is summarized by the eight-form I_8 through the *descent equations*: write $I_8 = dI_7$, and let the gauge variation of I_7 define $\delta_\lambda I_7 = dI_6^{(1)}(\lambda)$; then the anomalous variation of the effective action is $\delta_\lambda \Gamma = 2\pi \int_{\mathcal{M}_6} I_6^{(1)}(\lambda)$. (2) If, and only if, the polynomial factorizes as in (6.12), the descent of $I_8 = \frac{1}{2} X_4 \wedge X_4$ can be arranged as $I_6^{(1)} = X_2^{(1)}(\lambda) \wedge X_4$, with $X_2^{(1)}$ the descent of X_4 itself. (3) Assigning the tensor the anomalous transformation $\delta_\lambda B_2 = -X_2^{(1)}(\lambda)$ makes the inflow term shift by $\delta_\lambda S_{\text{GS}} = -\int X_2^{(1)} \wedge X_4 = -\delta_\lambda \Gamma$: the two variations cancel identically, to all orders, for any gauge parameter. The rank-one criterion of GG-A3-BI6 is precisely the statement that step (2) is possible — the anomaly matrix M_{ij} (the coefficient of $\text{tr } F_i^2 \text{ tr } F_j^2$ in I_8) factorizes as $M_{ij} = x_i x_j$ if and only if it has rank one, and the vector $x_i \propto (1, -1, -\frac{2}{5}, -\frac{2}{5})$ is then unique up to overall normalization.

Physical picture. The anomaly is a leak: gauge charge seeps out of the fermion sector at one loop. The Green–Schwarz tensor is a drainpipe fitted to the leak: it accepts exactly the seepage ($\delta B_2 = -X_2^{(1)}$) and discharges it through the inflow coupling. A drainpipe only works if all the

leaks run in proportion — one pipe, one cross-section, hence rank one. The Kaxion, as the dual of this pipe, inherits its plumbing: the proportions $(1, -1, -\frac{2}{5}, -\frac{2}{5})$ are not decorations but the leak rates the pipe was built to match.

6.4 GDOS: from the bulk to the boundary observer

6.4.1 Why a projection is constitutional

Everything above lives in the six-dimensional bulk. No experiment does. Every measurement ever performed — every spectral line, every collider event, every haloscope scan — is an event registered by an observer confined to the four-dimensional boundary. The OGGEP axiom elevates this truism to a constitutional demand: the physical content of the theory *is* its boundary registration, full stop. The machinery that executes the demand is GDOS — Geometric Dynamics of Open Systems — constructed in GG-A5-GDOS (Arisaka, 2026, A5-GDOS).

6.4.2 A one-paragraph open-systems primer

A quantum system that is part of a larger whole does not evolve unitarily. Tracing out the unobserved remainder leaves the subsystem with *open* dynamics: a completely positive, trace-preserving (CPTP) channel, the general framework of decoherence theory, Lindblad equations, and quantum noise. The familiar slogan is that the subsystem sees a reduced, noisy, dissipative image of the whole. GDOS applies this with the bulk as the whole and the boundary as the subsystem: the boundary observer does not see the bulk’s wavefunction; the observer sees its CPTP shadow.

6.4.3 What GDOS proves, and what the Kaxion takes from it

Two results of GG-A5-GDOS carry this paper. First, the canonical-boundary-law theorem (T6): for any theory satisfying the constitution, the boundary content is forced into one canonical law, expressible in three equivalent dress codes — operator (a Lindblad-type channel), mesoscopic (a probability-flow equation), and Langevin (stochastic trajectories). One law, three notations; the cosmological equation of motion (12.2) of Part III is its homogeneous classical limit. Second, the four-currents Ward package (T9.7): the boundary carries exactly four conserved currents —

$$\partial_\mu J_a^\mu = 0 \quad (a = \text{em}, B, Qi), \quad \nabla_\mu T^{\mu\nu} = 0, \quad (6.18)$$

— protected against anomalous breaking by the BI-6 inflow of Section 6.3. The bulk’s anomaly bookkeeping thus reappears on the boundary as conservation law, which is why the *ratios* written into X_4 are measurable in a laboratory at all.

Two derived objects knit the package together and will be referred to by name. The Total Geometric–Gauge object TGG is the unified stress tensor (matter + gauge + Green–Schwarz + Kaluza–Klein) with one master conservation law: it is what places inflation, dark matter, and dark energy inside a single conserved cosmic packet in Part III. The Covariant Conservation Identity CCI is the seam law — BI-6 inflow plus the four Ward identities — governing how any bulk-descended field, the Kaxion included, may lawfully couple on the boundary. Beyond these, GG-A7-Qi4 (Arisaka,

2026, A7-Qi4) proves that the four currents of (6.18) are not independent but locked into one operational identity (the Arisaka equation, $J_E/Q_B = J_e = J_B = J_{Q_i}$); the dark-sector reading of that lock is the transport of the ledger’s portal weights to the boundary, used in Part V.

Physical picture. Think of the bulk as the full orchestral score and the boundary as what reaches a listener in the hall: filtered, reverberant, noisy — but lawful, with conserved quantities surviving the passage. GDOS is the acoustics of that hall. The Kaxion is a bulk instrument; everything we will ever measure of it — its oscillation, its portals, its one tone — is its GDOS sound.

6.4.4 The boundary law in its three dress codes, displayed

GG-A5-GDOS’s canonical-boundary-law theorem (T6) asserts one law in three equivalent notations; for later use we display all three. In the *operator* dress code the boundary state evolves by a Lindblad-type CPTP generator,

$$\dot{\rho} = -i[H_{\text{eff}}, \rho] + \sum_k \left(L_k \rho L_k^\dagger - \frac{1}{2} \{L_k^\dagger L_k, \rho\} \right); \quad (6.19)$$

in the *mesoscopic* dress code the probability density on configuration space obeys a Smoluchowski flow,

$$\partial_t P = \nabla \cdot \left(\frac{P \nabla F}{\gamma} \right) + \frac{T_{\text{eff}}}{\gamma} \nabla^2 P; \quad (6.20)$$

and in the *Langevin* dress code individual trajectories descend the same landscape with noise,

$$\gamma \dot{x} = -\nabla F + \sqrt{2\gamma T_{\text{eff}}} \xi(t), \quad \langle \xi(t) \xi(t') \rangle = \delta(t - t'). \quad (6.21)$$

One law, three costumes: (6.19) for quantum readout chains (the detector physics of Part VII), (6.20) for ensembles, (6.21) for histories.

The payoff for this paper is a derivation the reader can now perform in one line. Apply (6.21) to the homogeneous compact-angle mode in an expanding background: the configuration variable is a_K , the landscape is the potential energy, the friction coefficient supplied by the expanding background is $\gamma = 3H$ (the standard redshift drain of field momentum), and the noise temperature after inflation is utterly negligible ($T_{\text{eff}} \sim H/2\pi \lesssim 10^{-13}$ eV at onset, thirty orders below the barrier scale). Restoring the inertial term, the boundary law reads

$$\ddot{a}_K + 3H\dot{a}_K + V'(a_K) = 0$$

— exactly Eq. (12.2), the equation on which all of Part III runs. The cosmological pendulum is not an analogy bolted onto the theory; it *is* the GDOS boundary law, worn in its Langevin costume, with the universe itself as the bath.

6.5 GRIP: the runtime that selects vacua

6.5.1 Why a Lagrangian is not enough

A Lagrangian lists possibilities; the universe realizes one of them. Which vacuum, which symmetry-breaking pattern, which Wilson-line phase? In most frameworks these are chosen by hand or by anthropics. The series instead supplies dynamics: GRIP, the Gauge-Runtime Invariance Pipeline of GG-A6-GRIP (Arisaka, 2026, A6-GRIP), a gradient flow on the space of boundary configurations,

$$\dot{x} = -g^{-1} \nabla F_{\text{GFF}}(x), \quad (6.22)$$

descending a single scalar landscape until it stops. Everything the runtime needs is packed into that landscape, so we build it first.

6.5.2 A Morse-theory primer in one figure’s worth of words

Let F be a smooth function on an n -dimensional configuration space. Its *critical points* are where $\nabla F = 0$; at each, the Hessian matrix $\partial^2 F$ has some number k of negative eigenvalues, called the *Morse index* — the number of independent downhill directions. Index 0 is a minimum (nowhere to fall: stable); index n is a maximum (everywhere to fall); intermediate indices are saddles. Under the gradient flow (6.22), almost every trajectory leaves high-index critical points along their downhill directions and terminates at an index-0 minimum. Morse theory is the statement that this skeleton of critical points and connecting flows captures the global topology of the landscape. That is the entire toolkit needed here: *the index counts the ways to fall, and generic histories end at index zero.*

A worked example on the cylinder. Abstractions stick best with a computation, so run the toolkit once on the simplest landscape this paper will actually use: $F(\theta, y) = \chi(1 - \cos \theta) + \frac{1}{2}M^2 y^2$ on the cylinder $S^1 \times \mathbb{R}$ (one compact angle, one heavy transverse direction — a cartoon of the compact-angle stratum). Critical points: $\nabla F = (\chi \sin \theta, M^2 y) = 0$ gives $(\theta, y) = (0, 0)$ and $(\pi, 0)$. Hessians: $\text{diag}(\chi, M^2)$ at the origin — both eigenvalues positive, Morse index 0: this is the **GA**. At $(\pi, 0)$: $\text{diag}(-\chi, M^2)$ — one negative eigenvalue, index 1: a **GSSB** saddle, the watershed of the circle. (The DGI is the $\chi \rightarrow 0$ limit in which the whole circle is flat and fully symmetric — the configuration before the potential turns on.) The gradient flow $\gamma \dot{\theta} = -\chi \sin \theta$ even integrates in closed form,

$$\tan \frac{\theta(t)}{2} = \tan \frac{\theta_i}{2} e^{-\chi t / \gamma}, \quad (6.23)$$

exhibiting the whole story in one formula: release anywhere but exactly on the saddle and the trajectory relaxes exponentially into the attractor, at a rate set by the curvature χ — the same curvature that will reappear in Eq. (10.1) as the Kaxion’s mass squared. Add inertia and weaken the friction and the same descent arrives *ringing*; that ringing, on the real stratum, is Part III.

6.5.3 The landscape: a free energy, not an energy

The landscape GRIP descends is not the bare energy. GG-A6-GRIP (structure theorem GFF-ST2) proves it splits as

$$F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}, \quad (6.24)$$

where F_{GEP} — the Geometric Energetic Part — is the bulk-to-boundary descent of the bulk TGG action (the ordinary energy cost of a configuration), and F_{CEP} — the Curvature-Entropic Part —

$$F_{\text{CEP}} = -T\mathcal{S} + F_{\text{diff}} \quad (6.25)$$

adds the entropic and diffusive contributions of the boundary’s open-system (CPTP) dynamics, at an effective temperature set by the bulk scale profile. The structure is exactly that of a thermodynamic free energy, and the intuition transfers whole: the runtime does not seek the configuration of least energy but of least *free* energy — energy and entropy bargaining, as in any statistical system. A reader who has minimized a Landau free energy has done this before.

6.5.4 The cascade and the uniqueness theorem

On each stratum of configuration space, the critical points of F_{GFF} organize into a strict descent, which the series names once and uses everywhere:

$$\text{DGI} \longrightarrow \text{GSSB} \longrightarrow \text{GA}. \quad (6.26)$$

At the top, the DGI (Dynamical Gauge Invariance): the critical point of *maximal* Morse index, the fully symmetric configuration in which nothing has committed — every gauge equivalence intact, every direction downhill. In the middle, the GSSB chain (Geometric Spontaneous Symmetry Breaking): the sequence of saddles of decreasing index, each a partial commitment, each step shedding two indices and carrying a quantized action ($S = \pi k$; GG-A6-GRIP). At the bottom, the GA (Geometric Attractor): the index-zero minimum, stable against all small perturbations — in the series’ vocabulary, a *stable boundary observable*. The keystone is the equilibrium-uniqueness theorem (GFF-ST7):

$$[\Pi_*] = \arg \min_{[\Pi] \in \text{Bdy}} F_{\text{GFF}}([\Pi]) \quad \text{exists and is unique,} \quad (6.27)$$

and is empirically identified with the Standard-Model vacuum. The same pipeline fixes the Wilson-line breaking $\text{SU}(5) \rightarrow \text{SM}$ (Hosotani, 1983), the Higgs vev of Table 2, and — the business of Section 10 — the birth and stabilization of the Kaxion condensate.

Physical picture. A terraced funnel. The DGI is the rim; the GSSB saddles are the terraces; the GA is the floor. GFF-ST7 says the funnel has one floor. Release the universe at the rim and its destination is fixed; its *history* — which terraces, what ringing on arrival — is the physics. Part III is the story of the ringing.

6.5.5 From GA to GGA: the ground-state refinement

One refinement of the cascade vocabulary, supplied by the capstone paper GG-A11-Origin (Arisaka, 2026, A11-Origin), matters enough for this paper to adopt throughout. A Morse landscape generally admits *many* stable basins — every index-zero minimum is a GA — but the cascade output is not “some stable basin”: it is the *unique deepest* one, the global minimum of F_{GFF} , which GG-A11-Origin names the *Ground Geometric Attractor*,

$$\boxed{\text{GGA} := \text{GA}^* = \arg \min_{x \in \mathcal{S}} F_{\text{GFF}}(x)} \quad (\text{GG-A11-Origin, T-GGA-as-Ground-State}), \quad (6.28)$$

with the selection of GA^* from among the GAs enforced by the four Qi-4 admissibility pillars of GG-A7-Qi4 (T-QUA, T-EP, T-CON, T-HAL; consolidated there, with the pillar-by-pillar division of labor, as the proposition “Qi-4 as the GGA selector”). In this vocabulary the full runtime chain reads $\text{GRIP} := \text{DGI} \rightarrow (\text{GFF}) \rightarrow \text{GSSB} \rightarrow \text{GGA}$, and GG-A6-GRIP’s equilibrium-uniqueness theorem GFF-ST7 (6.27) is recognized as precisely the GGA statement *avant la lettre* — the existence and uniqueness of the ground attractor, before it had its name.

Physical picture. Chemistry made this distinction a century ago. A molecule’s electronic landscape has many stationary configurations, but chemistry is built on the *ground state*: every water molecule is identical because its electrons do not merely occupy “a” stable orbital configuration but *the* lowest one. The GGA is that idea promoted to a constitutional principle: cosmos-scale structures, like molecules, are reproducible and stable because the runtime terminates not at an attractor but at the ground state of the attractor manifold. GG-A11-Origin runs the principle through eighteen layers (T-Atomic-Ground-State, T-Molecular-Closure, ...); GG-A8-QM uses it to prove that every electron is the same electron. Section 10.4 will run it on the Kaxion — with consequences for both its stability and its coherence that the weaker statement “the Kaxion is a GA” could not deliver.

6.5.6 The two regimes: GRIP-In and GRIP-Out

One further piece of runtime vocabulary, defined in GG-A6-GRIP (Definitions def:GRIPin and def:GRIPout) with the regime theorems proved in GG-A11-Origin (T-GRIP-In-Lyapunov-Descent, T-GRIP-Out-Necessity), completes the review. The runtime operates in exactly two regimes, split by one question: *does the landscape change?* GRIP-In is all dynamics on a *fixed* landscape — gradient flow within a basin, instanton steps across the GSSB saddles between basins — halting uniquely at the stratum’s GGA. GRIP-Out is the complementary regime: a transient out-of-equilibrium event that *reorganizes the landscape itself* and supplies the initial datum from which the next descent begins. In one sentence: GRIP-Out opens the landscape; GRIP-In finishes the journey. (The runtime paper now carries this picture as a full-page figure — the cascade column with its way-station GA tier and the two regime branches — to which the reader is referred once: GG-A6-GRIP, figure “the GRIP cascade and its two regimes.”)

Physical picture. On the matter strata, Nature implements GRIP-In through the weak charged current — every decay chain $\tau \rightarrow \mu \rightarrow e$ is GRIP-In on a fixed family landscape — while the one

GRIP-Out of that story was the electroweak transition that built the landscape (GG-A9-SM). The dark sector will turn out to be the purest GRIP-In specimen of all: the Kaxion is released in a GRIP-Out (the lock crossing; Section 10), relaxes on a fixed compact-angle landscape with *Hubble friction itself* as the within-basin flow, and — because its exit-channel set is structurally empty (Section 10.4) — its halt at the GGA is permanent. The entire thirteen-billion-year biography of Part III is one uninterrupted GRIP-In.

6.6 The tour in one paragraph

GG-6 supplies a rigid stage: four spacetime dimensions, a scale dimension of depth $K = 35$, a non-spatial phase circle of unit radius, a four-factor gauge group, and five integers with a documented track record. BI-6 forces onto that stage a tensor bookkeeper B_2 with one spare degree of freedom and its Abelian entries locked at $-2/5$. GDOS projects everything to the boundary where observers live, preserving exactly four currents. GRIP then decides occupancy: a unique free-energy minimum, reached by a named cascade. The remainder of this Part assembles these four facts into a theorem: the stage, the bookkeeper, the projection, and the runtime jointly force a light pseudoscalar — and Parts III and V will show that this pseudoscalar, with one calibrated scale and no further input, is the dark matter.

7 The five-link existence chain

With the framework in hand, the existence argument is a chain of five links,

$$\begin{array}{ccccccc}
 \underbrace{\mathcal{M}_4 \times I_r \times S_\theta^1}_{\text{Link 1: GG-A2-GG6 T1-T2}} & \implies & \underbrace{(1, 3, 5; 7, 35)}_{\text{Link 2: GG-A2-GG6 T6}} & \implies & \underbrace{B_2, H_3}_{\text{Link 3: GG-A3-BI6}} & \implies & \underbrace{\theta_{GS}, \theta_W}_{\text{Link 4: boundary duals}} \implies \underbrace{\text{light eigenstate}}_{\text{Link 5: GG-A6-GRIP}} \\
 & & & & & & (7.1)
 \end{array}$$

each link a numbered theorem of the series, none proved for dark-matter reasons. We walk the chain once, link by link, extracting from each exactly what the Kaxion needs.

7.1 Link 1: the stage provides the circle

From Section 6.2: the bulk contains the non-spatial phase circle S_θ^1 , rigid ($G_{\theta\theta} = 1$, no modulus) and integer-quantized through its holonomy (6.4). Two features carry all the downstream weight. First, the circle is *non-spatial*, so a field living on it is automatically a boundary scalar — neutral under boundary translations, carrying no spatial momentum tower of its own. Second, the circle is *rigid* — nothing about it can be tuned to accommodate a desired phenomenology. Whatever the circle produces, it produces at fixed gear ratios. Both features are MSMF at work: a tunable radius would be a continuous dial, and the constitution forbids it.

7.2 Link 2: the ledger provides the gears

From Section 6.2: the five integers (6.6) with the dual lock (6.8). We display the ledger as its own link because the Kaxion will be read off it repeatedly, and the reader should recognize each appearance as the *same* five integers rather than as new inputs. In this paper: $K = 35$ is the scale

depth that sets the hierarchy (6.3) *and* fences the laboratory window of Part V; $L = 5$ fixes the unification embedding whose anomaly ratio $E/N = 8/3$ sets the photon coupling of Part V; $J = 7$ closes the dyadic coupling $\alpha_{\text{lock}}^{-1} = 2^J = 128$ of Table 2; and the pair $(L - N, L) = (2, 5)$ is the trace ratio (6.17) behind the Kaxion’s baryonic entry. One ledger, many readouts. A dark-matter candidate whose every observable is a gear ratio of $(1, 3, 5; 7, 35)$ is either right in a highly testable way or wrong in a highly testable way — exactly the position a predictions paper wants to occupy.

7.3 Link 3: anomaly freedom provides the tensor

From Section 6.3: the rank-one criterion forces B_2 with field strength (6.13), Bianchi identity (6.14), inflow term (6.15), and locked coefficients (6.17). The link to retain: the tensor’s existence is precisely as negotiable as the consistency of the quantum theory — that is, not at all — and it brings exactly one light pseudoscalar degree of freedom to the boundary (counted in Section 9.1; the anomaly companion’s mode analysis also registers a decoupled vector zero mode, Section 9.4) plus a ledger of locked weights.

7.4 Link 4: the boundary provides its pseudoscalars — and exactly one stays light

Projected through GDOS to the boundary, Links 1 and 3 leave behind periodic, parity-odd directions — the dual of the tensor field strength (the *tensor remnant* θ_{GS} , the tensor shadow) and the circle’s holonomies promoted to fields (written θ_W for the direction the two-field Lagrangian below tracks) — of which exactly one stays light. Their joint boundary-effective Lagrangian, derived in GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos),

$$\mathcal{L} = \frac{1}{2} K_{ij} \partial_\mu a_i \partial^\mu a_j - V(c_i a_i / f_i) + \sum_F \frac{C_F^i g_F^2}{8\pi^2} \frac{a_i}{f_i} F \tilde{F}, \quad i \in \{\text{GS}, \theta\}, \quad (7.2)$$

contains nothing postulated: the field content is Links 1 and 3, the periodicity is Link 1, the portals are Link 3 transported by the CCI, and the absence of any non-derivative, non-topological coupling is the shift symmetry. The two directions appear symmetrically at this level; the asymmetry — which of them stays light — is dynamical, and Section 9 settles it: the holonomy directions are lifted at the circle’s curvature scale by the Hosotani mechanism (GG-A9-SM registers the flat zero mode of the coset holonomy near 2×10^{15} GeV, PS-SM-22, and leaves the identity of the electroweak Higgs open, O₂₂), and the light spectrum below the lift has exactly one member from the compact-angle sector, the tensor’s dual. The detailed construction of both directions — the duality transformation, the degree-of-freedom count, the parity and periodicity arguments, and the lift — is the subject of Section 9.

7.5 Link 5: the runtime provides occupancy

Diagonalizing (7.2) yields a light eigenstate protected by the residual shift symmetry, the holonomy directions having been lifted at the circle’s curvature scale (Section 9.3). The runtime of Section 6.5 then converts the light eigenstate from a Lagrangian entry into a physical occupant:

by GFF-ST7 the late-time boundary configuration is the unique free-energy minimum, and on the compact-angle stratum that minimum — the **GA** — is the stabilized light eigenstate, with its cosmological displacement and relaxation forming a **GRIP** event chain (Section 10). With all five links in place we can state the existence theorem. Table 3 assembles the chain in one view — each link’s parent theorem, the object it delivers, and where this paper inspects the delivery; the theorem below is precisely the conjunction of its five rows.

Table 3: The five-link existence chain. Each link is a published theorem of the GG-A series; none was proved for dark-matter reasons. Each link is a theorem that can be checked on its own, and the last column says what the link buys. The chain is what makes the Kaxion an existence claim rather than a proposal, because no link was added in order to reach it. If any one of the five fails the Kaxion does not exist, which is the sense in which this is also a falsification map.

Link	Parent result	Statement	What it delivers to the Kaxion
1	GG-A2-GG6 T1–T2	Bulk $\mathcal{M}_4 \times I_r \times S_\theta^1$; canonical metric, no free radius	The non-spatial phase circle: a rigid, periodic, integer-quantized home for a compact-angle field.
2	GG-A2-GG6 T6	OGN-5 = (1, 3, 5; 7, 35), $K = LJ = M^2 + N^2 + L^2$	The integers behind every later observable: $K = 35$ anchors the benchmark loading; N, L, J furnish the ladder (an ansatz, Part V).
3	GG-A3-BI6 rank-one criterion	B_2/H_3 forced; BI-6 with $(1, -1, -\frac{2}{5}, -\frac{2}{5})$	The tensor sector exists; portal weights locked at the trace ratio 2/5.
4	Boundary reduction (GG-A5-GDOS, GG-A10-Cosmos)	H_3 dual θ_{GS} ; circle holonomies θ_W	Periodic pseudoscalar directions, shift-symmetric, parity-odd; the holonomies are lifted and the tensor dual stays light (count graded conditionally, T-Kaxion-5).
5	GG-A6-GRIP GFF-ST7	$[\Pi_*] = \arg \min F_{\text{GFF}}$	The light eigenstate is selected and stabilized as the Geometric Attractor.

8 The Kaxion existence theorem

Definition 8.1 (Kaxion). The *Kaxion* is the single light compact-angle pseudoscalar of the boundary theory: the canonically normalized boundary dual of the forced Green–Schwarz field strength H_3 — the tensor remnant, with its full X_4 couplings — surviving at energies far below the throat scales, selected and stabilized as the *Ground Geometric Attractor* (**GGA**, the unique deepest attractor of the stratum; Section 10.4) of the **GRIP** cascade on the compact-angle stratum.

Theorem 8.2 (T-Kaxion-1: existence of the compact-angle pseudoscalar). *Assume the constitutional axioms and postulates of GG-A1-Constitution, and accept the following published results of the series: the bulk and metric theorems GG-A2-GG6 T1–T2, the gauge-content theorem GG-A2-GG6 T3, the OGN-5 cascade GG-A2-GG6 T6, the rank-one criterion and BI-6 coefficient theorem of GG-A3-BI6, the boundary projection of GG-A5-GDOS (T6, T9.7), and the equilibrium-uniqueness theorem GFF-ST7 of GG-A6-GRIP. Then the four-dimensional boundary theory contains exactly one light periodic*

pseudoscalar field with the following properties, none of which involves a new field, a new symmetry, or a new continuous parameter:

1. **(Existence and ancestry)** *It is the boundary dual of the forced Green–Schwarz field strength H_3 — the one light pseudoscalar the anti-self-dual tensor leaves on the boundary — and within the compact-angle sector it is alone: the holonomies of the non-spatial circle S_θ^1 are Hosotani-lifted and supply no light QCD-coupled mode. Whether it is the only light QCD-coupled pseudoscalar of the boundary theory as a whole is the count, stated separately and graded conditionally on the parents’ registered exposures (Section 9.4, T-Kaxion-5).*
2. **(Protection)** *Its potential vanishes to all orders in perturbation theory: the shift symmetry inherited from (6.15) and (6.4) is broken only by topological (anomaly-descended) effects.*
3. **(Locked portals)** *Its non-derivative couplings descend exclusively from the BI-6 ledger X_4 , and therefore occur in the fixed proportions of the coefficient vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$; the baryonic entry carrying the trace-ratio coefficient $-\frac{2}{5}$ of hypercharge’s under the anomaly companion’s normalization, with no open channel below M_{lock} .*
4. **(Stability of occupation)** *The configuration in which the light eigenstate sits at a displaced compact angle and relaxes is a GRIP runtime event — a GRIP-In descent on the fixed compact-angle landscape — terminating on the stratum’s Ground Geometric Attractor (GGA); the field is dynamically selected — and uniquely so — not merely kinematically allowed.*

Proof. Existence and ancestry: by GG-A2-GG6 T1–T2 the bulk contains the rigid circle S_θ^1 , whose holonomy (6.4) promotes to the boundary field (9.3); by the rank-one criterion of GG-A3-BI6 the bulk must contain B_2 with field strength (6.13), whose single physical degree of freedom dualizes on the boundary to the pseudoscalar (9.2). The circle’s holonomies are those of the rank-six group $\text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}$ and no others: GG-A2-GG6 T1 forbids further compact factors, and GG-A2-GG6 T3 forbids further gauge factors. Of these directions only the tensor dual stays light: the Cartan holonomies acquire Hosotani potentials at the circle’s curvature scale (Section 9.2), and the mode analysis of GG-A3-BI6 (its section on the self-duality of the Green–Schwarz tensor; O13/O14) leaves the tensor exactly one light pseudoscalar — one chirality’s zero mode, the opposite chirality absent — together with a decoupled vector zero mode and, under the boundary-condition termination, two interval Wilson lines, all listed with their grades in Section 9.4.

Protection: B_2 enters the bulk action only through H_3 and S_{GS} , and A_θ enters boundary observables only through W ; both statements forbid any perturbative potential for the duals, leaving only the topological breakings carried by X_4 (the explicit argument is Section 9.5). Locked portals: the boundary couplings of both directions descend from (6.15) by the descent procedure of GG-A5-GDOS (T9.7 anomaly-inflow package), which transports the coefficient vector of (6.14) unchanged; the Abelian entries carry the trace-ratio coefficient $-2/5$ of (6.17).

Stability: GFF-ST7 of GG-A6-GRIP guarantees a unique global minimizer of F_{GFF} on the boundary moduli space; restricted to the compact-angle stratum the minimizer is a discrete OGN-5 phase (Section 10), and the light eigenstate relaxing toward it is by definition a $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ cascade, terminating — by T-GGA-as-Ground-State of GG-A11-Origin (Arisaka, 2026, A11-Origin) — on the stratum’s unique deepest minimum, its GGA, the attractor from which no admissible exit

channel remains (T-GGA-No-Exit, *ibid.*). Uniqueness of the light *pseudoscalar within the tensor sector* follows because the opposite chirality is absent; uniqueness across the whole boundary theory is the count, T-Kaxion-5, conditional on the registered exposures of Section 9.4. ■

Corollary 8.3 (Automatic dark-matter credentials). *The field of Theorem 8.2 is electrically neutral and colorless (it is a gauge singlet of the boundary group); its every coupling is suppressed by one power of a decay constant f_K far above the electroweak scale; and it is therefore cosmologically long-lived. (The quantitative lifetime, $\tau_{K \rightarrow \gamma\gamma} \gg t_U$ by many tens of orders of magnitude, is computed in Part V.)*

Remark 8.4 (What is claimed, and what is not). Theorem 8.2 is deliberately silent about three numbers: the decay constant f_K , the absolute mass m_K , and the normalization of the physical eigenstate. Those are fixed in Parts III and V by the π -locked abundance closure and the susceptibility relation, and they rest on one calibration (the bulk normalization, O14) and the named conjectures of the deck (Table 41), fenced as such. The theorem-level content is the existence, the protection, and the ledger-locked portals; the count — exactly one light field, hence one line — is stated separately at derived-conditional grade (T-Kaxion-5, §9.4). The reader should hold this division firmly: the experiments of Parts VII–VIII test the count even if the benchmark shifts, and the window moves with the calibration.

Remark 8.5 (Least-new sector: from “may” to “must”). The constitutional rule MSMF (Section 6.1) does more than permit the Kaxion: it *orders the search*. A new dark sector — a WIMP multiplet, a sterile fermion, a hidden gauge group — would introduce fields and continuous parameters that the constitution forbids unless absolutely required. A sector already forced by anomaly freedom introduces nothing. MSMF therefore converts “the compact-angle pseudoscalar may be the dark matter” into “the compact-angle pseudoscalar must be the *first* candidate examined, and a new sector may be contemplated only if it fails.” Parts III and V show that it does not fail.

9 The clock and the lifted dial: constructing the light mode

Theorem 8.2 leans on the claim that the boundary inherits the tensor’s dual and the rank-six holonomy angles of the circle, and that exactly one light state survives from that sector. This section constructs the two representative directions explicitly, at the level of detail a graduate student can audit: degree-of-freedom counts, the duality transformation, parity, periodicity, the lift of the holonomy direction, and the mass-protection argument with numbers.

9.1 The tensor remnant: dualizing B_2

Counting first. A massless p -form gauge field in D dimensions carries $\binom{D-2}{p}$ physical polarizations — the transverse little-group count, exactly as the photon’s $D - 2$. For the two-form on the four-dimensional boundary: $\binom{2}{2} = 1$. One degree of freedom, and a scalar one. The dualization below is therefore not a sleight of hand but a change of variables between two descriptions of the *same* single degree of freedom.

The duality transformation. Start from the first-order form of the boundary tensor action, treating H_3 as an independent field with a Lagrange multiplier a_{GS} enforcing its Bianchi identity:

$$\mathcal{L}_{1\text{st}} = -\frac{1}{12} H_{\mu\nu\rho} H^{\mu\nu\rho} + \frac{1}{6 f_{\text{GS}}} a_{\text{GS}} \epsilon^{\sigma\mu\nu\rho} \partial_\sigma H_{\mu\nu\rho} \longrightarrow \begin{cases} \delta a_{\text{GS}} : & dH_3 = 0 \Rightarrow H_3 = dB_2 \text{ locally,} \\ \delta H_3 : & H^{\mu\nu\rho} = f_{\text{GS}}^{-1} \epsilon^{\mu\nu\rho\sigma} \partial_\sigma a_{\text{GS}}. \end{cases} \quad (9.1)$$

Varying the multiplier returns the original description; varying H_3 instead and substituting back yields $\mathcal{L} = \frac{1}{2}(\partial a_{\text{GS}})^2$: a canonical massless scalar. The two descriptions are exactly equivalent; we record the dictionary in the form used throughout the series,

$$H_{\mu\nu\rho} = f_{\text{GS}}^{-1} \epsilon_{\mu\nu\rho\sigma} \partial^\sigma a_{\text{GS}}. \quad (9.2)$$

Parity. The Levi-Civita symbol in (9.2) is parity-odd, while H_3 (a field strength of a tensor gauge field) is parity-even; consistency forces $a_{\text{GS}} \rightarrow -a_{\text{GS}}$ under spatial inversion. The dual is a *pseudoscalar* — which is why it can couple to $F\tilde{F}$ and why its condensate violates no parity test.

Periodicity. The large gauge transformations of B_2 are integrally quantized (the differential-cohomology refinement of GG-A3-B16; dually, the flux of H_3 through closed three-cycles is quantized, which fixes the winding charge); through (9.2) the first statement becomes the statement that a_{GS} is an *angle*: $a_{\text{GS}} \rightarrow a_{\text{GS}} + 2\pi f_{\text{GS}}$ is a gauge identification. We write $\theta_{\text{GS}} = a_{\text{GS}}/f_{\text{GS}} \in [0, 2\pi)$ and call it the tensor remnant (the tensor shadow): the four-dimensional trace of the bulk bookkeeper.

Why no potential. B_2 enters the bulk action only through H_3 (derivatives of B_2) and through the inflow term (6.15), which is topological — independent of the metric, hence contributing no energy density for constant fields. In the dual language: nothing in perturbation theory depends on the *value* of θ_{GS} , only on its gradients. The shadow’s potential is exactly flat until topology (instantons; Part V) intervenes.

9.2 The holonomy echo: the Wilson line as a field

The Aharonov–Bohm template. An electron transported around a solenoid acquires the phase $\exp(iq \oint A \cdot dl)$ — physical, gauge-invariant, and nonlocal: no local field along the path detects it. The Wilson loop (6.4) is the identical object wound around the phase circle of the bulk. Its logarithm,

$$\vartheta(x) = \oint_{S_\theta^1} A_\theta(x, \theta) d\theta, \quad a_\theta(x) = f_\theta \vartheta(x), \quad (9.3)$$

promoted to depend on the boundary point x , is a *bona fide* four-dimensional field: the holonomy echo $\theta_W = \vartheta$.

Periodicity and protection, again. Large gauge transformations — those that wind once around the circle — shift $\vartheta \rightarrow \vartheta + 2\pi$ while leaving every observable invariant; the echo is therefore an angle, exactly as the shadow is. And a *constant* Wilson line is locally pure gauge: no local

operator built from field strengths can give it a potential, since $F_{\theta\mu}$ vanishes for constant ϑ . Its potential, too, arises only from effects that wrap the circle — holonomy-sensitive, topological effects. The same mechanism, run with the bulk matter content, is the Hosotani lift (Hosotani, 1983): GG-A9-SM computes the coset holonomy’s radial profile — flat and ultraviolet-weighted — and the mass its holonomy would register, near 2×10^{15} GeV (PS-SM-22); the Hosotani potential itself is expected at the circle’s curvature scale and is not computed in any parent, and it is what removes these directions from the light spectrum — the holonomies acquire their potentials at the circle’s curvature scale, not at QCD’s, and are no light modes. Which mode is the 246 GeV electroweak field GG-A9-SM leaves open (O_{22}); nothing here depends on the answer. The Kaxion’s own potential arrives by a different route — the tensor’s inherited gluonic portal of Part V.

The Higgs and the Kaxion. It is tempting to place the two in one family — a scalar whose vacuum the holonomy selects and a pseudoscalar the tensor’s shift symmetry protects, both supplied by the compact geometry. The codebook paper’s present state does not support the first half: GG-A9-SM’s PS-SM-22 finds the coset holonomy’s zero mode registers far above the electroweak scale, and the identity of the 246 GeV field is its Open Problem O_{22} . What survives is the weaker statement, which is all this paper needs: the cohort expects few light scalars, each with a geometric reason. The Kaxion’s reason — a shift symmetry the bulk action cannot break perturbatively — is the one this Part establishes, and it does not lean on the Higgs’s.

Parity. Anomaly-descended $\vartheta F\tilde{F}$ operators fix the parity assignment of a compact angle as odd; both circle directions are pseudoscalars (the interval Wilson lines of §9.4 are a separate question). That they share this structure does not make them share a fate: parity is kinematics, and the lift is dynamics.

9.3 The kinetic matrix and the light eigenstate

Whether the directions mix. The Chern–Simons improvement in (6.13) contains the gauge potential whose holonomy is θ_W , so dimensional reduction of the $|H_3|^2$ kinetic term *can* generate an off-diagonal $\partial\theta_{GS}\partial\theta_W$ term alongside the diagonal ones. On the intended class — no gauge-field background — the quadratic action contains no such bilinear: the term that survives at that order mixes the tensor’s vector zero mode with A_μ , and a scalar–scalar overlap arises only with a background or at cubic order. The overlap integral of the two profiles (earlier revisions’ “gearing” $1/K$, now recorded in Appendix F) is therefore an upper bound on a mixing that need not be present at all. Nothing below depends on it, because the two directions do not both stay light.

The count: why one light QCD-coupled angle, and at what grade. The compact-angle sector supplies exactly one light QCD-coupled pseudoscalar, for two reasons that can be read off the parents. *The Hosotani lift:* the circle’s holonomies acquire their potentials from gauge and matter loops winding the circle at the circle’s curvature scale (GG-A9-SM PS-SM-22 for the coset direction; expected for the Cartan directions), so they are lifted far above the susceptibility scale and are no light modes. *Single inflow:* the boundary spectrum’s anomaly is carried once, by the

single tensor’s inflow, so there is no second anomalous current for a second angle to relax. One direction of the sector remains — the tensor’s dual, with its full X_4 couplings: the Kaxion. Whether it is the only light QCD-coupled pseudoscalar of the *whole* boundary theory depends on the five registered exposures — nine modes — the parents carry with open masses; they are listed, with the reason each is not expected to be a light QCD-coupled pseudoscalar and the grade of that reason, in Section 9.4, and the count is graded conditionally on them.

The cap diagonalization. The diagonalization that remains is light-against-heavy. Canonical normalization and a rotation bring the system to mass eigenstates:

$$\begin{pmatrix} a_K \\ a_H \end{pmatrix} = \begin{pmatrix} \cos \beta & \sin \beta \\ -\sin \beta & \cos \beta \end{pmatrix} \begin{pmatrix} a_{\text{GS}} \\ a_\theta \end{pmatrix}. \quad (9.4)$$

Here a_H is the heavy partner — dominated by the lifted holonomy direction, at the circle’s curvature scale (Appendix C.2) — and a_K is the canonically normalized light eigenstate, the Kaxion of Definition 8.1, which keeps the flat potential and the locked portals. The lift removes the heavy mode; it does not set *where* the line sits — a rotation preserves canonical normalization, and the absolute decay constant of a_K is the open normalization of O14, parametrized in §26.3 by the benchmark cap loading Ξ_{cap} — and it does not populate the light spectrum with a partner: below the lift there is one field from this sector.

Physical picture. One clock hand on one dial. The bulk Green–Schwarz normalization (a_{GS}) and the canonical light eigenstate (a_K) the apparatus couples to are one field in two normalizations; the benchmark parametrization relates them by the integer loading Ξ_{cap} , and the physical decay constant, mass, and frequency are those of a_K , the single propagating field — the bulk normalization is an intermediate of the normalization chain, not a second particle.

9.4 The registered exposures of the parents

The count claims more than the compact-angle sector alone can give, and the parents register the modes against which it must be checked. Five registered exposures — nine modes beyond the tensor remnant — are massless at tree level on the admissible class and carry open masses in GG-A3-B16 (its mode analysis, O13/O14, and clause C9) and GG-P1-KK (O-KK-67). Each is listed here with the reason it is not expected to be a light QCD-coupled pseudoscalar, at the grade the parent supports.

(1) *The Cartan holonomies of SU(5)* (four angles) and (2)–(3) *the Abelian circle holonomies* $A_{B,\theta}$, $A_{Q_i,\theta}$: flat, ultraviolet-weighted zero modes of A_θ . The SU(5) directions are expected to be Hosotani-lifted at the circle’s curvature scale — GG-A9-SM computes the coset zero mode’s radial profile and the mass its holonomy would register, near 2×10^{15} GeV (PS-SM-22), but the Hosotani potential that lifts the Cartan directions is not computed in any parent (grade: expectation). The Abelian holonomies have no bulk $G\tilde{G}$ coupling — a six-dimensional theory carries no bulk $A \wedge F \wedge F$ Chern–Simons term for them — but boundary-localized five-dimensional terms at the interval ends, which inflow to endpoint-localized chiral matter can require, are not excluded by any parent; they

are lifted by loops of the charged fields at the circle scale (grade: expectation; mass registered open in GG-A3-B16).

(4) *The two interval Wilson lines* $\int_0^K A_{B,r} dr$ and $\int_0^K A_{Q_i,r} dr$, present under the boundary-condition termination (GG-A3-B16 clause C9): real (non-complex) four-dimensional scalars whose parity is open, massless at tree level and lifted only at loop level by the baryonic current — the interval analogue of a Hosotani mode, “the one new light state the choice implies,” with mass and couplings registered open (GG-P1-KK, O-KK-67; GG-A3-B16 C9). A one-loop estimate places the mass parametrically at $g_B M_{\text{lock}}/4\pi \sim 10^2$ GeV before order-one factors, so the count’s conditional is live for these two modes in particular. Their parity under the interval’s reflection is not fixed by any parent; they carry no anomaly-descended coupling and so no tree-level $G\tilde{G}$ term (grade: open in both parents).

(5) *The decoupled vector zero mode* $B_{\mu\theta}$ of norm $2\pi K$: a vector, not a pseudoscalar; it couples to X_4 only through Kaluza–Klein modes, and whether it mixes kinetically with hypercharge at loop level is the dark-photon question Part IV must carry as open (grade: registered by GG-A3-B16 as an open exposure).

The honest statement of the count is therefore: the tensor sector leaves exactly one light pseudoscalar (the mode analysis of GG-A3-B16, its § *Self-duality of the Green–Schwarz tensor*, with open items O13/O14); the six holonomies are expected to lift at the circle’s curvature scale, of which GG-A9-SM’s coset zero mode near 2×10^{15} GeV is the one computed marker, none of the six computed; the vector is not a pseudoscalar, and the two interval lines’ parity is open. “Exactly one light QCD-coupled pseudoscalar” is the conclusion *conditional* on the eight registered scalar masses (five groups, nine modes counting the vector) lying above the electroweak scale and the vector staying decoupled, which no parent has computed. T-Kaxion-5 is graded accordingly — derived-conditional, not theorem — and the count test of Part VI is what decides it: a confirmed second light QCD-coupled line inside the window would not merely kill the count, it would identify which registered mode was light.

9.5 Why the Kaxion is light: protection, quantified

The shift symmetries established above forbid any perturbative potential. What finally breaks them is topology: instanton-like field configurations, whose contributions carry the exponential $e^{-S_{\text{inst}}}$ of their Euclidean action. Such effects generate a periodic potential of the generic form $V = \chi [1 - \cos \theta]$, where χ — the *topological susceptibility* of whatever gauge sector the angle talks to — absorbs the exponential. The mass that results is not a geometric scale but the ratio of a topological scale to a geometric one:

$$m_K = \frac{\sqrt{\chi}}{f_K}. \quad (9.5)$$

For the Kaxion the relevant susceptibility is QCD’s (the strong-CP alignment of Part V), measured on the lattice as $\chi_{\text{QCD}}^{1/4} = 75.5$ MeV (Borsanyi et al., 2016). The arithmetic preview — done honestly in Part V, but worth seeing now —

$$m_K = \frac{(75.5 \text{ MeV})^2}{7.1 \times 10^{11} \text{ GeV}} = \frac{5.70 \times 10^{15} \text{ eV}^2}{7.1 \times 10^{20} \text{ eV}} \simeq 8.0 \times 10^{-6} \text{ eV}, \quad (9.6)$$

shows how a theory whose native scales are 3.85 TeV and 2.3×10^{16} GeV comes, without embarrassment, to predict a dark-matter particle at micro-electron-volts. Dividing by the two-figure 7.1×10^{11} GeV returns $8.0 \mu\text{eV}$; the benchmark quoted in Part V, $m_K = 8.1 \mu\text{eV}$ (26.1), carries the unrounded $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$ of (26.3), and the two agree to within the rounding.

9.6 The zero-mode among three towers

A six-dimensional theory hands the boundary three families of non-spatial states, and placing the Kaxion correctly among them keeps this paper sharply separated from its sister (Arisaka, 2026, P1-KK).

First, the *angular tower*: integer Fourier modes $e^{in\theta}$, $n \neq 0$, of bulk fields on S_θ^1 . Because the circle is rigid with unit metric, these quanta sit at the ultraviolet end of the scale interval, with masses of order $|n|$ times the circle’s curvature unit — which GG-A9-SM places at 10^{17} – 10^{18} GeV, above $E_{\text{GUT}} = 2.3 \times 10^{16}$ GeV — thirty-one to thirty-two orders of magnitude above the Kaxion, frozen out since the dawn of time. Second, the *radial lock cluster*: the discrete resonances of the scale interval pinned at its infrared endpoint $r = K$, with masses at $M_{\text{lock}} \approx 3.85$ TeV (6.3) — nearly eighteen orders above the Kaxion, and the collider business of P1-KK. Third, the *compact-angle zero-modes*: the $n = 0$ angle variables themselves — the tensor remnant and the holonomies — carrying no tower quantum number at all.

The hierarchy among the three families is not numerical accident but structural, and the structure is worth one paragraph of intuition. A tower state owes its mass to a *gradient*: it costs energy to wiggle a field around the circle ($n \neq 0$) or along the interval (the lock cluster). A zero-mode is not a wiggle; it is the position of the dial itself, and rotating the whole dial costs nothing to any finite order — that is Theorem 8.2(ii) in mechanical language. Only topology knows the absolute dial position, and topology pays in exponentially suppressed coin (9.5). Hence three families at 10^{16} GeV, 10^3 GeV, and 10^{-14} GeV, all from one geometry, none tuned. Table 4 sets the three families side by side — support, mass scale, observable window, and the cohort paper that owns each — making visible that the three energies are one geometry read at three depths, not three assumptions.

Table 4: The three non-spatial mode families of GG-6 and their division of labor across the phenomenology branches. The three towers exhaust the non-spatial modes the geometry permits, and each has a different owner: the angular tower is frozen out at unification and inaccessible, the radial lock cluster is the collider signature owned by GG-P1-KK, and the compact-angle zero mode is this paper’s subject. Reading the middle column down is the quickest way to see why dark matter had to come from the third and not from either of the others.

Family	Mass scale	Paper	Physical character
Angular tower ($e^{in\theta}$, $n \neq 0$)	$\sim n E_{\text{GUT}} \sim 10^{16} \text{ GeV}$	— (inaccessible)	Quantized wiggles around the phase circle; frozen out at unification.
Radial lock cluster	$M_{\text{lock}} \approx 3.85 \text{ TeV}$	P1-KK	Resonances of the scale interval at $r = K$; the collider signature.
Compact-angle zero-mode	μeV corridor	this paper	The tensor’s dial position itself; protected by shift symmetry, massed only by topology, Eq. (9.5).

9.7 The KK in the name

The name *Kaxion* can now be parsed exactly. The field is axion-*like* in its low-energy kinematics — periodic pseudoscalar, shift-symmetric, anomaly-coupled — and it discharges the strong-CP duty that the constitution assigns (GG-A1-Constitution P7; Part V). But it is Kaluza-like in its ancestry: it remembers a compact dimension, a throat depth, and an integer ledger, and that memory is precisely what equips it with the fenced window, the fixed location on the DFSZ line, and the single-line count that no Peccei–Quinn construction shares (Peccei & Quinn, 1977; Weinberg, 1978; Wilczek, 1978). The “K” is the throat-depth integer $K = 35$ wearing its own name.

10 Born, selected, stabilized: the GRIP cascade on the compact-angle stratum

Theorem 8.2 establishes that the Kaxion exists in the Lagrangian. Dark matter, however, is not a Lagrangian entry; it is an *occupation* — a particular configuration realized by the actual universe. Section 6.5 built the general machine; here we run it on the Kaxion’s home stratum.

10.1 The stratum and its landscape

The compact-angle stratum Σ_θ is spanned by the tensor remnant and the holonomy directions, $(\theta_{\text{GS}}, \theta_W)$, with the free energy $F_{\text{GFF}}|_{\Sigma_\theta}$ inherited from (6.24): the energetic part carries the (topology-generated) periodic potential and the gradient costs; the entropic part carries the open-system temperature of the boundary. GG-A6-GRIP proves that the critical points of this restricted landscape are not a continuum but sit at the *discrete OGN-5 phases* — ledger-rational angles such as $2\pi M/L = 2\pi/5$ for the holonomy directions, a discreteness inherited from the integer quantization of the Wilson loop (6.4) — while along the light tensor direction the QCD potential of Part V

has its unique minimum at $\bar{\theta}_{\text{eff}} = 0$ (Vafa–Witten), which is where the attractor selection reads the CP-conserving point. MSMF smiles on this: even the *vacua* of the compact sector carry no continuous freedom.

10.2 The cascade, run explicitly

On the stratum, the cascade (6.26) reads off the cosine landscape directly. The DGI is the symmetric configuration — the angles undisplaced, the full compact-angle symmetry intact; on the potential it is the degenerate point whose stability the susceptibility decides. The GSSB chain passes through the saddle at the hilltop ($\theta = \pi$, Morse index 1 on the circle direction): the configuration of maximal banked energy and the watershed between basins. The GA is the locked phase at the basin floor — Morse index 0, Hessian positive definite, with curvature

$$\left. \frac{\partial^2 V}{\partial a_K^2} \right|_{\text{GA}} = \frac{\chi}{f_K^2} = m_K^2 > 0, \quad (10.1)$$

the explicit Hessian computation, with the lifted direction’s admixture bounded, being Appendix C.2. Equation (10.1) is worth a pause: *the dark-matter mass is the stability of the attractor* — the same number read twice, once by a cosmologist as m_K and once by the runtime as the confirmation that the cascade has genuinely terminated.

10.3 The landscape, quantified

Numbers anchor intuition, so evaluate the stratum landscape the cascade descends. The barrier between adjacent vacua of the cosine is $\Delta V = 2\chi_{\text{QCD}} = 2 \times (75.5 \text{ MeV})^4 = 6.5 \times 10^{-5} \text{ GeV}^4$ — converted to mass density, $1.5 \times 10^{13} \text{ g/cm}^3$, about six percent of nuclear density: the rim of the Kaxion’s funnel is, by everyday standards, an extraordinarily stiff structure. The curvature at the floor is $m_K^2 = \chi_{\text{QCD}}/f_K^2$ (Eq. 10.1); the dimensionless steepness $m_K/f_K = \chi^{1/2}/f_K^2 \sim 10^{-26}$ tells the other half of the story — a funnel of nuclear-stiff rim and almost inconceivably gentle slope, which is exactly the combination that stores a cosmologically long-lived oscillation. Today’s displacement energy, $\rho_{\text{loc}} \simeq 3.5 \times 10^{-6} \text{ eV}^4$, stands to the barrier as 5×10^{-38} : the dial that once leaned nearly half the way up the rim now rattles thirty-eight orders of magnitude below it. Every stability statement of this paper — the positivity of (10.1), the lifetime of Part V, the coldness of Part III — is a restatement of these three numbers.

10.4 The Kaxion as Ground Geometric Attractor

The discrete OGN-5 phases of the stratum are, each of them, index-zero minima — each a legitimate GA. The statement that elevates the Kaxion’s locked phase above the others is the ground-state refinement of Section 6.5: by T-GGA-as-Ground-State (GG-A11-Origin), exactly one of the stratum’s attractors passes the Qi-4 admissibility filter and is the global minimum of $F_{\text{GFF}}|_{\Sigma_\theta}$,

$$\text{GGA}_\theta = \arg \min_{(\theta_{\text{GS}}, \theta_W) \in \Sigma_\theta} F_{\text{GFF}}, \quad (10.2)$$

and the Kaxion is the light excitation about *that* configuration — the GGA of the compact-angle stratum, not merely one of its GAs. The upgrade is not vocabulary; it buys two physical theorems the weaker statement cannot.

Universality: every Kaxion is the same Kaxion. Chemistry’s deepest convenience is that ground states are *types*: every water molecule is identical because each one’s electrons sit in the same lowest configuration. GG-A8-QM (Arisaka, 2026, A8-QM) runs exactly this logic on the electron itself — T-GGA-as-Ground-State specialized to the charged-lepton stratum proves that every cascade realization, anywhere in spacetime, selects the *same* unique minimum $(m_e, -e, \frac{1}{2}\hbar)$, which is why every electron is rigorously interchangeable. The identical argument on Σ_θ delivers the Kaxion’s version: every quantum of the condensate, at every point of the halo, is an excitation of one and the same GGA. This is the microscopic warrant for two things this paper elsewhere assumes: the condensate’s universality (a population of excitations of a single type, Part III; its phase coherence is finite, set by the halo dispersion, Part VI), and the universality of its constants — the window fence and the portal weights are properties of *the* Kaxion, exactly as the Rydberg is a property of *the* hydrogen atom, not averages over a population that could drift or disperse.

Stability: nowhere deeper to go. A GA is stable against small perturbations; only a GGA is stable against *tunneling*, because there is no deeper basin on the stratum to tunnel to. This is the same statement chemistry makes about a filled shell, and the same one GG-A9-SM makes about the proton (no lower-lying baryonic configuration once $U(1)_B$ is gauged). For the Kaxion it closes the last conceivable loss channel, and the regime vocabulary of Section 6.5 lets the closure be stated categorically rather than numerically: the Kaxion’s exit-channel set is *structurally empty* — the gauge-protected shift symmetry deletes any inter-basin channel, so within its stratum the field has literally nowhere to decay, not merely nowhere convenient. (What this paper argued structurally is now a named result of the series: T-GGA-No-Exit, the corollary of the capstone’s descent theorem (Arisaka, 2026, A11-Origin), states that GRIP-In halts exactly where the exit-channel set is empty — and the Kaxion is its cleanest physical instance, the empty set being enforced by a gauge redundancy rather than by an energy accident.) Its only escape is the portal leak *off* the stratum — the two-photon channel computed in Part V, thirty-two orders of magnitude slower than the age of the universe for the laboratory eigenstate — and that leak is portal radiation, not a basin exit: it does not move the condensate to another attractor, and its rate could be multiplied a billionfold without changing the categorical statement. The stability ledger of the dark matter is thus three lines, all ground-state statements: stable against perturbation (Hessian, Eq. 10.1), stable against tunneling (GGA with empty exit set, this section), stable against decay (lifetime, Part V).

A cohort note. The GGA concept is defined and proved in the capstone (GG-A11-Origin) and carried in GG-A6-GRIP’s runtime vocabulary, where it natively belongs (Definition def:GGA), together with the two-regime vocabulary GRIP-In/GRIP-Out (Definitions def:GRIPin, def:GRIPout) used throughout this paper; the cascade reads uniformly $DGI \rightarrow GSSB \rightarrow GGA$ across the series, with GFF-ST7 recognized as the GGA statement’s runtime content.

10.5 The Kaxion on the two fibers: matter of the circle, occupancy of the interval

The cohort's recent campaign (Arisaka, 2026, W1-Symmetry; Arisaka, 2026, P5-CPT; Arisaka, 2026, P3-Electron) compresses the bulk's non-spatial pair into one exact statement, and this subsection reads everything Part II has built on it: *holonomy requires closed loops, and an interval has none*. The compact circle S_θ^1 is nothing but closed loops — everything on it is holonomy, quantized, gauge-protected, its transport composing to a group: the fiber that *remembers*. The scale interval I_r can hold no holonomy at all — its transport composes only forward, a semigroup: the fiber that *forgets*. The cohort has already drawn two portraits on this anatomy. The electron (Arisaka, 2026, P3-Electron) is matter *on* the circle: a holonomy class — charge, spin, representation — wearing a mass that the descent assigns. The proton (GG-P6-Proton; Arisaka, 2026, P6-Proton) is matter *guarded by* the circle: two holonomy ledgers whose joint integrality welds the state shut. The Kaxion completes the triptych, and it is different in kind from both: it is matter *of* the throat's tensor sector — the boundary dual of the tensor the anomaly ledger forces, promoted to a boundary field by OGGEF, occupied by the runtime, and left ringing by the descent. In one sentence: *the electron is a holonomy class wearing a mass; the proton is two holonomy ledgers wearing a weld; the Kaxion is the anomaly ledger's own phase, wearing an occupation*. Dark matter, in GG-Theory, is the sixth dimension relaxing into its ground attractor after the fifth has carried the universe below the barrier. Figure 6 draws the claim; the paragraphs below audit it, structure by structure, and every clause carries the grade of the parent theorem it reorganizes.

Circle-side: everything the Kaxion is. Run down Theorem 8.2 and mark the fiber doing the work. *Existence and ancestry* (item 1): the Kaxion is a creature of the closed fiber — the tensor remnant θ_{GS} is the circle's *bookkeeping made field*, the boundary dual of the tensor that exists because the anomaly ledger must balance (GG-A3-BI6); and its would-be sibling, the holonomy echo θ_W — the Wilson line of S_θ^1 promoted to a field — is the fiber's other creature, lifted at the circle's curvature scale and absent from the light spectrum. *Protection* (item 2) is the circle's memory wearing a Lagrangian: a periodic direction admits no polynomial potential, so the shift symmetry that keeps the Kaxion light to all orders in perturbation theory is not a symmetry imposed but a topology remembered — *the Kaxion is light because the circle cannot forget*. *Locked portals* (item 3, and Theorem 24.1): the coefficient vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$ and the unit color anomaly — all are the BI-6 ledger transported unchanged, which is to say circle-side accounting; the Kaxion couples *only through topology because it is made of topology*, and that single clause is the entire origin of its darkness (§10.7). The *discrete vacua* of the stratum landscape — ledger-rational angles such as $2\pi/5$ — are the Wilson loop's integer quantization reaching down into the vacuum manifold itself: even the floor the cascade selects is a holonomy datum. The *window* of Part V is a discrete ladder of cap loadings around a calibrated benchmark — the ladder from the ledger by ansatz, the placement by calibration — and the *count* of Part VI — one line, alone — is scale-free and immune to every residual uncertainty in this paper, conditional on the registered exposures of §9.4. The *strong-CP alignment* of Part V (§24) is the same attractor selection reading the CP-conserving point off the groove. And the *universality* established below — every Kaxion the same Kaxion — is

the condensate analog of the electron’s sameness: one frozen angle is one holonomy configuration for the entire observable universe, and holonomy configurations, like classes, do not age.

Interval-side: everything the Kaxion *has*. The other column of the ledger is occupancy, and every entry is the open fiber’s. The bulk normalization scale $f_{\text{bulk}} \simeq 3.04 \times 10^{13}$ GeV is a scale of the throat — descended, an address. The *freeze* is the interval’s grip: the Hubble friction that holds the displaced angle in place for the first three hundred nanoseconds is the expansion itself — the semigroup — damping the circle’s coordinate (§12). The *release* at $T_{\text{osc}} \simeq 1$ GeV is the moment the circle’s clock outruns the interval’s (quantified in §12). The *mass* is delivered by the descent: the susceptibility $\chi(T)$ that stiffens the potential rises as the universe cools through the QCD scale, so the number $m_K = \sqrt{\chi_{\text{QCD}}}/f_K$ is the boundary’s gift to the circle, handed over at a temperature the one-way history fixes. The *relic abundance* is the memory of the crossing — an adiabatic invariant integrated down the descent, $\Omega_K \propto f_K^{1.165} \theta_i^2$ (§15.3) — and the *coldness*, the $\theta_0 \simeq 4.6 \times 10^{-19}$ radian amplitude surviving today, is what the fiber that forgets does to an amplitude across $\sim 10^{27}$ oscillations while never touching the frequency. The division of labor could not be cleaner: *the circle supplies what the Kaxion is; the interval supplies what the Kaxion has*. Identity is topology; occupancy is history.

The normalization: the one place the fibers multiply. Exactly one structure in this paper belongs to both fibers at once, and it is where the open numbers live. The tensor’s boundary norm and the Wilson line’s (§9.3, §26.3) differ by e^{2K} — the *interval’s depth appearing inside the circle sector’s normalization*, $e^{70} \simeq 2.5 \times 10^{30}$ — the infrared-enhanced norm against the order-one one, and the ratio of the bulk to the canonical normalization is parametrized, by ansatz, as the integer loading of (26.3) (the absolute value being the calibration of Part V): $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$. Note the *shape* of that formula: a calibrated scale divided by a circle integer. It is the same grammar the codebook paper reads across all twenty-eight Standard-Model parameters (Arisaka, 2026, A9-SM) — a holonomy-class label correcting a descent address — operating here inside the Kaxion’s own normalization. The anatomy thereby does something for the honesty ledger that no amount of computation could: it *localizes* the residual uncertainty. The circle-side half of the formula — that the loading is an integer, that the window of Part V is discrete — is an ansatz (§26.3); the *which* integer is the compact-angle mass-uniqueness conjecture; and the seam between the two is where the open items of this paper live: the calibration of the absolute scale (O14), the ladder ansatz, and, on the cosmology side, Conjecture C-Kaxion-d (§13.5). The first two sit at the one point where the two fibers mesh — which is exactly where a two-fiber theory should keep its unknowns.

Why the Kaxion is the cleanest GGA exhibit in the cohort. Two features distinguish this stratum’s attractor from the electron’s and the proton’s, and both deserve to be said plainly. *First, the Hessian is the observable*. For the electron, the GGA’s stability manifests as an absence — no decay, no drift; for the proton, as a lifetime bound. For the Kaxion, the attractor’s stiffness *is the laboratory number*: Eq. (10.1) says m_K^2 is the curvature of the free-energy floor, so the 1.95 GHz line of Part V is that curvature read as a frequency. A haloscope detection would be the first direct *measurement of a GGA Hessian* — the runtime’s termination certificate, heard as a tone.

Second, the cascade is still finishing. The electron's and the proton's cascades terminated within the first second of cosmic history; their attractors are ancient. The Kaxion's relaxation is happening *now*: the coherent oscillation about the floor is the final GRIP-In of the observable universe, caught mid-ring. Dark matter is the one place in nature where the GRIP cascade is observable in real time rather than forensically — and the dark-matter search is, in this precise sense, the cohort's only live broadcast from the runtime. This subsection, like the campaign it joins, is an organization of proved structures, not a new theorem; every clause above carries its parent's grade.




Figure 6: The Kaxion on the two fibers. (a) *The circle: identity.* The compact fiber S_θ^1 carries the two compact directions (θ_{GS} light, θ_W lifted; the throat norm $\sim e^{2K}$ sets the normalization chain); its vacua sit at discrete ledger-rational phases (dots); the GRIP cascade runs DGI (center) $\rightarrow$ GSSB (the $\theta = \pi$ saddle) $\rightarrow$ GGA (the floor), and the curvature at the floor *is* the dark-matter mass, Eq. (10.1). The portals and the address on this panel are exact and unrunnable; the window is a ladder by ansatz around one calibrated pitch, one frozen angle, one line. (b) *The interval: occupancy.* The one-way descent holds the angle through primary inflation and deposits it at the lock crossing ($R \equiv m_K/3H \approx 0$), grips it through the frozen era ($3H \gg m_K$), releases it exactly once when the circle's clock outruns the interval's ($R = 1$ at $T_{\text{osc}} \simeq 1 \text{ GeV}$), and has since spent the amplitude across $\sim 10^{27}$ oscillations while never touching the frequency. The equation of motion (12.2) is the anatomy in one line: a_K is the circle's coordinate, $3H\dot{a}_K$ the interval's friction, $m_K^2 a_K$ the circle's stiffness delivered by the descent. *The circle supplies what the Kaxion is; the interval supplies what the Kaxion has.*

10.6 Birth, selection, stabilization

Three runtime statements now carry forward into Part III, and they are the entire interface between this Part and cosmology.

Birth. The Kaxion's cosmological history begins when the global GRIP cascade crosses the lock saddle of its two-basin descent, releasing the compact-angle sector displaced from its eventual attractor — the angle having been held through primary inflation (the conditional C-Kaxion-d of Part III). The field is released displaced *by the offset at which it was held* — the lock crossing releases the angle; it does not set the displacement — rather than by accident.

Selection. The cascade releases the compact-angle sector displaced from its attractor by the offset at which the inflation-era pinning held it — a structured $\mathcal{O}(1)$ angle, not a tuned one, and an inflation-era datum rather than a lock-crossing one. The runtime does not predict its precise

value; Part III reads it from the observed abundance at the calibrated benchmark ($\theta_i \simeq 1.1$ rad), the generic misalignment expectation.

Stabilization. The endpoint of the relaxation is the GA: index zero, Hessian (10.1) positive. The condensate oscillating about it is therefore stable, coherent, and cold — the three adjectives that Part III converts, respectively, into a lifetime, a wave description, and a structure-formation verdict.

Physical picture. Release a marble high on a terraced funnel. The terraces are the GSSB saddles; the funnel floor is the GA. The marble’s destination was never in doubt (GFF-ST7); the *history* of its descent — where it was released, which terraces it crossed, how it rings as it settles — is physics. The Kaxion is the ringing: the compact-angle sector, released displaced at the lock crossing, still settling today, at a frequency we compute in Part V and propose to detect in Parts VII–VIII.

10.7 Why the Big Bang cannot fail to make Kaxions

The electron paper closes its identity question with a theorem; the proton paper closes its stability question with a guard. The parallel question for this paper is *inevitability*: is the Kaxion condensate an accident of initial conditions, or a forced outcome of the Big Bang, on the same footing as the electrons and protons that the descent cannot fail to deposit? On the two-fiber anatomy the answer assembles in four steps, each carrying its grade.

Step 1 (the circle is forced; theorem). By GG-A2-GG6 T1–T2 any admissible bulk carries the rigid circle S_θ^1 with quantized holonomy, and by Theorem 8.2 its light periodic descendant on the boundary — the tensor remnant — is then not optional, its holonomy sibling lifted. A universe built on this constitution comes with the Kaxion’s parents installed.

Step 2 (the descent is forced and one-way; theorem plus the derived history). A hot start must descend; r decreases; the transport along I_r is a semigroup and cannot re-ascend (GG-A5-GDOS; Arisaka, 2026, A10-Cosmos). There is exactly one pass through every scale.

Step 3 (a descent from above the barrier cannot leave the angle on its floor; named conditional). The lock-crossing saddle crossing deposits the compact-angle sector displaced — the GSSB event of the *same* cascade that produces the inflationary exit and the inventory, with the inflation-era pinning of Conjecture C-Kaxion-d holding one angle coherently across the whole future observable universe. This is the one conditional step, and it is named, owned, and falsifiable: a Kaxion minicluster population, a line at the randomized-angle mass, or a detection of axion-sector isocurvature kills the pinning outright (§13.5). The step is the one conjecture of the cosmology side — not a new hypothesis in kind, since it is a segment of the event chain already required upstream, but one neither paper has computed; “one marble descending one funnel.”

Step 4 (a displaced circle coordinate must relax as cold, dark, coherent matter; theorem given 1–3). Displaced, the coordinate is gripped while $3H \gg m_K$ and released exactly once when the descent hands it its mass (§12); thereafter it *must* oscillate about the GGA and redshift as pressureless matter. It is *dark* because its portals are the anomaly ledger’s, suppressed by f_K — darkness is not an assumption but the absence of any non-topological coupling for a pure holonomy mode to have; a field made of topology can talk to matter only through topology. It is *cold* because it is born as a zero-momentum condensate — one pinned angle, no thermal history. It is *coherent* — one angle, one phase, one line at 1.95 GHz.

Table 5 places the conclusion beside its two siblings. Given the constitution, the Big Bang cannot end without electrons (the ladder must terminate), without protons (the weld must close and the guard must hold), and without a Kaxion condensate (the displaced angle must relax). The three are the three ways a one-way descent deposits permanence on a compact fiber: a labeled ground state, a guarded weld, and the fiber’s own coordinate ringing at its floor. *Dark matter is as inevitable as matter, and for the same reason: the universe has one fiber that must forget and one that cannot.* The conditionality is honest and localized — Step 3 rides on the inflation-era pinning (Conjecture C-Kaxion-d, killable by isocurvature or miniclusters), and the quantitative address rides on the cap conjecture (§26.3; the discrete window of Part V) and on the calibration of the bulk normalization; the count rides on the parents’ registered exposures, and the closure on GG-A10-Cosmos’s target (C-Kaxion-c) and on misalignment completeness — and nothing else in the chain is open.

Table 5: The triptych: three Ground Geometric Attractors, one anatomy. The electron (Arisaka, 2026, P3-Electron), the proton (GG-P6-Proton), and the Kaxion (this paper) as the three ways the one-way descent deposits permanence on the compact fiber. Read across, the three differ in stratum and in scale, and not in mechanism.

	Electron	Proton	Kaxion
Relation to the circle	matter <i>on</i> it (labeled)	matter <i>guarded by</i> it (welded)	matter <i>of</i> it (the coordinate itself)
Stratum	charge stratum	baryon stratum	the compact-angle stratum Σ_θ
Identity mechanism	one ground basin: classes do not age	$Q + B \in \mathbb{Z}$ double guard	one frozen angle: patch-uniform pinning
Mass origin (interval)	terminus address of the lepton descent	binding at the QCD floor	$\chi(T)$ switched on by the descent; $m_K = \sqrt{\chi_{\text{QCD}}}/f_K$
Labels (circle)	charge winding, spin, representation	B, Q ledgers	portals $(1, -1, -\frac{2}{5}, -\frac{2}{5})$; single line; ledger-rational vacuum
GSSB event	family cascade	confinement weld	lock-crossing release
GGA certificate	no exit from the ground basin	no-exit via the double guard	Hessian $= m_K^2 > 0$: the 1.95 GHz line
Why inevitable	the ladder must terminate	the weld must close	the displaced angle must relax

10.8 What Part II has established

A reader who began this Part with no exposure to GG-Theory now holds the following, with no reference to any astronomical observation: a rigid six-dimensional stage with a non-spatial phase circle and a five-integer ledger of documented accuracy (Section 6); a forced tensor bookkeeper with ledger-locked entries (Section 6.3); a constitutional projection to the boundary preserving four currents (Section 6.4); a runtime with a unique free-energy minimum (Section 6.5); and, assembled from these, a unique light periodic pseudoscalar of the compact-angle sector (Theorem 8.2) — the

tensor's remnant, its holonomy sibling lifted (Section 9), exponentially below the geometric scales for a structural reason (Section 9.6), born, selected, and stabilized by the same cascade that selects the Standard-Model vacuum (this section). What remains is to let the universe act on it. That is Part III.

11 Part II in plain English: why the geometry has no choice but to produce a light pseudoscalar

This Part is the tour for a reader who has not read the parent papers, and its job is to make one step feel inevitable rather than clever.

The framework puts a four-dimensional boundary on a six-dimensional bulk, and the two extra directions are not alike. One is a compact circle. A field can wind around a circle, and the phase it accumulates going around is a quantity with no local definition — there is no point at which you can say what it is, only a statement about the loop as a whole. That non-locality is the whole argument. A quantity with no local definition has no local counterterm, so nothing in the theory can give it a large mass. It is therefore light, and it is light for a reason rather than by arrangement.

The rest follows from bookkeeping the parent papers already did. The circle is there because the anomaly identity requires it. The integers that set the gear ratios are fixed. The coupling to gluons is the only portal the charge assignment allows. None of this was arranged with dark matter in view, which is the point worth holding onto: the pseudoscalar is a consequence the framework could not avoid, not an addition made to explain something.

A reader who wants to attack the paper should attack it here. If the circle is optional, or if the field could have acquired a mass by some route the argument has missed, then everything downstream is a calculation about nothing.

Part III

Cosmology of the Kaxion

One event chain: inflation exit, Kaxion birth, dark energy

12 The cosmic event chain

12.1 One ledger, three epochs

Part II ended with a field; cosmology begins with a history. The cosmology paper GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) insists — and this Part adopts as its organizing principle — that the Kaxion’s history is not a separate hypothesis appended to inflation and dark energy, but one segment of a single GRIP event chain read off one ledger:

$$\text{BI-6} \implies \text{TGG} \implies \text{CCI} \implies \text{GFF} \implies \text{GRIP} \implies \text{parameter block.} \quad (12.1)$$

The reading is as follows. BI-6 fixes which fields exist (Part II). The total geometric–gauge object TGG places the inflationary exit, the compact-angle displacement, and the late vacuum residue inside *one* conserved cosmic packet, so their budgets are not independent. The Covariant Conservation Identity CCI makes the compact-angle mode’s four-dimensional response lawful on the boundary seam. The GFF Morse landscape ranks the available branches, and GRIP executes the descent: the lock saddle of the two-basin descent is crossed, the compact-angle sector — held through primary inflation under C-Kaxion-d — is released displaced from its attractor, and the relaxation of that displacement is — from that moment to this — the dark matter. Inflation, dark matter, and dark energy are one marble descending one funnel, observed at three epochs.

12.2 The equation of motion and its two regimes

All of the cosmology of this Part flows from a single equation. Varying the canonical action of the light eigenstate a_K (Part II, Eq. (7.2)) on a flat Friedmann–Robertson–Walker background $ds^2 = dt^2 - R^2(t) d\vec{x}^2$ gives, for the homogeneous mode,

$$\ddot{a}_K + 3H \dot{a}_K + V'(a_K) = 0, \quad H = \frac{\dot{R}}{R}, \quad (12.2)$$

and for small displacements about the attractor, $V'(a_K) \simeq m_K^2(T) a_K$. Equation (12.2) is a pendulum with friction, and the friction is the expansion of the universe itself: the $3H\dot{a}_K$ term removes energy from the homogeneous mode exactly as a damping rod removes energy from a

laboratory oscillator. During radiation domination the Hubble rate is

$$H(T) = 1.66 \sqrt{g_*(T)} \frac{T^2}{M_{\text{Pl}}}, \quad M_{\text{Pl}} = 1.22 \times 10^{19} \text{ GeV}, \quad (12.3)$$

with $g_*(T)$ the effective number of relativistic degrees of freedom ($g_* \simeq 73$ at $T \simeq 1$ GeV, falling to 3.36 today, with the entropy counterpart $g_{*s,0} = 3.91$).

The pendulum has exactly two regimes, and the boundary between them is the single most important moment in the Kaxion's life.

Regime 1: frozen ($3H \gg m_K$). When the friction dominates, the field cannot fall. Quantitatively, the terminal-velocity solution of (12.2) is $\dot{a}_K \simeq -V'/3H$, so the fractional displacement lost per Hubble time is

$$\left. \frac{\Delta a_K}{a_K} \right|_{1/H} \sim \left(\frac{m_K}{3H} \right)^2 \ll 1. \quad (12.4)$$

The field is a pendulum released in honey: it leans where it was born and stays there, its potential energy banked and waiting. This is why the initial angle released at the lock crossing survives unmodified across the four orders of magnitude in temperature between release and onset, the integrated slip at release being $\sim 10^{-39}$ — the universe holds the Kaxion still until it is ready to let go.

Regime 2: oscillating ($3H \lesssim m_K$). When the Hubble rate falls to

$$3H(T_{\text{osc}}) \simeq m_K(T_{\text{osc}}), \quad (12.5)$$

the friction can no longer hold the pendulum, and the field begins coherent oscillation about the attractor. From this moment the energy density of the oscillating condensate redshifts as pressureless matter (proved in Section 16), and the Kaxion *is* cold dark matter. Everything between Eq. (12.5) and the present is bookkeeping — exact, and done below — but no new physics.

The pendulum, read on the anatomy. Equation (12.2) deserves one more reading, because no other equation in the cohort exhibits the two fibers of §10.5 so compactly. Term by term: a_K is the *circle's coordinate* — the light remnant of the two compact angles; $3H\dot{a}_K$ is the *interval's friction* — H is the descent rate read in boundary time (Arisaka, 2026, C2-Hubble: cooling *is* descent), so this term is the semigroup itself, damping the circle's swing inside a classical field equation; and $m_K^2(T) a_K$ is the *circle's stiffness*, delivered by the boundary as the descent cools through the QCD scale. The two regimes above are then two rates against each other: define the dimensionless ledger $R(T) \equiv m_K(T)/3H(T)$, the circle's frequency against the interval's rate. $R \approx 0$ at the pinning (the mass is off and H is enormous); $R = 1$ at $T_{\text{osc}} \simeq 1$ GeV, where the closed form (25.3) sits and the release occurs; and $R \approx 2 \times 10^{27}$ today — which is, as it must be, the order of the number of oscillations the line has performed since release ($\nu_K t_0 \approx 8 \times 10^{26}$). The two-fiber point is the *monotonicity*: because the descent is one-way, R crosses unity exactly once, in one direction, with no possibility of re-crossing. The release of the Kaxion is not an event that happened to occur; it is

an event that could not fail to occur and could not occur twice — the same gated-interval logic by which the CPT companion’s descent passes the lock threshold once and buys the baryon asymmetry outright (Arisaka, 2026, P5-CPT). That is why the misalignment condensate is *primordial, unique, and unreplenished* — the defining phenomenology of cold dark matter, here a corollary of the fiber anatomy rather than an assumption of the cosmology.

12.3 The timeline at a glance

Anticipating the numbers derived in Sections 13–16, the Kaxion’s biography is short to state and long to honor. Table 6 runs it from the lock epoch to today; every number in it is derived, not asserted, at the section its row cites, and the two decisive rows — the onset at $T \simeq 1.0$ GeV and the present half-attoradian amplitude — anchor everything the laboratory Parts build on:

Table 6: The Kaxion’s cosmic timeline (onset numbers from Section 13). The rows run from the lock crossing to the present, each with its temperature, its clock time, and what happens to the condensate there. The onset row carries the physics: the field begins to oscillate when the Hubble rate falls to the mass, and everything downstream — the abundance, the present amplitude, the swing of the dial — is fixed by where that happens. The initial angle is read back from the observed abundance at the calibrated benchmark.

Epoch	Clock	Event
Lock crossing ($T_{\text{lock}} \simeq 1.6$ TeV, $H = H_2$, $t \simeq 10^{-13}$ s)	—	The descent’s second saddle releases the compact-angle sector displaced by a fixed, patch-independent $\mathcal{O}(1)$ angle, read from the abundance as $\theta_i \simeq 1.1$ rad (Section 13.3); patch independence rides on the inflation-era pinning, Conjecture C-Kaxion-d.
$T \simeq 1.0$ GeV	$t \simeq 0.3 \mu\text{s}$	Onset: $3H = m_K(T)$; the condensate starts to oscillate, with a mass still only $\sim 0.04\%$ of its final value.
$T \simeq 0.15$ GeV	$t \simeq 20 \mu\text{s}$	The QCD susceptibility saturates; the mass reaches its zero-temperature value; the spring is fully stiffened.
$T \simeq 0.80$ eV	$t \simeq 5 \times 10^4$ yr	Matter–radiation equality: the condensate, already $\sim 3 \times 10^{21}$ oscillations old, begins to dominate structure growth.
Today	$t = 13.8$ Gyr	The dial still rings: $\sim 8 \times 10^{26}$ cycles elapsed; present swing amplitude $\theta_0 \simeq 5 \times 10^{-19}$ (Section 14).

Physical picture. It is worth absorbing how asymmetric this biography is. The Kaxion spent less than a microsecond being born and has spent 13.8 billion years ringing down — and the ring-down is so gentle (the only drain is Hubble friction, Section 16) that today’s oscillation is, cycle for cycle, the *same* coherent oscillation that began before the universe was a millisecond old. In regime language (Section 6.5): the birth was the tail of a GRIP-Out; everything since is one uninterrupted GRIP-In, with Hubble friction as the within-basin flow — the dark sector’s counterpart of the weak channel that drains the matter strata (GG-A9-SM). A haloscope is, quite literally, listening to the late stages of an event that occurred at $t \simeq 0.3 \mu\text{s}$.

13 Misalignment production and the initial angle

13.1 When does the pendulum let go? The onset temperature, twice

The onset condition (12.5) can be solved in two instructive approximations, and the contrast between them teaches the key feature of the Kaxion’s potential.

First pass: constant mass. Pretend the mass is temperature-independent at its present value, $m_K = 8.1 \mu\text{eV} = 8.1 \times 10^{-15} \text{ GeV}$ (Part V). Then (12.5) with (12.3) gives

$$T_{\text{osc}}^{\text{const}} = \left(\frac{m_K M_{\text{Pl}}}{3 \times 1.66 \sqrt{g_*}} \right)^{1/2} = \left(\frac{(8.1 \times 10^{-15})(1.22 \times 10^{19})}{42.6} \right)^{1/2} \text{ GeV} \simeq 48 \text{ GeV}, \quad (13.1)$$

using $g_* \simeq 73$. A field this light would, on this reckoning, have begun oscillating when the universe was at a few tens of GeV.

Second pass: the real, temperature-dependent mass. But the Kaxion’s mass is not constant. Its potential is generated by the QCD topological susceptibility (the strong-CP alignment of Part V), and above the QCD scale the susceptibility is instanton-suppressed. The lattice determination (Borsanyi et al., 2016) gives

$$\chi_{\text{QCD}}(T) \simeq \chi_{\text{QCD}}(0) \left(\frac{T_{\text{QCD}}}{T} \right)^{2b}, \quad b \simeq 4.08, \quad T_{\text{QCD}} \simeq 150 \text{ MeV}, \quad (13.2)$$

so the mass turns on as $m_K(T) = \sqrt{\chi_{\text{QCD}}(T)}/f_K = m_K(0) (T_{\text{QCD}}/T)^b$ for $T > T_{\text{QCD}}$. Inserting this into (12.5) yields a polynomial condition,

$$m_K(0) T_{\text{QCD}}^b T_{\text{osc}}^{-b} = \frac{3 \times 1.66 \sqrt{g_*}}{M_{\text{Pl}}} T_{\text{osc}}^2 \implies T_{\text{osc}}^{b+2} = \frac{m_K(0) T_{\text{QCD}}^b M_{\text{Pl}}}{4.98 \sqrt{g_*}}. \quad (13.3)$$

Solving numerically for the benchmark ($b = 4.08$, $g_* = 73$):

$$T_{\text{osc}} \simeq 1.0 \text{ GeV}, \quad \frac{m_K(T_{\text{osc}})}{m_K(0)} = \left(\frac{0.15}{1.0} \right)^{4.08} \simeq 4.4 \times 10^{-4}, \quad t_{\text{osc}} \simeq \frac{1}{2H_{\text{osc}}} \simeq 0.3 \mu\text{s}. \quad (13.4)$$

The two passes differ by a factor of fifty in temperature, and the difference is physics, not bookkeeping: because the susceptibility is still rising when the pendulum lets go, the field starts oscillating with a spring constant only $\sim 0.04\%$ of its final stiffness, and the spring continues to stiffen underneath the oscillation until $T \simeq T_{\text{QCD}}$. This “stiffening spring” is precisely what the adiabatic-invariant bookkeeping of Section 15 is built to handle, and it is the origin of the famous non-integer exponent in the abundance formula.

Physical picture. Imagine the pendulum’s restoring spring being wound stiffer while it swings. Each oscillation is faster than the last; the amplitude shrinks (the adiabatic invariant — the number of quanta — is conserved, so energy per quantum grows with the frequency); and the final energy

bank depends on *when* the swinging started relative to the winding. That interplay, and nothing else, is where the $f_K^{1.165}$ of Eq. (15.7) below comes from.

13.2 The exact solution: Bessel functions and the envelope law

The two regimes of Eq. (12.2) are not merely qualitative limits; for constant mass in the radiation era the equation solves *exactly*, and the solution is worth one subsection because it converts every hand-wave of the misalignment story into a formula. In radiation domination $H = 1/2t$, so

$$\ddot{a}_K + \frac{3}{2t} \dot{a}_K + m_K^2 a_K = 0. \quad (13.5)$$

Substituting $a_K = (m_K t)^{-1/4} u(m_K t)$ reduces (13.5) to Bessel's equation of order $\frac{1}{4}$; the solution regular at $t \rightarrow 0$ is

$$a_K(t) = a_i \Gamma(\frac{5}{4}) \left(\frac{2}{m_K t} \right)^{1/4} J_{1/4}(m_K t) \quad (13.6)$$

normalized with $J_\nu(x) \simeq (x/2)^\nu / \Gamma(\nu + 1)$ so that $a_K \rightarrow a_i$ at early times. The two asymptotics deliver the two regimes with their precise coefficients:

$$a_K(t) \rightarrow \begin{cases} a_i \left[1 - \frac{(m_K t)^2}{5} + \dots \right], & m_K t \ll 1 \quad (\text{frozen, with the leading slip}), \\ 0.86 a_i (m_K t)^{-3/4} \cos(m_K t - \frac{3\pi}{8}), & m_K t \gg 1 \quad (\text{oscillating}), \end{cases} \quad (13.7)$$

the envelope coefficient being $\Gamma(\frac{5}{4}) 2^{1/4} \sqrt{2/\pi} = 0.86$. Three standard statements become theorems of (13.6). *First*, the onset convention: $3H = m_K$ is $m_K t = \frac{3}{2}$, which sits precisely in the knee of the Bessel function — the conventional onset is the mathematical turn-over, not an arbitrary marker. *Second*, the redshift law: the envelope $a_K \propto t^{-3/4} \propto R^{-3/2}$ gives $\rho_K = \frac{1}{2} m_K^2 a_K^2 \propto R^{-3}$ *exactly* — the dust behavior of Section 16 is not an average over assumptions but the asymptotics of a special function. *Third*, the $\mathcal{O}(1)$ honesty: the coefficient 0.86 (and its analog for the temperature-dependent mass, which shifts the Bessel order) is exactly the class of factor that separates back-of-envelope misalignment estimates from calibrated pipelines; the normalization of Part V is obtained by integrating the equation of motion through the onset rather than from any such coefficient, and the 2.5% row that the $\pm 6\%$ budget on the closure angle (Table 12) books for the matching to the precision literature is that residual.

Physical picture. A pendulum released in honey that thins with time. While the honey is thick ($m_K t \ll 1$) the bob creeps by the tiny slip in the first line of (13.7) — note the slip is *quadratic* in $m_K t$, which is why the whole frozen era leaves the angle untouched. When the honey thins past the knee, the bob swings, and thereafter the universe takes its three-quarters power of amplitude per e-fold of time, no more, no less.

One population, one onset. A note before proceeding. The cosmological history is computed once, with the single physical decay constant f_K and mass $m_K = 8.1 \mu\text{eV}$ (Section 14); the mass that sets the oscillation onset is the same one a detector reads at 1.95 GHz.

13.3 The misalignment angle

The single initial condition the misalignment mechanism requires is the displacement angle at birth, θ_i . In generic axion cosmology θ_i is the one continuous dial that turns the abundance into a band; here the band collapses for the opposite reason to the usual one. Because the bulk normalization is calibrated at the benchmark so that a generic angle saturates the inventory (Section 14), the observed abundance returns θ_i rather than scanning it — and what is claimed for the returned value is only that it is generic.

The runtime constrains the angle only qualitatively. The *vacua* of the compact-angle stratum are the discrete OGN-5 phases (GG-A6-GRIP); the *displacement at birth* is the offset between the value at which the angle was held through primary inflation and the QCD minimum it later relaxes toward, released at the lock crossing of GG-A10-Cosmos’s two-basin descent ($T_{\text{lock}} \simeq 1.6$ TeV, $t \simeq 10^{-13}$ s; Section 13.5 states the conditional this rides on, and the release must be the same in every horizon patch for a single angle to be meaningful). What the runtime guarantees is that this offset is a structured $\mathcal{O}(1)$ angle — neither the unbroken point $\theta = 0$ (which would yield no dark matter) nor the unstable hilltop $\theta = \pi$ — not that it takes any particular value. Its value is then read off the abundance closure of Section 15.3,

$$\theta_i \simeq 1.1 \text{ rad}, \quad (13.8)$$

the generic misalignment angle, no more tuned than the $\mathcal{O}(1)$ angle every standard axion-dark-matter scenario assumes.

Two remarks keep this honest. First, (13.8) presumes the *pre-inflationary* (single-patch) history, in which one value of θ_i is homogenized across the observable universe; the post-inflationary alternative would replace θ_i^2 by the circle average $\langle \theta^2 \rangle = \pi^2/3$ and add minicluster substructure, a contingency fenced in Part IV. Second, at $\theta_i \simeq 1.1$ rad the cosine potential is mildly anharmonic, so the harmonic estimate (15.7) overstates the closure angle by 4%; the integration of Part V carries the full cosine. Neither remark touches the mass at fixed f_K ; both move the angle.

13.4 How violently is the angle tested?

It is worth quantifying how stable the frozen angle is against the perturbations the early universe supplies. During the frozen regime the slip per Hubble time is (12.4); integrating from the lock crossing ($T_{\text{lock}} \simeq 1.6$ TeV) to onset, the accumulated classical drift is dominated by the last Hubble time before onset and is of order $(m_K/3H)^2 \rightarrow \mathcal{O}(1)$ only *at* onset — before that, the integrated slip is negligible (at $T = 10$ GeV, $(m_K(T)/3H)^2 \sim 10^{-12}$). The lock crossing itself is prompt — GG-A10-Cosmos carries no second inflation there ($N_2 = 0$) — so it supplies no de Sitter kicks at all. The only epoch that could randomize the angle is the primary inflation, and it would: at $H_1 \simeq 1.5 \times 10^{13}$ GeV (from $r = 0.0038$ and $A_s = 2.1 \times 10^{-9}$, as GG-A10-Cosmos now prints) a light angle receives $\delta\theta = H_1/(2\pi f_K) \simeq 3.4$ rad per e-fold. The angle therefore reaches onset untouched *only if* it was not light then — the pinning conditional of Section 13.5.

13.5 The isocurvature question, quantified

A pre-inflationary spectator field acquires de Sitter fluctuations, $\delta\theta = H_{\text{inf}}/(2\pi f_K)$ per horizon mode, which after onset become *isocurvature* perturbations — dark-matter density fluctuations uncorrelated with the photons — of fractional power

$$\beta_{\text{iso}} \simeq \left(\frac{2\delta\theta}{\theta_i} \right)^2 / \mathcal{P}_\zeta = \left(\frac{H_{\text{inf}}}{\pi \theta_i f_K} \right)^2 \frac{1}{2.1 \times 10^{-9}}. \quad (13.9)$$

Planck bounds $\beta_{\text{iso}} < 0.04$ (Planck Collaboration, 2020); at the benchmark ($\theta_i = 1.18$, $f_K = 7.1 \times 10^{11}$ GeV) that bound would tolerate a light angle only below $H_{\text{inf}} \lesssim 2.3 \times 10^7$ GeV. The question must be asked at the right epoch, and earlier revisions of this paper asked it at the wrong one. GG-A10-Cosmos’s history has one inflation, the primary GUT-adjacent epoch with $H_1 \simeq 1.5 \times 10^{13}$ GeV, reheating to $T_{\text{rh}} \simeq 6 \times 10^{10}$ GeV; the later lock crossing of the two-basin descent, at which the compact angle is released, is prompt ($N_2 = 0$ e-folds) and imprints no super-horizon perturbation of any kind. So the lock epoch contributes nothing to (13.9), and the budget is decided entirely by whether the angle was *light during primary inflation*. If it was, $\delta\theta = H_1/(2\pi f_K) \simeq 3.4$ rad per e-fold: the angle is randomized on every scale that exited primary inflation, and the observable universe is in the post-inflationary class — no CMB-scale isocurvature, because the coarse-grained density of a uniformly distributed angle is the same in every patch, but an order-unity minicluster spectrum at the onset horizon scale, and the circle-averaged abundance, which closes the inventory at $\gtrsim 6$ GHz from misalignment alone (Section 36.2) — so that the window of this paper is gone, not the Kaxion. That is this paper’s reading of the unpinned branch, and GG-A10-Cosmos does not share it. GG-A10-Cosmos holds that an angle randomized during inflation carries an independent draw on every observable scale and hence an order-unity cold-dark-matter isocurvature, ten orders above the *Planck* bound, so that the unpinned branch is excluded rather than relocated; and that at the bulk normalization f_{bulk} , where the step is only 0.08 rad per e-fold, an unpinned angle carries the ordinary pre-inflationary fluctuation, eight orders above the bound. The second clause is not in dispute. The first is: the draw an observable scale receives at horizon exit is not the angle the scale carries at the end of inflation, because the diffusion continues for fifty more e-folds at 11 rad² per e-fold, and a periodic density evaluated on an angle whose variance so far exceeds $(2\pi)^2$ is uncorrelated on every scale that left the horizon more than half an e-fold before the end (the step per e-fold sets that length; the uniformity of the one-point distribution is what makes the mean the same in every patch) — white noise at the onset scale, and nothing on the microwave sky. Which of the two readings is right is a computation neither paper prints, and it decides only what the unpinned branch would have looked like, not whether the paper’s history holds; it is registered here as the one live disagreement between the two papers, for the cosmology companion to close. If instead the angle carried a mass well above H_1 through those e-folds, it was pinned, its fluctuations were suppressed, and β_{iso} vanishes identically.

The paper’s claim is the second, and it is a conditional, not a theorem. GG-A10-Cosmos carries the pinning as this paper’s conjecture and does not compute it; this paper does not compute it either. We therefore declare it:

Conjecture C-Kaxion-d (inflation-era pinning). *During the observable e -folds of primary inflation the compact angle carries an effective mass $m_{\text{pin}} \gg H_1 \simeq 1.5 \times 10^{13} \text{ GeV}$ — a pinning potential $V_{\text{pin}} \gtrsim H_1^2 f_K^2 \simeq 10^{50} \text{ GeV}^4$ — whose minimum sits displaced by the misalignment angle, $\Delta\theta \simeq 1.1 \text{ rad}$, from the CP -conserving point of the QCD potential (so it is not the strong- CP mechanism of GG-A9-SM); which releases the angle at the lock crossing identically in every horizon patch; and which is absent below the lock crossing to the level $V_{\text{res}} \lesssim 4 \times 10^{-15} \text{ GeV}^4$ that the neutron bound $|\bar{\theta}| < 10^{-10}$ allows against the QCD-generated $\chi_{\text{QCD}}[1 - \cos\theta] \lesssim 3 \times 10^{-5} \text{ GeV}^4$ of Part V. No sector of the ledger is presently known that could supply such a potential.*

The conjecture is in tension with the flatness clause of Theorem 8.2(ii) as stated for the late universe, and the tension is the point: a potential fifty-five orders of magnitude larger than QCD’s must exist during inflation and vanish afterward, and neither this paper nor GG-A10-Cosmos exhibits its origin. The single-patch history, the isocurvature silence, and the absence of miniclusters all ride on C-Kaxion-d; they are carried in the ledger as *assumed* (Table 12), not as forced. What kills the conjecture is observable: a Kaxion minicluster population, or a confirmed line at the randomized-angle mass of Section 36.2; a CMB isocurvature detection attributable to this sector kills the conjecture outright, as GG-A10-Cosmos also states — the angle was not held; partial pinning, $m_{\text{pin}} \sim H_1$, is not a branch this paper computes. What would establish it is a computation of the angle’s inflation-era potential on GG-A10-Cosmos’s primary-inflation background, registered in Part XI as the sharpest open item on the cosmology side.

14 The decay constant and the misalignment angle

Two quantities must be pinned before the abundance, mass, and couplings of Parts III–V can be evaluated: the Kaxion’s decay constant f_K and its initial misalignment angle θ_i . The logic must be stated once, plainly, because it is easy to overstate. The geometry supplies the *discrete* structure — the count, and a ladder of admissible cap loadings — but not, as this paper stands, the absolute normalization of the decay constant: the computation that would fix it is the anomaly companion’s Open Problem O14, and the bulk scale is *calibrated* through the relic closure. The closure is therefore one equation in two unknowns. The paper fixes the benchmark loading $\Xi_{\text{cap}} = 43$, calibrates the bulk scale by requiring the closure at a generic misalignment angle, and thereafter reads the angle as the value at which the calibrated benchmark saturates the inventory. The frequency is not independent of the abundance; what is independent of it is the window’s discrete ladder and the count.

One decay constant, calibrated. The light eigenstate that survives the lift of the holonomy directions (Section 9.3) is a single canonically normalized field with one decay constant f_K . Its absolute value is not derived in this paper or its parents: the tensor’s boundary kinetic norm and the Wilson line’s differ by e^{2K} (Appendix C.1), and what would place the canonical field’s own normalization relative to the bulk one is the anomaly companion’s Open Problem O14, restated in §26.3. We therefore write $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$ with a bulk normalization $f_{\text{bulk}} \simeq 3.04 \times 10^{13} \text{ GeV}$

calibrated through the relic closure of Section 15.3 at the benchmark loading $\Xi_{\text{cap}} = 43$ (Section 26.3). The benchmark gives

$$f_K = 7.1 \times 10^{11} \text{ GeV}, \quad (14.1)$$

with the discrete ladder of admissible loadings $\Xi_{\text{cap}} \in \{35, \dots, 47\}$ around it that Part V converts into $\nu_K \in [1.59, 2.13]$ GHz (benchmark 1.95 GHz). The susceptibility relation

$$m_K f_K = \sqrt{\chi_{\text{QCD}}} \simeq (75.5 \text{ MeV})^2 \quad (14.2)$$

(Part V) then fixes the mass $m_K = 8.1 \mu\text{eV}$, the laboratory frequency, and the photon coupling from this one number. There is exactly one physical line; the bulk normalization scale ($\sim 3 \times 10^{13} \text{ GeV}$) is an intermediate of the diagonalization, not the decay constant of any propagating field.

The angle follows from the abundance. With f_{bulk} calibrated at the benchmark, the displacement at birth θ_i is the value at which the condensate saturates the derived inventory Ω_{DM} (Section 15.3): $\theta_i \simeq 1.1$ rad by direct integration of the equation of motion. This is the standard $\mathcal{O}(1)$ misalignment angle, fitted to the abundance exactly as in every axion-dark-matter calculation; what the paper claims for it is only that it is generic — neither tuned toward zero nor toward the hilltop — and not that it is derived.

Two invariants used downstream. Two ratios, independent of any normalization choice, become experimental statements later. (i) *The local angular amplitude*, fixed by the local energy density $\rho_{\text{loc}} \simeq 0.45 \text{ GeV}/\text{cm}^3$ through $\rho = \frac{1}{2} m_K^2 a_K^2$ and the invariant (14.2):

$$\theta_0 = \frac{\sqrt{2\rho_{\text{loc}}}}{m_K f_K} = \frac{2.6 \times 10^{-3} \text{ eV}^2}{(75.5 \text{ MeV})^2} \simeq 4.6 \times 10^{-19}, \quad (14.3)$$

a half-attoradian swing oscillating at $\nu_K = 1.95 \text{ GHz}$ — the 10^{-19} -radian bell detection must hear (Part VII). (ii) *The line count* of Part VI carries no scale dependence whatsoever: the light spectrum has exactly one member, at whatever frequency the cap loading selects inside the window — which is why Part VI makes the count test — a reanalysis of any scan crossing the window for a second line anywhere inside it — the cheapest identification measurement the paper asks for.

15 Abundance and the π -locked inventory

15.1 The target is inherited, derived-conditionally

In conventional cosmology the cold-dark-matter density Ω_c is a measured input. In GG-A10-Cosmos it is an output: the late-time inventory is locked to π by the boundary cosmology of the bulk,

$$\Omega_m = \frac{1}{\pi} = 0.3183, \quad \Omega_\Lambda = 1 - \frac{1}{\pi} = 0.6817, \quad \Omega_b = \frac{1}{2\pi^2} = 0.05066, \quad (15.1)$$

whence the dark-matter share is the difference of two locked numbers,

$$\boxed{\Omega_{\text{DM}} = \frac{1}{\pi} - \frac{1}{2\pi^2} = 0.26765} \quad (15.2)$$

— the *target* that any dark-matter candidate of the program must saturate. Two cautions of principle travel with it. The fraction Ω_{DM} is a present-epoch quantity; the timeless content the closure below actually uses is the comoving $\Omega_{\text{DM}}h^2 = 0.1215$ and the constant species ratio $\rho_{\text{DM}}/\rho_b = 2\pi - 1$ (Eq. 6.9). And the π -locked inventory, like the $\alpha_s = 1/(\pi e)$ entry of Table 2, is among the *weaker* (π/e -fraction) members of the track record rather than its rigid integer locks — a point made honestly there. **GG-A10-Cosmos**’s compact-angle dark-matter theorem sharpens this into an obligation: *if* the light compact-angle mode is the dominant cold relic, its abundance must be matched to (15.2), not to an adjustable Ω_c . Against observation the chart performs as follows:

Table 7: The π -locked inventory versus Planck 2018 (Planck Collaboration, 2020) (TT,TE,EE+lowE+lensing, $h = 0.674$). The baryon entry carries the known percent-level tension, flagged in **GG-A10-Cosmos**. The comparison is printed at the precision Planck quotes, and the honest row to read is the baryon one rather than the average of the five.

Quantity	π -lock	Value	Planck 2018	Deviation
Ω_m	$1/\pi$	0.3183	0.3153 ± 0.0073	+1.0% (+0.4 σ)
Ω_Λ	$1 - 1/\pi$	0.6817	0.6847 ± 0.0073	−0.4%
Ω_b	$1/(2\pi^2)$	0.05066	0.0493 ± 0.0003	+2.8%
Ω_{DM}	$1/\pi - 1/(2\pi^2)$	0.26765	0.2645 ± 0.0070	+1.2%
$\Omega_{\text{DM}}h^2$	—	0.1215	0.1200 ± 0.0012	+1.3%

15.2 The misalignment abundance, derived

We now compute what the condensate of Section 13 actually delivers, carefully enough that the closure of Section 15.3 can lean on it. The computation is three lines of physics — bank, count, redeem — plus one piece of care about the stiffening spring.

Bank. At onset the field has banked the potential energy of its displacement:

$$\rho_{\text{osc}} = \chi_{\text{QCD}}(T_{\text{osc}})[1 - \cos \theta_i] \simeq \frac{1}{2} m_K^2(T_{\text{osc}}) f_K^2 \theta_i^2, \quad (15.3)$$

the small-angle form carrying a 4% anharmonic correction to the angle at (13.8), which the integrated closure of Section 15.3 includes.

Count. Once oscillating, the comoving *number* of quanta is an adiabatic invariant even while the mass grows — this is the standard harmonic-oscillator adiabatic theorem, and it is the correct bookkeeping through the stiffening-spring epoch. The number density at onset is

$$n_{\text{osc}} = \frac{\rho_{\text{osc}}}{m_K(T_{\text{osc}})} = \frac{1}{2} m_K(T_{\text{osc}}) f_K^2 \theta_i^2, \quad \frac{n}{s} = \text{const thereafter}, \quad (15.4)$$

with $s = \frac{2\pi^2}{45} g_{*s} T^3$ the entropy density.

Why the number is invariant (proof, and its one caveat). The claim $n/s = \text{const}$ through the stiffening-spring epoch deserves its two-line proof. For an oscillator of slowly varying frequency the action variable $J = \oint p dq = 2\pi E/\omega$ is adiabatically invariant when $\dot{\omega}/\omega^2 \ll 1$; for the condensate, E/ω per comoving volume *is* the comoving quantum number, so $N = \rho R^3/m_K(T)$ is conserved — energy is not (the spring stiffens under the oscillation and does work on it), but quanta are. The caveat is quantitative and honest: at onset itself $\dot{m}/m^2 = bH/m_K(T_{\text{osc}}) = b/3 \simeq 1.4$, so adiabaticity is *marginal* for the first oscillation or two before becoming excellent ($\dot{m}/m^2 \propto t^{-1}$ thereafter). The $\mathcal{O}(1)$ transient this induces is precisely what the instantaneous-onset estimate omits and the direct integration of Part V resolves: it is the factor 1.49 between the estimate 0.090 and the normalization 0.135.

For completeness, the two cosmographic constants used at the “redeem” step, evaluated rather than quoted: today’s entropy density is

$$s_0 = \frac{2\pi^2}{45} g_{*s,0} T_0^3 = 0.4386 \times 3.91 \times (2.349 \times 10^{-4} \text{ eV})^3 = 2.22 \times 10^{-11} \text{ eV}^3 = 2891 \text{ cm}^{-3},$$

using $1 \text{ eV}^3 = 1.302 \times 10^{14} \text{ cm}^{-3}$; and the critical-density unit is $\rho_{\text{crit}}/h^2 = 1.0537 \times 10^{-5} \text{ GeV cm}^{-3}$. Their ratio, $2.744 \times 10^8 \text{ GeV}^{-1}$, is the single conversion factor between early-universe bookkeeping and the observed Ωh^2 .

Redeem. Today each quantum is worth its full zero-temperature mass:

$$\rho_{K,0} = m_K(0) n_{\text{osc}} \frac{s_0}{s_{\text{osc}}} = \frac{1}{2} m_K(0) m_K(T_{\text{osc}}) f_K^2 \theta_i^2 \frac{g_{*s,0} T_0^3}{g_{*s}(T_{\text{osc}}) T_{\text{osc}}^3}. \quad (15.5)$$

The scaling, transparently. Insert $m_K(T_{\text{osc}}) = 3H(T_{\text{osc}}) \propto T_{\text{osc}}^2$ and the onset solution $T_{\text{osc}} \propto m_K(0)^{1/(b+2)}$ from (13.3) into (15.5):

$$\rho_{K,0} \propto m_K(0) f_K^2 \theta_i^2 T_{\text{osc}}^{-1} \propto m_K(0)^{\frac{b+1}{b+2}} f_K^2 \theta_i^2 \propto \theta_i^2 f_K^{\frac{b+3}{b+2}}, \quad (15.6)$$

where the last step used the frame-invariant $m_K(0) \propto 1/f_K$ of (14.2). For the lattice exponent $b = 4.08$ this gives $(b+3)/(b+2) = 7.08/6.08 = 1.165$ — the celebrated near-7/6 power, here derived rather than quoted. Restoring the normalization by direct numerical integration of the equation of motion through the onset (Part V, where the instantaneous-onset estimate is also shown for comparison) gives the master formula

$$\Omega_K h^2 \simeq 0.135 \theta_i^2 \left(\frac{f_K}{10^{12} \text{ GeV}} \right)^{1.165}. \quad (15.7)$$

(The series carries one misalignment normalization: GG-A10-Cosmos prints the same integrated value at its misalignment equation, and Section 25.5 records the one point on which the two papers still differ.)

Physical picture for the odd exponent. Why $f_K^{1.165}$ and not f_K^2 , as the banked energy (15.3) would suggest? Because a larger f_K means a lighter field, which lets go *later* (T_{osc} lower), when the universe has less entropy left to dilute it — but also banks each quantum at a smaller onset mass. The competition between “later release” and “cheaper quanta” shaves the naive f_K^2 down by $T_{\text{osc}}^{-1} \propto f_K^{1/(b+2)} \cdot (\dots)$, landing at $(b+3)/(b+2)$. A reader who can reproduce this paragraph owns the misalignment mechanism.

15.3 The closure: from the π -chart to the misalignment angle

Now run the logic in the direction unique to this program. Equation (15.7) has *one* unknown left once the calibration convention is adopted: with the decay constant at its calibrated benchmark $f_K = 7.1 \times 10^{11}$ GeV (14.1), the exponent fixed by lattice QCD, and the target fixed by the derived chart (15.2) — converted to $\omega = \Omega_{\text{DM}} h^2$ with $h = 0.674$ — the Planck-fit CMB-branch convention shared with GG-A10-Cosmos’s two-stage Hubble corridor, i.e. from the *same series* that derived the inventory — the only quantity not yet pinned is the initial displacement θ_i . Setting $\Omega_K h^2$ equal to the π -locked value $\Omega_{\text{DM}} h^2 = 0.1215$ (Table 7) and solving for the angle:

$$\theta_i \simeq 1.22 \text{ rad (harmonic)}, \quad \theta_i = 1.18 \pm 0.04 \text{ rad (full cosine, tabulated } \chi_{\text{QCD}}(T)). \quad (15.8)$$

Both numbers are read from the integration of Section 25 on the tabulated equation of state and susceptibility, the first with a harmonic restoring force and the second with the full cosine; the closed-form estimate (25.2) is not the calculation of record and returns 1.16 for the first of them. The direction of inference must be stated without overclaiming. The abundance was not used to fit a free f_K ; but neither did the geometry deliver f_K . The bulk normalization was calibrated so that the benchmark row of the cap ladder saturates the inventory at a generic angle, and the angle printed above is that generic value read back. What the closure establishes is a consistency statement — the benchmark line and the π -locked inventory are compatible at an untuned angle — and what it does not establish is a derivation of the frequency from the integers. At $\theta_i \simeq 1.18$ rad the cosine potential is mildly anharmonic; the integration carries the full cosine, which lowers the angle by 4.0% against the harmonic force on the same tables.

Honesty requires naming the load-bearing assumption: that the compact-angle condensate is the *sole* saturator of (15.2) — no second component, no early decay, no late entropy injection diluting it. That assumption is GG-A10-Cosmos’s *misalignment-completeness conjecture*, carried here as the named conjecture it is: Part V propagates its uncertainty, and Part XI states what would close it. Within the conjecture the map is simple: each rung of the discrete window carries its own closure angle, and integrating the calculation of record rung by rung gives $\theta_i \simeq 1.05$ rad at $\Xi_{\text{cap}} = 35$, 1.18 at the benchmark $\Xi_{\text{cap}} = 43$ and 1.23 at $\Xi_{\text{cap}} = 47$ — every rung generic — so the closure disfavors no rung and, equally, cannot select one.

15.4 What the condensate does to the rest of cosmology

Three consistency checks close the section, each cheap and each passed. *Radiation budget:* at BBN temperatures the condensate, oscillating since $T \simeq 1$ GeV, carries $\rho_K/\rho_{\text{rad}} \simeq 6 \times 10^{-7}$ (Table 8)

— the Kaxion does not perturb N_{eff} or light-element yields; before onset its frozen energy is smaller still. *Equality*: a component scaling as a^{-3} from $t \sim \mu\text{s}$ onward with today’s share (15.2) reproduces the standard matter–radiation equality at $z_{\text{eq}} \simeq 3400$; the Kaxion changes nothing in the linear growth history that GG-A10-Cosmos already matched. *Phase-space*: the condensate is bosonic and classical (Section 16); no Tremaine–Gunn-type bound applies, and micro-eV occupation numbers face no degeneracy obstruction whatsoever — in sharp contrast to light *fermionic* dark matter, which the same phase-space argument forbids.

16 Why the Kaxion is cold

“Cold” is the load-bearing adjective of ΛCDM , and a μeV -mass particle earns it in a way that deserves explicit proof rather than assertion. Three independent computations — pressure, velocity, wavelength — each return the same verdict.

16.1 Pressure: the oscillating condensate is dust

For the homogeneous oscillating field, the energy density and pressure are

$$\rho_K = \frac{1}{2}\dot{a}_K^2 + V(a_K), \quad p_K = \frac{1}{2}\dot{a}_K^2 - V(a_K). \quad (16.1)$$

Once $m_K \gg H$, the oscillation is fast compared with the expansion, and the virial theorem for the (effectively quadratic) potential equates the cycle-averages $\langle \frac{1}{2}\dot{a}_K^2 \rangle = \langle V \rangle$. Hence

$$\langle p_K \rangle = 0 \quad \implies \quad \langle w \rangle = 0 : \quad (16.2)$$

the condensate is pressureless dust on any timescale longer than an oscillation period (10^{-8} s today — every cosmological timescale qualifies). Equivalently, inserting (16.2) into the continuity equation $\dot{\rho} = -3H(\rho + p)$ gives $\rho_K \propto R^{-3}$, the matter redshift law already used in (15.4). The instantaneous equation of state oscillates wildly between $+1$ and -1 at $2\nu_K$; cosmology only ever sees its average.

16.2 Velocity: the coldest matter in the universe

A thermal relic remembers the temperature of its bath. The condensate has no bath to remember; its momentum spread is bounded by the only scale in the problem at onset — the horizon. Under the single-patch history the field is homogeneous up to the adiabatic perturbations, so gradients at onset carry momenta at most $p_{\text{osc}} \sim H_{\text{osc}}$, which subsequently redshift as $1/R$:

$$p_0 \sim H_{\text{osc}} \frac{R_{\text{osc}}}{R_0} = H_{\text{osc}} \frac{T_0}{T_{\text{osc}}} \left(\frac{g_{*s,0}}{g_{*s,\text{osc}}} \right)^{1/3} \simeq (1.2 \times 10^{-18} \text{ GeV}) \times (8.9 \times 10^{-14}) \simeq 1.0 \times 10^{-22} \text{ eV}, \quad (16.3)$$

using the onset values of (13.4). The primordial velocity dispersion today is therefore

$$v_{\text{prim}} = \frac{p_0}{m_K} \sim 1 \times 10^{-17} \hat{=} 4 \times 10^{-9} \text{ m/s} : \quad (16.4)$$

a few nanometers per second. For comparison, a 100-GeV WIMP decouples with $v \sim 0.3$ and cools to $\sim 10^{-12}$ today; a keV warm-dark-matter fermion retains $v \sim 10^{-8}$; the Kaxion sits ten further orders below the warm benchmark. Of course, infall into the Galaxy subsequently endows the local condensate with the virial velocity $v_{\text{vir}} \sim 10^{-3}$ — but that velocity is *gravitational*, identical for every cold species, and carries no memory of the production mechanism. Intrinsically, the Kaxion is the coldest matter the early universe can manufacture.

The corresponding free-streaming damage to structure is nil. The comoving free-streaming length is dominated by the epoch around equality, $\lambda_{\text{fs}} \sim v_{\text{eq}}(ct_{\text{eq}})(1 + z_{\text{eq}})$ with $v_{\text{eq}} = v_{\text{prim}}(1 + z_{\text{eq}}) \simeq 4 \times 10^{-14}$, giving $\lambda_{\text{fs}} \sim 10^{-6}$ pc — sub-solar-system, nine orders of magnitude below the dwarf-galaxy scales where warm dark matter is constrained. Every structure Λ CDM predicts, the Kaxion predicts.

16.3 Wavelength: cold waves, not fuzzy waves

A light boson invites the question: does its wave nature smooth structure? The de Broglie wavelength in the virialized halo ($v \sim 10^{-3}$) is

$$\lambda_{\text{dB}} = \frac{2\pi}{m_K v} \simeq 154 \text{ m}, \quad (16.5)$$

seventeen orders of magnitude below the kiloparsec scales of galactic cores. The contrast class is fuzzy dark matter (Hu, Barkana, & Gruzinov, 2000), which achieves kpc-scale wave effects only by pushing the mass to 10^{-22} eV — seventeen orders below the Kaxion. The wave pressure (quantum pressure) suppression scale, $k_J \sim \sqrt{m_K H}$, lies correspondingly far beyond any observable cosmological wavenumber. The Kaxion is a wave, but a *short* wave: within each de Broglie volume it packs

$$\mathcal{N}_{\text{occ}} = \frac{\rho_{\text{loc}}}{m_K} \lambda_{\text{dB}}^3 \simeq 2.0 \times 10^{26} \quad (16.6)$$

quanta — a magnificently classical field, which is precisely what a resonant detector wants (Section 17), while remaining utterly incapable of erasing kiloparsec structure. In one sentence: *the Kaxion behaves as ordinary cold dark matter everywhere ordinary cold dark matter is tested, and as a coherent wave only inside an instrument built to notice.*

16.4 The condensate’s diary: the amplitude through cosmic history

The Bessel envelope of Section 13.2 plus number conservation let us write the dial’s angle at every epoch — a diary worth keeping, because each entry is an independent consistency check against an era of known physics:

$$\theta(T) = \frac{\sqrt{2\rho_K(T)}}{m_K f_K}, \quad \rho_K(T) = m_K \left(\frac{n}{s}\right) s(T), \quad \frac{n}{s} = 5.5 \times 10^4 \text{ (the closure value)}. \quad (16.7)$$

Table 8 evaluates (16.7) across cosmic history — the dial’s amplitude at the onset, through the QCD epoch and recombination, down to today — exhibiting the nearly nineteen orders of quiet decline that make the present oscillation both minute and perfectly coherent.

Table 8: The dial’s swing amplitude through history (local-halo value in the last row). Four epochs, each with its temperature, the swing amplitude there, and what that amplitude means for the physics of that epoch. The table answers an objection before it is raised: a field starting at an order-one angle might be expected to disturb something, and the last column says at each stage why it does not. The final row is the local-halo value, which is what an experiment on Earth actually sees.

Epoch	T	θ	Check against known physics
Onset	1.0 GeV	1.1 (θ_i ; the diary formula applies below T_{QCD})	banked energy $\ll$ radiation: no expansion-history disturbance
BBN	1 MeV	3.6×10^{-7}	$\rho_K/\rho_{\text{rad}} \sim 6 \times 10^{-7}$: light-element yields untouched
Equality	0.80 eV	1.5×10^{-16}	$\rho_m = \rho_{\text{rad}}$ by definition, $\rho_K = 0.84 \rho_{\text{rad}}$: standard $z_{\text{eq}} = 3400$
Today (halo)	—	4.6×10^{-19}	the haloscope target amplitude of Part VII

The diary reads as a nineteen-order decrescendo, and its moral is the detection problem in one line: the same field that once leaned 1.1 radians now whispers at half an attoradian, and the whisper is all that was ever going to reach a laboratory thirteen billion years downstream. Nothing about the decrescendo is adjustable; it is Hubble friction applied to Eq. (13.6), era by era.

16.5 Wave pressure, quantified: the Jeans scale

The qualitative statement that wave effects are irrelevant deserves its formula. Linear perturbations of a coherent condensate obey a fluid equation with an effective, scale-dependent sound speed $c_s^2 = k^2/(4m_K^2 R^2)$ — the quantum pressure — which balances gravity at the Jeans wavenumber

$$\frac{k_J}{R} = (16\pi G \bar{\rho} m_K^2)^{1/4}. \quad (16.8)$$

Evaluating today at the mean dark-matter density ($\bar{\rho} = 9.8 \times 10^{-48} \text{ GeV}^4$, $G = 6.71 \times 10^{-39} \text{ GeV}^{-2}$) at the physical mass $m_K = 8.1 \mu\text{eV}$: $16\pi G \bar{\rho} m_K^2 = 2.2 \times 10^{-112} \text{ GeV}^4$, so $k_J/R = 1.2 \times 10^{-28} \text{ GeV}$ and

$$\lambda_J = \frac{2\pi R}{k_J} \simeq 1.0 \times 10^{13} \text{ m} \simeq 3 \times 10^{-4} \text{ pc} : \quad (16.9)$$

wave pressure smooths structure below a third of a *milliparsec* ($\sim 60 \text{ AU}$) — solar-system scale, six orders below the smallest galaxy. Run the same formula at $m = 10^{-22} \text{ eV}$ and λ_J inflates by $(m_K/10^{-22} \text{ eV})^{1/2} \sim 3 \times 10^8$ into the tens-of-kiloparsec regime — which is the entire content of the fuzzy-dark-matter phenomenology, and the precise quantitative sense in which the Kaxion is not it. One formula, two regimes, seventeen orders of mass between them.

16.6 Verdict

Pressureless on cosmological timescales (Eq. 16.2); primordial velocities at the $\sim 10^{-17}$ level (Eq. 16.4); free streaming at the microparsec level; wave pressure below a third of a milliparsec; a de Broglie wavelength of ~ 150 m. By every operational definition in the structure-formation literature, the Kaxion is cold, and colder than any thermal candidate can be. The adjective in “cold dark matter” is, for once, a theorem rather than a hope.

17 Interactions with matter

17.1 The portal sheet, previewed

Part V derives the couplings; this section needs only their structure. Every interaction of the Kaxion descends from the BI-6 ledger through the anomaly-inflow package of GG-A5-GDOS, and therefore (i) is suppressed by one power of f_K , and (ii) arrives in ledger-locked ratios:

$$\mathcal{L}_{\text{int}} \supset \frac{\alpha_s}{8\pi} \frac{a_K}{f_K} G\tilde{G} + \frac{g_{K\gamma\gamma}}{4} a_K F\tilde{F} + \frac{\partial_\mu a_K}{2f_K} (C_p \bar{p}\gamma^\mu\gamma_5 p + C_e \bar{e}\gamma^\mu\gamma_5 e). \quad (17.1)$$

Four channels, one ledger: gluonic (the mass-making channel), electromagnetic (the haloscope channel), nucleon-spin (the NMR channel), and electron-spin (the magnon channel); the ledger’s baryonic entry is real at the anomaly level but opens no channel below M_{lock} (Part VI). Gravity, universal and unsuppressed, is the fifth channel and the one astronomy has been reading all along.

17.2 Why every recoil experiment is blind

The great direct-detection program of the WIMP era — xenon time-projection chambers, cryogenic crystals — is structurally incapable of seeing the Kaxion, for two independent reasons, each fatal alone.

Kinematics. A halo Kaxion carries kinetic energy

$$E_{\text{kin}} = \frac{1}{2} m_K v^2 \simeq 4 \times 10^{-12} \text{ eV}, \quad (17.2)$$

fourteen orders of magnitude below the $\sim \text{keV}$ thresholds of recoil detectors. There is no recoil to detect: an elastic “collision” with a nucleus transfers momentum $\sim m_K v$, displacing the nucleus by less than a zeptometer-scale velocity kick.

Dynamics. Even waiving kinematics, the per-particle cross sections scale as $g^2 \sim 1/f_K^2 \sim 2 \times 10^{-24} \text{ GeV}^{-2}$ times further phase-space suppression — the same $1/f_K$ that guarantees cosmological stability guarantees per-quantum invisibility.

17.3 Why resonant conversion works anyway

The rescue is the classicality computed in (16.6). A detector need not catch quanta one at a time; it can couple to the *amplitude* of a classical field that is coherent across the entire apparatus ($\lambda_{\text{dB}} \gg$ any cavity) and in time for $Q_a \sim v^{-2} \sim 10^6$ oscillations. In a static magnetic field B_0 , the

electromagnetic portal in (17.1) converts the oscillating $a_K(t) = \theta_0 f_K \cos(m_K t)$ into an oscillating electromagnetic source at exactly $\nu_K = m_K/2\pi$; a cavity tuned to ν_K integrates that drive coherently, and the signal power scales as $g^2 \rho_{\text{loc}}/m_K \times B_0^2 V Q$ — linear in the *density*, not quadratic in a wavefunction overlap. Coherence converts a hopeless per-particle problem into a merely heroic radio-engineering problem. The single tone rides on this same carrier: the haloscope of Part VII is a radio tuned to a 10^{-19} -radian bell, and the bell rings at one pitch inside the fenced band.

17.4 What Part III has established

The Kaxion is born displaced at an $\mathcal{O}(1)$ angle by the same runtime that selects the vacuum (Sections 12–13); it begins oscillating at $t \simeq 0.3 \mu\text{s}$ and redshifts as dust thereafter; with its decay constant at its calibrated benchmark (Section 14), demanding it saturate the derived-conditional target $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ returns the misalignment angle $\theta_i \simeq 1.1$ rad (Section 15); it is provably cold by pressure, velocity, and wavelength (Section 16); and it touches matter only through ledger-locked, $1/f_K$ -suppressed portals that single out resonant conversion as the detection strategy (this section). Part IV places this candidate among its competitors; Part V converts the geometry into the mass and the coupling.

18 Part III in plain English: where the dark matter came from, in one chain of events

Cosmology is where a light field stops being a curiosity and starts being dark matter, and the account here is a single chain with one named branch point — whether the angle was held still through inflation, which the paper assumes and names.

Once the early universe has cooled to the lock crossing, the compact-angle field is released sitting at some angle — not at the bottom of its potential, because what held it through inflation was displaced from it there. It then stays where it is, because the expansion is faster than the field can respond. When the expansion rate falls to the field’s mass, the field starts to oscillate, and a coherently oscillating scalar dilutes exactly as cold matter does. What was a frozen displacement becomes a density of non-relativistic matter, and it has been diluting as matter ever since.

The one number this chain needs is the angle the field started at. It is not derived here, and the paper says so plainly: it is fixed backwards, by requiring that the resulting abundance match the measured one. That is one input, and the honest description of this Part is that it converts one observation — how much dark matter there is — into a statement about where the field started, and calibrates the one scale the geometry has not supplied. Everything else in the chain is fixed by the geometry and by standard cosmology.

What the Part invokes beyond the field itself is one thing, named and killable: a potential that held the angle during inflation and is gone today. There is no other new species and no new interaction below the lock crossing. The dark matter is a field that was already required, doing the only thing it could do.

Part IV Context

The Kaxion among candidates, and its astrophysical reach

This Part places the Kaxion among the field’s candidates and states, plainly, the strongest claim the framework supports. We proceed in four steps. First the *candidate landscape* (§19): the empirical status of every leading candidate as of 2026 — the ground has shifted decisively, and every shift has been a null — together with the structural comparison on which honest preference rests. Then the *constitutional verdict* (§20): that within GG-Theory the Kaxion is not merely the best candidate but the only admissible one, required and sufficient, with every rival excluded by a clause the theory adopted for other reasons. Then the *astrophysical reach*, in two parts — early, massive structure in the JWST era (§21) and the cusp–core problem (§22). Finally a summary of what the Part establishes (§22.9).

19 The candidate landscape: empirical status and structural comparison

The honest case for any candidate rests on two readings, and we give both before drawing the verdict: where each candidate actually stands in the laboratory and the sky (the empirical status), and whether its existence is forced, its mass predicted, and its abundance derived (the structural comparison).

19.1 The empirical status as of 2026

We record first where each candidate stands, because the last few years have been decisive — and uniformly negative for the front-runners of the previous decade.

The WIMP, squeezed between the neutrino fog and the LHC. The weakly interacting massive particle dominated the search for thirty years on the strength of the “WIMP miracle” — a weak-scale annihilation cross-section freezing out near the observed abundance. The tonne-scale xenon program has now driven the spin-independent WIMP–nucleon cross-section to world-leading bounds without a signal: LUX-ZEPLIN (LZ) excludes $\sigma_{SI} \lesssim 2 \times 10^{-48} \text{ cm}^2$ near 40 GeV and finds no WIMP above 9 GeV (LZ Collaboration, 2024), with XENONnT and PandaX-4T corroborating within a small factor (XENON Collaboration, 2023; PandaX Collaboration, 2024). Two ceilings are now closing on the candidate at once. From below, XENONnT and PandaX-4T have begun to register solar ^8B neutrinos — the onset of the irreducible “neutrino fog” at which a xenon recoil detector loses discriminating power. From above, the WIMP’s natural ultraviolet home, weak-scale supersymmetry, is excluded by the LHC to gluino and first-generation squark masses $\gtrsim 2 \text{ TeV}$,

stops $\gtrsim 1.3$ TeV, and electroweakinos ~ 1 TeV, with no superpartner found (ATLAS & CMS Collaborations, 2024). The thermal WIMP is not refuted in every corner — compressed and pMSSM spectra survive — but its miracle has lost both its abundance coincidence and its UV completion.

The axion and the wider ALP family. Through the opening left by the WIMP nulls, the axion has moved from alternative to front-runner. The QCD axion solves strong CP and is a cold-dark-matter candidate, a pseudo-Goldstone of a spontaneously broken Peccei–Quinn symmetry; its mass and photon coupling lie on a line, but the decay constant — hence the mass — is free over roughly ten decades, and the haloscope program (ADMX, CAPP, HAYSTAC) has only in the last few years begun to touch DFSZ depth in narrow bands. The generic axion-like particle (ALP) relaxes even the QCD line: (m, g) unconstrained, it is the maximally flexible and minimally falsifiable option, the null hypothesis of the field.

Sterile neutrinos and the 3.5 keV line. A keV-scale gauge-singlet (“sterile”) neutrino, produced by mixing, was the warm-dark-matter front-runner, and the unidentified 3.5 keV X-ray line from clusters was read as the decay of a 7 keV state. Deeper observation has not been kind: the dark-matter origin of the line is now broadly disfavored (blank-sky and dwarf non-detections, plausible atomic explanations), and warm dark matter is in growing tension with small-scale structure; laboratory mixing is being probed by KATRIN/TRISTAN (Boyarsky et al., 2019).

Ultralight (fuzzy) dark matter. A pseudoscalar of mass $\sim 10^{-22}$ eV with a kiloparsec de Broglie wavelength was invoked to soften galactic cores. The Lyman- α forest now excludes $m \lesssim 2 \times 10^{-20}$ eV, closing the canonical window by roughly two orders of magnitude; a few dwarf-galaxy analyses still favor $\sim 1.5 \times 10^{-22}$ eV, but the tension is severe (Hu, Barkana, & Guzinov, 2000; Rogers & Peiris, 2021).

Primordial black holes. Only the asteroid-mass window, $\sim 10^{17}$ – 10^{22} g, remains open for primordial black holes to be *all* of the dark matter; every other mass is constrained to a few percent or less by lensing, evaporation, and dynamical bounds (Green & Kavanagh, 2021). Even the open window requires a tuned, enhanced primordial power spectrum to form the holes in the first place.

Self-interacting and hidden-sector candidates. Self-interacting dark matter (SIDM) addresses cores with a velocity-dependent, tuned σ/m ; kinetically-mixed dark photons, SIMPs, and asymmetric dark matter each add a new sector with its own couplings. None predicts its mass; none predicts its abundance; each adds parameters rather than removing them.

The empirical summary. Four decades of magnificent experiments have excluded vast tracts of parameter space without identifying a particle. The field has, in consequence, returned to its widest question in fifty years — *not where is it, but what is it* — and to a landscape of candidates none of which is forced, predicts its mass, or derives its abundance. That is the backdrop against which the structural comparison, and then the constitutional verdict, should be read.

19.2 The structural comparison

The Kaxion competes in a crowded field, and the comparison is most honest when made on structural questions rather than on taste: Is the candidate’s existence forced or chosen? Is its mass predicted? Is its abundance derived or fitted? Does it solve an independent problem? Can a detection *identify* it?

The WIMP. A stable weak-scale particle, motivated by the elegant coincidence that a weak-scale annihilation cross-section freezes out at roughly the observed abundance. But its existence is postulated, its mass spans orders of magnitude, its abundance coincidence has been eroded by null results down to $\sigma_{SI} \lesssim 10^{-47} \text{ cm}^2$, and a detection would identify a mass and cross-section — two numbers no theory pinned in advance.

The QCD axion, KSVZ and DFSZ. The closest relatives, and the comparison that matters most. Both solve strong-CP, as the Kaxion does; both are cold misalignment condensates, as the Kaxion is. But in both, the Peccei–Quinn symmetry is *installed by hand*, the decay constant — hence the mass — is a free parameter spanning 10^{-12} – 10^{-2} eV, the initial misalignment angle is scanned, and the quality of the global symmetry against gravitational violation is an acknowledged open wound. The Kaxion differs on every count while sharing the virtues: the field is forced (T-Kaxion-1), the shift symmetry is gauge-protected (no quality problem), the decay constant’s benchmark is calibrated once through the π -locked inventory with the angle read back as generic, the ladder around it is read from the integers by ansatz, and — since $E/N = 8/3$ — the photon coupling lands exactly on the DFSZ line. Operationally: *the Kaxion is what a DFSZ axion becomes when every one of its free parameters is replaced by a theorem, a declared calibration, or a fenced conjecture*, plus two fingerprints DFSZ does not have.

The generic ALP. An axion-like particle with unconstrained (m, g) is the null hypothesis of the haloscope field — maximally flexible, minimally falsifiable. The Kaxion is its antithesis: one laboratory window with a marked benchmark, one coupling, one count.

The gauged-baryon rival, built in four dimensions. One comparison deserves more than a scorecard row, because it is the only rival that answers *this paper’s own structural question* — what does it cost to gauge baryon number? — and answers it the orthodox way. Benakli, Goodsell and Goudeau (Benakli et al., 2026) gauge $U(1)_B$ in four dimensions with $\tilde{B}|_{\text{SM}} = B$ exactly, cancel the resulting anomalies by adding a secluded sector of chiral fermions, and give the baryonic vector its mass by the Higgs mechanism at a scanned breaking scale. The construction of this cohort cancels the same anomaly by BI-6 inflow, adds no matter at all, and removes the zero mode by the Dirichlet clause of its admissible class (GG-A3-BI6 C9; GG-A9-SM). These are the two mathematically legitimate answers to one consistency demand, and they are experimentally distinguishable, because the added-matter route necessarily populates the multi-TeV range with colored and charged exotics and the inflow route necessarily does not.

Their paper’s most quotable structural finding is the one that bears hardest here, and it is theirs rather than ours: anomaly cancellation is cheap, and *viability* is expensive. The anomaly equations admit many solutions; what actually fixes the model is that the exotics must decay, that no colored relic may be stable, that the dark matter needs a discrete parity, and that the couplings should converge at a high scale. Meeting those demands forces a secluded sector, an enlarged scalar potential, and a light pseudoscalar that then does real work in the relic calculation. Consistency, in short, drags in an entire apparatus. The cohort pays for the same consistency with one topological term already present for other reasons — which is the architecture argument of GG-A3-BI6 made by someone else’s construction.

For *this* paper the discriminant is the dark sector, and it is the cleanest in the whole comparison. Their candidate is a Majorana state of the secluded sector at 0.1–1 TeV, with the relic set by annihilation into the light pseudoscalar, by co-annihilation, or into W^+W^- , and it is testable by direct detection. Ours is the compact-angle pseudoscalar at 8.1 μeV . *Two answers to one question place dark matter sixteen to seventeen orders of magnitude apart.* A nuclear recoil from a TeV-scale state coupled through a baryonic vector would be evidence for their class and cannot be produced here; a haloscope line inside (36.1), standing alone at the DFSZ address, would be evidence for this one and has no counterpart there. Either experiment discriminates between the completions outright, which is more than most dark-matter comparisons can say.

Two riders keep the comparison honest. First, one of their non-unified benchmarks places a baryonic Z' at 3.85 TeV — the locked mass of P1-KK to the printed digit (a placement P1-KK states on a vanishing seam kinetic term, $\rho_K = 0$; a seam term of the size GG-A3-BI6 registers as O15 would move the whole cluster to ~ 22 TeV). That is a scan output landing on our number by coincidence, and it is *staging, never evidence*: what it buys is a clean head-to-head. At equal mass the coupling-squared ratio is $(0.565/0.0275)^2 \approx 420$ — the 0.565 being the value GG-A3-BI6 derives (T-BI6-6.4) — but a coupling ratio is not a rate, and here it is not a rate at all: under the boundary-condition termination the cohort’s baryonic vector is a Dirichlet mode that endpoint-localized quarks do not produce at tree level — the quark depths being open in GG-A9-SM’s cascade-depth lemma, so that the tree-level dijet rate on our side is not printed rather than zero — and there is no dijet line on our side to set against theirs (Arisaka, 2026, P1-KK). The observable discriminant is the pattern — their lone, quark-universal line, its dark sector below it and most of its anomaly-canceling spectrum above it, against the cohort’s color octet at the lock mass with its angular partners, and no exotics anywhere. Second, and against us: their existence proves that a gauged baryonic symmetry can be built with a *scanned* coupling. What the laboratory therefore discriminates is the *locking* of that coupling to ledger integers — and by no means the mere existence of a baryonic gauge field.

Sterile neutrinos, fuzzy dark matter, SIDM, PBHs. A keV sterile neutrino is warm — its mass bounded from below by phase-space packing (Tremaine & Gunn, 1979) — in growing tension with small-scale structure — and its 3.5 keV candidate line has not survived deeper observation. Fuzzy dark matter ($m \sim 10^{-22}$ eV) buys kpc-scale cores at the price of suppressing the small-scale power that high-redshift observations increasingly demand (Section 21); the Kaxion, seventeen orders heavier, is immune to those constraints and forgoes the cores. Self-interacting dark matter

addresses the same cores with a tuned cross-section — a dial, where the Kaxion offers a computed, negligible portal and a cuspy prediction (below). Primordial black holes survive only in narrowing mass windows and require a tuned primordial spectrum. None of the four solves an independent particle-physics problem; none predicts its own abundance. Table 9 scores the rivals and the Kaxion against the five structural questions just posed; the pattern to notice is columnar — every rival fails through a *dial*, and the Kaxion’s column contains none.

Table 9: Structural scorecard. “Forced” means the candidate’s existence follows from consistency of a framework adopted for other reasons. The columns ask structural questions rather than phenomenological ones: is the candidate forced, is its scale forced, is its initial condition forced, does it sit on a line, and does it predict anything beyond its own mass and cross-section. A row of noes is not a criticism of a candidate’s physics — several of them are excellent physics — but a statement about how much of each candidate has to be supplied by hand.

Candidate	Forced?	Mass predicted?	Abundance derived?	Strong CP?	Identification handle
WIMP	no	no	freeze-out (eroded)	no	none beyond (m, σ)
KSVZ/DFSZ axion	no	no (f free)	no (θ_i free)	yes	QCD-line slope only
Generic ALP	no	no	no	no	none
Sterile ν	no	no	model-dep.	no	X-ray line
Fuzzy DM	no	no	no	no	core morphology
SIDM	no	no	no	no	tuned σ/m
PBH	no	window	tuned spectrum	no	lensing
Gauged- B secluded WIMP (Benakli et al., 2026)	no	no (scanned)	relic scan	no	Z' + secluded spectrum
Kaxion	yes (T- Kaxion- 1)	window (ladder by ansatz around a calibrated benchmark)	target inherited (GG-A10- Cosmos, cond.); matched at a read-back angle	yes (T- Kaxion-2)	fenced window + DFSZ address + count

The last column is the one this paper most cares about. Every other entry, if detected, would start an argument; the Kaxion, if detected in its window, at the DFSZ address, standing alone — would end one.

20 The GG-Theory verdict: within the framework, the Kaxion is required and sufficient

The comparison above was made on the field’s own terms — structural virtues weighed candidate by candidate. GG-Theory permits a stronger statement, which we make carefully and with its condition kept in plain view: *within the framework, the Kaxion is not merely the best candidate; it is the only admissible one.* Every competing candidate requires an ingredient the constitution forbids, and the one ingredient the constitution supplies — the boundary pseudoscalar already forced by anomaly freedom — is both *necessary* (it must exist) and *sufficient* (it carries the observed abundance at a generic misalignment angle once its one scale is calibrated to it). This is explicitly a conditional claim, conditional on the framework being correct; it is not a model-independent statement about nature, and we do not present it as one. What makes it worth stating is that the framework is sharply falsifiable (Parts VI–VIII), so the condition can be settled rather than merely assumed. Put as a testable proposition: *if GG-Theory is right, then within it the WIMP, the Peccei–Quinn axion, the sterile neutrino, fuzzy dark matter, and primordial black holes are structurally excluded — each by a clause the theory adopted for independent reasons — and a haloscope null across the Kaxion window retires the calibrated placement, its ladder and Conjecture C-Kaxion-d — the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinning branch, or the relic is not this sector’s — not the framework, and this paper does not claim otherwise.*

20.1 Required: the Kaxion must exist, and it must suffice

Two theorems make the positive half of the verdict.

Necessity. Anomaly freedom forces the Green–Schwarz two-form B_2 (GG-A3-B16); its boundary dual is forced, and with the circle’s holonomies lifted it is the one light, stabilized direction of the sector: the Kaxion (T-Kaxion-1, Theorem 8.2). The field is present *before* the dark-matter question is posed — it is the least-new sector the theory contains, not an addition to it.

Sufficiency. The π -locked inventory sets the target $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$, and the forced field saturates it at the calibrated benchmark with a generic angle (Part V), so it accounts for the *entire* dark-matter budget, not a fraction of it. Required and sufficient: the theory needs the Kaxion to exist, and once it exists it needs nothing else.

The ordering rule. Maximum Symmetry, Minimum Freedom (MSMF, GG-A1-Constitution AX-4) converts this into a search instruction: a sector already forced must be examined first, and a new sector may be contemplated only if the forced one fails. Parts III and V show it does not fail. The search therefore *ends* at the Kaxion; it does not pass through it on the way to something tunable.

20.2 Excluded: what each rival needs that the constitution forbids

The negative half of the verdict is candidate-by-candidate, and in each case the exclusion is not a preference but a constitutional clause.

The WIMP — no superpartner, no spare weak-scale matter. A thermal WIMP needs a new stable weak-scale multiplet, generically a superpartner (the neutralino) or extra electroweak matter. MSMF forbids introducing fields and continuous parameters while a forced sector stands available; the boundary matter content is fixed by GG-A9-SM; and supersymmetry — the WIMP’s ultraviolet home — doubles the spectrum with a tower of new states and soft parameters that the constitution does not admit. *There is no WIMP, because there is no supersymmetry and no spare weak-scale matter to be one.*

The Peccei–Quinn axion — no residual CP problem, and no PQ scalar. The QCD axion exists for one reason: to solve strong CP, by adding a global $U(1)_{PQ}$ and a complex scalar whose phase is the axion. GG-Theory removes both the need and the means. *The need:* strong CP is *already* addressed — the series carries GG-A9-SM’s structural protection and the Kaxion’s relaxation as one mechanism at GG-A9-SM’s grade (GG-A1-Constitution P7; Part V), so there is no residual CP problem for a separate axion to address. *The means:* there is no Peccei–Quinn scalar in the spectrum — the Kaxion is a gauge-topological angle (the dual of B_2 and a Wilson line), with no radial order parameter to break, so the very field a PQ axion is built from does not exist, and with it go the decay-constant dial, the scanned misalignment angle, and the global-symmetry quality problem. *There is no PQ axion, because there is no CP problem left to solve and no scalar to solve it with.* What remains is not an installed axion but a forced one: the Kaxion is the relaxation the theory *compels*, the DFSZ axion with every free parameter replaced by a theorem, a declared calibration or a fenced conjecture.

The sterile neutrino — no spare singlet, no warm relic. A keV sterile neutrino needs a gauge-singlet fermion at keV mass with a tuned active-sterile mixing. The lepton and neutrino sector is fixed by GG-A9-SM (three generations, the seesaw the ledger sets); there is no spare singlet slot at keV, and a *warm* relic would in any case collide with the *cold*, theorem-derived Kaxion that already saturates Ω_{DM} . *There is no sterile-neutrino dark matter.*

Fuzzy dark matter — the forced pseudoscalar is seventeen orders too heavy. Ultralight dark matter needs a pseudoscalar with decay constant $\sim M_{Pl}$ and mass $\sim 10^{-22}$ eV. GG-Theory places the boundary pseudoscalar at $8.1 \mu\text{eV}$ at its calibrated benchmark — seventeen orders of magnitude heavier — and contains no *second* light pseudoscalar to place in the ultralight regime. The candidate is not merely disfavored observationally (the canonical window is already closed by Lyman- α , §19); it has no field to be. *There is no fuzzy dark matter.*

Primordial black holes — no tunable primordial spectrum. Forming primordial black holes as the dominant dark matter needs a tuned, strongly enhanced small-scale primordial power spectrum. The inflationary history is fixed by GG-A10-Cosmos to the derived GUT-adjacent plateau, a smooth spectrum with no engineered small-scale enhancement and no spectral dial to turn; the reheating temperature falls far short of any restoration that could reshape the sector (Part VII, §36.2). *Primordial black holes cannot be the dominant dark matter in this cosmology.*

Dark photons and self-interacting sectors — no spare gauge factor, only a terminated one. A kinetically-mixed dark photon needs an extra $U(1)$; a tunable SIDM sector needs a free self-coupling. The bulk gauge group is forced to $SU(5) \times U(1)_B \times U(1)_{Q_i}$, and GG-A2-GG6 forbids further gauge factors — there is no further gauge factor to mix with — the tensor’s decoupled vector zero mode that GG-A3-BI6 registers is a massless vector whose loop-level mixing with hypercharge is an open exposure, not a dial — and the only dark–baryon physics the theory permits runs through the massive gauged-baryon sector at M_{lock} — parameter-free, and with no channel a laboratory condensate can drive. *There is no dark photon with a dial and no dial-equipped self-interacting sector; the one massless vector the parents register has open mixing and is carried as such.* Table 10 compresses the verdict clause by clause: each rival, the constitutional article that forbids it, and the measurement now cornering it independently — the double closure, structural and empirical at once, that makes the Kaxion the last candidate standing *within this theory*.

Table 10: The constitutional verdict: what each rival candidate *requires*, and the GG-Theory clause that forbids it. The exclusions are structural (conditional on the framework), not model-independent experimental exclusions; the framework is itself falsifiable at 1.95 GHz.

Candidate	Required new ingredient	GG-Theory clause that forbids it
WIMP	stable weak-scale multiplet / superpartner	MSMF (no new fields/dials); fixed matter content (GG-A9-SM); no supersymmetry
PQ axion (KSVZ/DFSZ)	global $U(1)_{\text{PQ}}$ + complex PQ scalar	strong CP already solved by the Kaxion (GG-A1-Constitution P7); no radial/PQ scalar exists (gauge-topological angle)
Sterile ν	keV gauge-singlet fermion + tuned mixing	lepton/neutrino sector fixed (GG-A9-SM); MSMF; warm relic vs. cold saturating Kaxion
Fuzzy DM	$f \sim M_{\text{Pl}}, m \sim 10^{-22} \text{ eV}$ pseudoscalar	forced pseudoscalar placed at $8.1 \mu\text{eV}$ (calibrated benchmark); no second light scalar
PBH (dominant)	tuned, enhanced primordial spectrum	GUT-adjacent inflation fixed (GG-A10-Cosmos); no spectral dial
Dark photon / SIDM	extra $U(1)$ / tuned self-coupling	gauge group forced $SU(5) \times U(1)_B \times U(1)_{Q_i}$ (GG-A2-GG6); no light dark–baryon channel
Kaxion	none — already forced	required by BI-6; saturates the π-lock at the calibrated benchmark

20.3 The shape of the claim, stated honestly

These are exclusions *within* GG-Theory — each rests on a constitutional clause (Maximum Symmetry/Minimum Freedom, the forced gauge group, the fixed matter content, the strong-CP assignment, the fixed inflationary sector), not on a model-independent experimental bound. If the framework is wrong, the exclusions lapse. But the framework does not hide behind that caveat: it is falsifiable, and sharply. A null across the window retires the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d; the Kaxion is then owed at $\gtrsim 6 \text{ GHz}$ on the unpinned

branch or the relic is not this sector’s. A second line inside the window refutes the count and names a registered mode. Neither is a kill of the framework, and this paper does not claim one. A confirmed compact-angle-like line at any other frequency retires the same objects as the null, or, at $\gtrsim 6$ GHz, the pinning conjecture (Parts VI–VIII). And it offers the converse hostages: a confirmed thermal WIMP, a Peccei–Quinn axion off the closed window, a keV sterile-neutrino dark matter, or a primordial-black-hole-dominated halo would each refute GG-Theory’s dark-sector closure. The theory does not merely predict the Kaxion; it stakes a falsifiable claim against the rest of the field, testable from both sides.

20.4 The message to the community

For forty years dark matter has been treated as an open slot to be filled by whichever particle was theoretically fashionable, and the search has been organized as a sweep through that particle’s parameter space. GG-Theory reframes the problem. In a framework that has already fixed the gauge group, the matter content, the cosmological history, and the resolution of strong CP, dark matter is not a slot to be filled but a *consequence* to be read off — forced, unique, and present before the question is asked. The practical content is a redirection, and we state it as plainly as the evidence allows: the dark matter is a μeV pseudoscalar at 1.95 GHz — a single line in a fenced window — and the multi-decade WIMP, sterile-neutrino, primordial-black-hole, and fuzzy-dark-matter programs are — if the framework holds — searching where the answer is not. We offer this not as a rhetorical flourish but as a falsifiable, two-sided proposition, stated plainly for the community to test: the same haloscope campaign that would confirm the Kaxion would, by a clean null, retire the calibrated placement and the ladder that fences it, not the framework — and send the search to $\gtrsim 6$ GHz.

21 Astrophysical significance I: early, massive structure in the JWST era

Beyond the laboratory, the sharpest small-scale puzzle of galaxy formation now in play is the early appearance of massive, luminous structure revealed by JWST. A candidate with fixed properties ought to say something definite about it; the Kaxion does, quantitatively, and we develop the case in full here. The complementary small-scale puzzle — the cusp–core problem and the diversity of rotation curves, where the Kaxion’s locked prediction is weaker and its one candidate lever is honestly magnitude-pending — is the subject of Section 22.

21.1 The puzzle, in brief

Deep JWST imaging has revealed bright, apparently massive galaxy candidates at $z \gtrsim 10$, assembling stellar mass earlier and faster than the most straightforward ΛCDM expectations (Boylan-Kolchin, 2023). The remainder of this section makes the surprise quantitative and shows that its *direction* — a demand for *more*, not less, early small-scale structure — selects cold dark matter and, among the light-boson candidates, the Kaxion uniquely.

21.2 The JWST surprise, quantified

Puzzle (i) deserves the detail the rest of this Part gives the laboratory, because it is the sharpest astrophysical lever the dark sector has, and because its *direction* discriminates among candidates. We state the surprise without Kaxion first, then read it through the framework.

Deep imaging with JWST has, since 2022, found luminous, ultraviolet-bright galaxies far earlier than the pre-launch consensus expected. The spectroscopic record has moved past JADES-GS-z14-0 at $z = 14.32$ (to $z \simeq 14.4$ as of 2025), and that galaxy remains the well-characterized case: it sits at $z = 14.32$ — under 300 Myr after the Big Bang — with $\sim 5 \times 10^8 M_\odot$ already in stars, a compact ~ 1.6 kly body, and a bright $M_{\text{UV}} \simeq -20.8$ (Carniani et al., 2024). Behind that record stands a population: photometric and spectroscopic samples of bright galaxies at $z \simeq 10$ –14, and massive candidates at $z \simeq 7$ –10 with inferred stellar masses $M_\star \sim 10^{10}$ – $10^{11} M_\odot$ (Labbé et al., 2023).

The tension with the straightforward Λ CDM expectation is quantitative and large. The inferred cumulative stellar-mass density runs roughly $20\times$ above standard expectations at $z \simeq 8$ and as much as $\sim 10^3\times$ at $z \simeq 9$ (Boylan-Kolchin, 2023); the rest-ultraviolet luminosity function and the cosmic star-formation-rate density at $z \simeq 12$ –14 exceed constant-efficiency models by factors of 4–10. The most extreme massive candidates, taken at face value, demand a baryon-to-stellar conversion efficiency $\epsilon_\star = M_\star/(f_b M_{\text{halo}})$ approaching — or formally exceeding — unity, against the $\epsilon_\star \sim 0.05$ – 0.1 of established galaxy formation: the “impossible early galaxies.” The arithmetic of the surprise is the exponential rare-halo tail of the Press–Schechter (Press & Schechter, 1974) mass function,

$$n(>M, z) \propto \text{erfc}\left(\frac{\delta_c}{\sqrt{2}\sigma(M, z)}\right), \quad n(M, z) \sim e^{-\nu^2/2}, \quad \nu \equiv \frac{\delta_c}{\sigma(M, z)}, \quad (21.1)$$

with $\delta_c \simeq 1.686$. At high redshift $\sigma(M, z)$ is small, so ν is large and massive halos are exponentially rare; producing $10^{11} M_\odot$ halos at $z \simeq 10$ therefore requires either a much higher star-formation efficiency or *more, earlier* massive halos than Λ CDM delivers.

21.3 Status in 2026: attenuated, not dissolved

Honesty requires recording that the tension has softened since its 2022–2023 peak. It is now substantially relieved by baryonic astrophysics — feedback-free and compact starbursts that raise $\epsilon_\star$ at high redshift (Dekel et al., 2023), bursty star formation and ultraviolet variability that lift the bright end, active-nucleus “little red dot” contamination inflating photometric masses — and by observational systematics: revised, lower masses from mid-infrared photometry, and cosmic variance in the small ($\sim$ tens of arcmin²) early fields. The consensus of 2025–2026 is that Λ CDM is *not* broken. What remains is nonetheless a real, sharp residual: the bright-end ultraviolet luminosity function at $z \gtrsim 12$ –14 and the star-formation-rate-density excess persist, and JADES-GS-z14-0 is spectroscopically secure. The surprise is best read today as a demand for a *head start* — a modest advance in the first-light chronology and efficiency — and, decisively for this paper, as a sharp discriminant *against* dark matter that suppresses small-scale power.

21.4 Direction matters: the surprise excludes the suppressing candidates

Structure at a given epoch is built from the small-scale matter power spectrum, and the candidates that ease the cusp–core problem do so precisely by *removing* small-scale power — which makes early structure *later* and *rarer*. Warm dark matter free-streams; fuzzy dark matter exerts a galaxy-scale quantum pressure. Both cut off the transfer function near the galaxy-halo scale and suppress the very high- ν tail that JWST has populated. The data have therefore turned into a one-sided constraint: the high-redshift ultraviolet luminosity function now sets *lower* bounds of $m_{\text{FDM}} \gtrsim 5.6 \times 10^{-22}$ eV on fuzzy dark matter and $m_{\text{WDM}} \gtrsim 1.5$ keV (and $\gtrsim 3.2$ keV in stricter analyses) on warm dark matter (JWST dark-matter bounds, 2024). The canonical fuzzy candidate at 10^{-22} eV, whose half-mode mass $M_{1/2} \simeq 1.6 \times 10^{10} (m/10^{-22} \text{ eV})^{-4/3} M_{\odot}$ lands right on the JWST galaxy-halo scale, is disfavored outright. The surprise, read as dark-matter physics, is thus a vote *for* cold and *against* fuzzy and warm — exactly the candidates GG-Theory independently excludes (§20).

21.5 The Kaxion clears the necessary condition by sixteen orders of magnitude

The first thing the Kaxion must do to accommodate the surprise is the thing fuzzy and warm dark matter cannot: leave the small-scale spectrum intact. It does, by an enormous margin. By the coldness results of Part III, the Kaxion’s de Broglie wavelength in a halo is $\lambda_{\text{dB}} \simeq 154$ m, Eq. (16.5), its free-streaming length is $\sim 10^{-6}$ pc and its Jeans length $\sim 3 \times 10^{-4}$ pc, its primordial velocities $\sim 4 \times 10^{-14}$ at equality; the corresponding wave-pressure (half-mode) cutoff lies near $10^{-13} M_{\odot}$, more than twenty orders of magnitude below the galaxy scale. In the terms of the bound above: the Kaxion mass $8.1 \mu\text{eV} = 8.1 \times 10^{16} \times (10^{-22} \text{ eV})$ clears the JWST fuzzy-dark-matter floor by some *seventeen* orders of magnitude, and the warm bound does not apply at all (the Kaxion is a non-thermal, pressureless misalignment condensate from $t \simeq 0.3 \mu\text{s}$, not a free-streaming thermal relic). The Kaxion therefore preserves the *full* ΛCDM small-scale power: every early halo ΛCDM forms, it forms. Among the light-boson dark-matter candidates it is the unique one that pays *no* early-structure penalty — the necessary condition, met with room to spare.

21.6 The head start on offer: what the cosmology companion still carries

Coldness alone only *recovers* ΛCDM — which is itself mildly tense at the bright end. Earlier revisions of this paper listed three further levers from the companion cosmology paper; the present state of that paper (Arisaka, 2026, A10-Cosmos) leaves one standing, and this paper carries what its parent carries and no more. The one lever that stands is a reduced effective collapse threshold for the relevant early halos, a baryonic-environment effect whose amplitude is a working value: $\delta_c^{\text{eff}} \simeq 1.43\text{--}1.33$ for the population average, corresponding by GG-A10-Cosmos’s chronology map to $\Delta z \simeq 2\text{--}3$ at $z \simeq 10$ (the full $\delta_c^{\text{eff}} = 1.18$, which the tail weighting $n_{\text{GG}}/n_{\Lambda\text{CDM}} \sim \exp[(\delta_c^2 - \delta_c^{\text{eff}2})/2\sigma^2]$ turns into a factor ~ 500 at $\sigma \simeq 0.34$, is claimed there for the rarest halos only and gives $\Delta z \simeq 4.7$). GG-A10-Cosmos grades it “stands as stated; baryonic and environmental.” The other two levers earlier revisions carried are gone with their parent: the early-time small-scale power boost ($\Delta P/P \simeq 3\%$) is withdrawn there — its derivation retracted, and its residue moving the rare tail the wrong way

— and the baryon–dark-matter streaming-velocity damping is unpriced there, every rectification channel failing by many orders; and under the boundary-condition termination this paper installs it has no mediator at all, the lightest baryonic vector being Z'_B at 3.85 TeV with a range of 5×10^{-20} m. The head start, where it exists, is given to the first seeds, not to the bulk, and leaves broad linear growth (hence S_8) essentially untouched.

21.7 Honest fences on the early-structure claim

Four fences keep the claim exact. (a) Much of the surviving tension is plausibly baryonic (high-efficiency early star formation, §21.3); the GG head start *combines* with that astrophysics rather than replacing it. (b) The one surviving figure, δ_c^{eff} , is GG-A10-Cosmos’s working value, not a first-principles number; it carries that paper’s kill criteria and awaits full hydrodynamical calibration, and the two levers GG-A10-Cosmos has withdrawn or left unpriced are not claimed here. (c) The head start here is the *single-patch ledger* effect; it is *not* the minicluster lever of §21.10 below, which belongs to the post-inflationary branch the series excludes (it would move the mass out of the 1.95 GHz window). (d) If the observational tension dissolves entirely into systematics, the Kaxion simply predicts standard cold Λ CDM early structure — still on the favorable side, still paying no penalty, and the layers above become a falsifiable bonus rather than a necessity.

21.8 Where the Kaxion stands on early structure

The early-structure verdict mirrors the candidate verdict of §20, and that mirroring is the content. The Kaxion (i) satisfies, by sixteen orders of magnitude, the very JWST bounds that exclude fuzzy and warm dark matter — the necessary coldness; and (ii) inherits, at most, the modest collapse-threshold head start that GG-A10-Cosmos carries as a working value — a possible boost, not a sufficient one. The candidates the JWST surprise disfavors are exactly the ones GG-Theory structurally excludes; the candidate GG-Theory forces is exactly the one the surprise favors. The observational arrow and the structural arrow point the same way — which is the strongest thing an astrophysical argument can do for a laboratory prediction.

21.9 The not-fuzzy fence, and the one early-structure contingency

Everything above stands behind the not-fuzzy fence of Part III: with $\lambda_{\text{dB}} \simeq 154$ m (Eq. (16.5)), the Kaxion’s wave nature operates seventeen orders of magnitude below galactic scales. It reproduces every Λ CDM success on linear and mildly nonlinear scales, and it does *not* solve small-scale problems by wave pressure. Its early-structure leverage therefore comes not from wave physics but from the inherited cosmological head start already developed in §21.6 and, as a separate contingency, from the miniclusters below.

21.10 Lever one: minicluster seeds (post-inflationary contingency)

If the compact-angle field is randomized across causal patches (the *post*-inflationary history), its order-unity inhomogeneities collapse at matter–radiation equality into dense miniclusters (Kolb & Tkachev, 1993, 1994) — small-scale power *added* to the inflationary spectrum, going nonlinear early.

Directionally, this places the Kaxion on the *favorable* side of the JWST tension, opposite to warm and fuzzy candidates.

With $N_{\text{DW}} = 1$ in the period convention Theorem 24.1 adopts, the scenario is cosmologically safe — no domain-wall catastrophe; in the alternative convention ($N_{\text{DW}} = 5$, open) the randomized history would carry stable walls and be excluded outright rather than relocated. The fence: the minimal reading is the *pre*-inflationary single-patch history, in which one $\mathcal{O}(1)$ angle is homogenized across the observable universe (read from the abundance as $\theta_i \simeq 1.1$ rad, Part III) and this lever is absent. The post-inflationary variant would instead replace θ_i^2 by the circle average $\langle \theta^2 \rangle = \pi^2/3$, which at the benchmark decay constant overproduces the relic by a factor of three to four (full cosine; 2.6 harmonic), and which closes the inventory only at $f_K \lesssim 2.3 \times 10^{11}$ GeV — $\gtrsim 25 \mu\text{eV}$, a line at $\gtrsim 6$ GHz from misalignment alone, outside the window, and higher still with any defect contribution — while adding minicluster substructure.

Within the series this variant is excluded conditionally: the Kaxion has no Peccei–Quinn scalar to restore thermally, and the post-inflationary reheating temperature, 6×10^{10} GeV, lies a factor $\simeq 12$ below f_K ; but de Sitter randomization during primary inflation is foreclosed only by the inflation-era pinning the paper names as Conjecture C-Kaxion-d (Section 13.5; Part VII, Section 36.2). The minicluster lever is therefore registered as the contingency that a failure of C-Kaxion-d would make live — and, since it would also move the line out of the window, as a falsifier of the pinning rather than a live prediction of the theory.

21.11 Closing: where the Kaxion stands on early structure

The early-structure verdict is the strong one. Warm and fuzzy candidates *worsen* the JWST tension (they suppress the very small-scale power the data demand) and are now bounded away by it; the Kaxion, seventeen orders heavier, preserves the full ΛCDM budget and inherits, at most, the modest collapse-threshold head start its cosmology companion carries as a baryonic-environment working value ($\Delta z \sim 2\text{--}3$, graded qualitative there). Astrophysics will not identify the Kaxion; the laboratory will (Parts VI–VIII). But on early structure the astrophysical and structural arrows point the same way, which is the most an astrophysical argument can do for a laboratory prediction.

22 Astrophysical significance II: the cusp–core problem and the diversity of rotation curves

The second standing small-scale puzzle is the cusp–core problem and its modern, sharper form, the diversity of rotation curves. Honesty requires a different posture here than in Section 21. On early structure the Kaxion’s contribution was quantitatively strong and theorem-anchored; on cusp–core it is weaker, and its one candidate lever is genuinely *magnitude-pending*. We develop the problem in full, state precisely what the Kaxion can and cannot do, and mark throughout what is theorem-grade and what is an open computation.

22.1 The problem, quantified

Dark-matter-only simulations of cold dark matter assemble halos with centrally divergent “cuspy” density profiles of Navarro–Frenk–White form (Navarro, Frenk, & White, 1997), $\rho \propto r^{-1}$ (inner logarithmic slope $\gamma \rightarrow 1$) inside the scale radius. Yet the rotation curves and stellar kinematics of many dwarf and low-surface-brightness galaxies prefer central “cores” — $\rho \rightarrow \text{const}$, $\gamma \rightarrow 0$, over $\sim 0.5\text{--}2$ kpc — with inner densities a factor of a few to ten below the NFW expectation (Bullock & Boylan-Kolchin, 2017).

The modern, sharper statement is the *diversity* of rotation curves (Oman et al., 2015). At a fixed outer (peak) circular velocity $V_{\text{max}} \sim 30\text{--}100 \text{ km s}^{-1}$, the circular velocity measured at 2 kpc — a clean proxy for the central dark-matter density — spans a factor of about *four*, from NFW-cuspy systems to strongly cored ones. Collisionless Λ CDM predicts a *tight* $V_{2\text{kpc}}\text{--}V_{\text{max}}$ relation and cannot reproduce this scatter. The puzzle is therefore not “cores versus cusps” in the abstract but the *coexistence* of both at the same halo mass: a successful theory must produce cores and cusps and a continuum between them, correlated with the right galaxy property.

22.2 The standard solutions and their limits

Baryonic feedback, and its energy floor. Repeated supernova-driven gas outflows make the central gravitational potential fluctuate; the time-dependent potential transfers energy to the dark matter and can flatten a cusp into a core. This works well for gas-rich dwarfs with $M_{\star} \sim 10^7\text{--}10^9 M_{\odot}$ (the FIRE and NIHAO simulations). But it has a hard floor set by energetics: turning a cusp into a kiloparsec core requires injecting an energy comparable to the central binding energy, of order $10^{53}\text{--}10^{54}$ erg for a classical dwarf, and the integrated supernova budget falls below this when $M_{\star} \lesssim 10^6 M_{\odot}$ (Bolchini, 2022). The gas-poor classical dwarf spheroidals and the ultra-faint dwarfs — the most dark-matter-dominated systems, with $M_{\star} \ll 10^6 M_{\odot}$ — are therefore expected to retain cusps, yet cores are inferred in several. *Feedback fails exactly where the problem is cleanest.*

Self-interacting dark matter, and its dial. Elastic dark-matter self-scattering thermalizes the inner halo into an isothermal core whose radius is set by where a particle has scattered roughly once in a Hubble time. Fitting dwarf and spiral rotation curves requires a cross-section $\sigma/m \sim 0.5\text{--}5 \text{ cm}^2 \text{ g}^{-1}$ at dwarf velocities ($v \sim 10\text{--}50 \text{ km s}^{-1}$), *velocity-dependent* so as to fall to $\lesssim 0.1\text{--}1 \text{ cm}^2 \text{ g}^{-1}$ at cluster velocities ($v \sim 10^3 \text{ km s}^{-1}$) and respect cluster and merging-system bounds (Kaplinghat et al., 2016; Tulin & Yu, 2018). SIDM is the leading particle solution to the *diversity* precisely because its core size depends on two things at once: the baryon distribution (more baryons deepen the well and shrink the core) and the halo’s gravothermal phase — early core expansion followed by late *core-collapse*, in which the inner halo re-contracts to high density. That two-parameter behavior reproduces a cusp-to-core spread comparable to the observed one (Fischer et al., 2024). The price is a tuned, velocity-dependent function $\sigma(v)/m$ — a curve, not a number.

Fuzzy dark matter, now disfavored. Wave pressure at $m \sim 10^{-22} \text{ eV}$ produces a solitonic core; but the same suppression is excluded by the JWST and Lyman- α bounds of Section 21. The wave-pressure core mechanism is closing as a viable option.

Degeneracy. In gas-rich systems feedback and SIDM are observationally degenerate as core-forming mechanisms; the discriminating regime is the feedback-starved one — the gas-poor and ultra-faint dwarfs.

22.3 The Kaxion is cold, and is not fuzzy

The first two statements about the Kaxion are honest concessions. *It is cold:* its dark-matter-only halos are NFW-cuspy, exactly like Λ CDM, so the Kaxion *inherits* the cusp-core problem rather than dissolving it by exotic kinematics. *It is not fuzzy:* with $\lambda_{\text{dB}} \simeq 154$ m in a halo (Section 21), fifteen-plus orders below the kiloparsec, the Kaxion forms no solitonic, wave-pressure core. The coldness that made the Kaxion win on early structure forecloses the fuzzy core mechanism here. The Kaxion does not get to borrow the wave-DM success.

22.4 The Kaxion is not self-interacting dark matter

A second concession, equally important to state plainly: the Kaxion is *not* SIDM. It is neutral under $U(1)_B$ — it couples to baryon number through the anomaly term $a F_B \tilde{F}_B$, not as a charge — so there is no Kaxion–Kaxion gauge force, and its cosine self-coupling is gravitationally negligible at $m_K = 8.1 \mu\text{eV}$. The Kaxion therefore has *no* dark-matter–dark-matter elastic-scattering channel and does not inherit SIDM’s gravothermal core mechanism. We will not borrow that success either.

22.5 The one candidate lever, and why the termination closes it

What the Kaxion’s ledger *does* carry is a baryonic entry: the gauged- $U(1)_B$ coefficient of X_4 (Section 27) — though the termination leaves it no channel below M_{lock} , so whatever lever it could supply is bounded by the massive sector’s reach. A nonzero baryonic coupling means dark-matter centers and their baryons are not perfectly decoupled: energy can be exchanged between the two fluids in the densest regions — a dark-matter–baryon interaction core mechanism of the effective-theory class (Boddy & Gluscevic, 2018), distinct from *both* feedback (purely baryonic) *and* SIDM (dark–dark). Two structural virtues follow *if* the magnitude is sufficient.

1. **Feedback-independence.** The portal acts wherever dark matter and baryons coexist at high density — including the gas-poor and ultra-faint regime where the supernova budget fails. It is a candidate precisely for the residual that defeats feedback, the cleanest part of the problem.
2. **No new dial.** Its strength relative to the photon channel is parameter-free (the ledger’s $-2/5$ entry). Unlike SIDM, which must fit a velocity-dependent function $\sigma(v)/m$, the Kaxion portal introduces only a single overall magnitude — and that magnitude is tied to a laboratory-measurable coupling (below).

22.6 The magnitude, now computed: the portal does not core

The threshold the portal must clear is set by the same energetics that floor the feedback mechanism: coring a classical dwarf in a Hubble time requires an effective dark-matter–baryon cross-section $\sigma_{\chi b}^{\text{thr}} \simeq 2 \times 10^{-25} \text{ cm}^2$ per nucleon (equivalently a SIDM-scale $\sigma/m \sim \mathcal{O}(1) \text{ cm}^2 \text{ g}^{-1}$ of energy

transfer). What this paper pre-registered as an open computation is now done (§73.4): the gauged- $U(1)_B$ coupling is $1/f_K$ -suppressed ($f_K \simeq 7.1 \times 10^{11}$ GeV) and the induced scattering momentum-suppressed besides (Boddy & Gluscevic, 2018), giving $\sigma_{\chi b} \sim m_K^2 v^2 / \pi f_K^4 \simeq 3 \times 10^{-110} \text{ cm}^2$; even an extreme coherent enhancement by the full de Broglie occupation ($N \simeq 2.0 \times 10^{26}$ at $\rho_{\text{loc}} = 0.45 \text{ GeV cm}^{-3}$, squared) reaches only $\lesssim 10^{-57} \text{ cm}^2$. The portal therefore falls *thirty-two to eighty-five orders of magnitude* short of the threshold. *The dark-baryon portal does not core dwarfs.* Under the termination the conclusion only hardens: with no channel below M_{lock} the coring magnitude is bounded yet further below the computed number, and that number was already negligible. The result is a clean, definite prediction in place of an open hedge: the Kaxion forms standard cuspy cold-dark-matter halos, and the cusp-core problem is, for the Kaxion, baryonic.

22.7 What the Kaxion predicts for galaxy centers

With the portal computed inert (§22.6), the Kaxion’s prediction for galaxy centers is definite and falsifiably distinct from the self-interacting alternative.

- **Cuspy halos, baryonic cores.** The Kaxion makes standard NFW-cuspy dark-matter halos; whatever cores are real are carved by baryonic feedback, exactly as in collisionless Λ CDM. The Kaxion neither eases nor worsens cusp-core.
- **No gravothermal core-collapse.** Having no dark-dark scattering, the Kaxion does not thermalize the inner halo and produces *none* of SIDM’s late-time re-cuspatation. A confirmed population of *core-collapsed* dwarfs (high-density re-cuspatation) would favor SIDM over the Kaxion — a clean discriminator that survives the magnitude computation.
- **A real ledger entry, read only at the collider.** The baryonic entry is genuine and ledger-locked, but it is astrophysically silent and — under the termination — silent in the haloscope too: its exposure is the massive sector’s, at 3.85 TeV (Arisaka, 2026, P1-KK). The galaxy is where the entry *fails* to act; the collider is where it is hunted.

22.8 Where the Kaxion stands on cusp-core

This is the one place the Kaxion’s astrophysical story is weakest, and the computation of §73.4 now makes the position definite rather than pending. The Kaxion (i) does not *worsen* cusp-core (it is cold, not fuzzy or warm); (ii) is *not* SIDM (it has no dark-dark lever); and (iii) does *not* supply a dark core lever either — its one candidate channel, the dark-baryon portal, was computed thirty-two to eighty-five orders of magnitude too weak to core a dwarf even before the termination closed it below M_{lock} . The Kaxion therefore predicts standard cuspy cold dark matter, leaving the cores to baryonic feedback. This is not a weakness to apologize for but a *closed prediction*: galaxy centers are cuspy, the observed cores are baryonic, and the rotation-curve diversity is astrophysical. Decisively, the cusp-core problem was never a requirement for the Kaxion to be the dark matter — the abundance, the coldness, the lifetime, and the two laboratory fingerprints all stand regardless of how galaxy centers turn out. The dark-baryon coupling that might have cored them is real and locked, but it is too feeble to move a galaxy and shows up only on resonance in the laboratory.

Which is the through-line of the whole paper: it is the laboratory, not the galaxy center, that identifies the Kaxion (Parts VI–VIII).

22.9 What Part IV has established

This Part placed the Kaxion in the field of dark-matter candidates and drew the strongest claim the framework supports. Four results stand.

1. **The field is empty of forced candidates (§19).** Every leading rival is being squeezed by null results — the WIMP between the emerging neutrino fog and the LHC’s exclusion of weak-scale supersymmetry, the sterile neutrino and fuzzy dark matter bounded away by X-ray and Lyman- α data, primordial black holes confined to the asteroid-mass window — and not one of them is *forced*, predicts its own mass, or derives its own abundance. The field has returned to its widest question in fifty years: not where dark matter is, but what it is.
2. **Within GG-Theory the Kaxion is the unique answer (§20).** It is *required* — forced by the Green–Schwarz Bianchi identity (Theorem T-Kaxion-1), present before the dark-matter question is posed — and *sufficient* — the π -locked inventory is saturated at the calibrated benchmark with a generic angle. Every rival is excluded by a constitutional clause adopted for other reasons: no supersymmetry for the WIMP, no Peccei–Quinn scalar and no residual strong-CP problem for the axion, a fixed lepton sector against the sterile neutrino, a forced gauge group against the dark photon, a fixed inflationary sector against primordial-black-hole dominance. The candidates the data disfavor are exactly the ones the constitution forbids.
3. **On early structure the Kaxion is uniquely favored (§21).** The JWST surprise demands *more* early small-scale structure, which bounds away the suppressing candidates (fuzzy and warm dark matter); the Kaxion clears those bounds by sixteen orders of magnitude and inherits, at most, the modest collapse-threshold head start its cosmology companion carries as a working value. The observational and structural arrows point the same way.
4. **On cusp–core the Kaxion makes a closed, computed prediction (§22).** The one candidate dark lever, the dark–baryon portal, was computed thirty-two to eighty-five orders of magnitude too weak to core a dwarf even before the termination closed it (§73.4); the Kaxion therefore predicts standard cuspy cold-dark-matter halos, with cores baryonic and no gravothermal core-collapse. Cusp–core is an opportunity the Kaxion declines, not a test it must pass.

The through-line of the Part is one sentence: astrophysics *constrains and contextualizes* the Kaxion — favorably on early structure, neutrally on cusp–core — but it does not *identify* it. That the remaining Parts hand to the laboratory.

23 Part IV in plain English: how the candidate compares, and what it would take to tell them apart

It is easy to write down a dark-matter candidate and hard to write down one that could be wrong. This Part is about the difference.

The scorecard asks structural questions rather than phenomenological ones. Is the candidate's existence forced by something adopted for other reasons? Is its mass scale forced, or chosen? Is its initial condition forced, or chosen? Does it sit on a line in the coupling-mass plane, or anywhere in a plane? A candidate can be excellent physics and answer no to all four — several on the list are excellent physics — and the answers say nothing about which candidate nature chose. They say how much of each candidate has to be supplied by hand before it can be compared with data.

That matters for a practical reason. A candidate occupying a plane can absorb a null result by moving; a candidate at a point cannot. The Kaxion is at a point, and the rest of the paper spends that fact: the frequency is stated, the depth required is stated, and the experiment that reaches both either sees it or ends it.

The astrophysical reach in the second half of the Part is the other side of the same coin. A light pseudoscalar with a photon coupling has consequences in stars and in magnetized environments, and the paper checks them because a candidate that survives the laboratory and fails in a star is dead anyway.

Part V

Derivation of the basic quantities

Mass, photon coupling, abundance, and lifetime — from the series theorems

This Part derives the Kaxion’s basic quantities — mass, photon and baryonic couplings, abundance, and lifetime — from the series theorems and one declared calibration, with no free parameter after it. Figure 7 previews the whole chain on a single page: from the constitution, through the bulk geometry, the integer ledger, and the runtime, down to the four falsifiable numbers. The sections that follow fill in each arrow in turn.

24 Strong-CP alignment and the legitimacy of the QCD susceptibility

24.1 The constitutional assignment

The constitution does not leave the strong-CP problem to taste. Problem P7 of GG-A1-Constitution (Arisaka, 2026, A1-Constitution) declares it *Resolved* within the series, and the series carries one mechanism, stated here at its grade. The primary statement is GG-A9-SM’s (Arisaka, 2026, A9-SM): the strong-CP angle is the protected zero of the topological ledger, suppressed by the scale profile, and GG-A9-SM protects it “geometrically rather than by an additional axion-like field inserted ad hoc” — graded there derived-conditional, bare angle only, on its lemma L9 — the sign, and the shallow-winding leakage beyond first order. The Kaxion adds to that structural protection a dynamical relaxation: its gluonic portal makes whatever residual the projection leaves roll to the CP-conserving point. The two are one story at GG-A9-SM’s grade, and this paper’s relaxation inherits that grade rather than raising it; the window and the single-line count are the constitutional falsifiers. This Part begins by showing that the relaxation is legitimate — that the Kaxion really does couple to the gluon topological charge with the right strength — because everything downstream (the mass, and through the mass every frequency in this paper) hangs on it.

24.2 The gluonic portal theorem

Theorem 24.1 (T-Kaxion-2: the gluonic portal). *Under the hypotheses of Theorem 8.2, the boundary effective Lagrangian of the Kaxion contains the gluonic operator*

$$\mathcal{L} \supset \frac{\alpha_s}{8\pi} \frac{a_K}{f_K} G_{\mu\nu}^a \tilde{G}^{a\mu\nu}, \quad (24.1)$$




Figure 7: Derivation flow of the Kaxion: from the constitution to four falsifiable numbers. Each downward step adds no free parameter beyond the one calibrated scale. The constitution (GG-A1-Constitution) forces the six-dimensional bulk and the BI-6 anomaly (GG-A2-GG6, GG-A3-BI6), which carry the five-integer OGN-5 ledger; the GDOS boundary law and the GRIP runtime (GG-A5-GDOS, GG-A6-GRIP, GG-A11-Origin) select and stabilize the deepest vacuum — the Ground Geometric Attractor — and with it the Kaxion (Theorem 8.2). The four key quantities then follow from named upstream ingredients: the **mass** from the π -locked abundance (GG-A10-Cosmos) and the gluonic portal $m_K f_K = \sqrt{\chi_{\text{QCD}}}$; the **photon coupling** from $E/N = 8/3$ on the DFSZ line; the **window** from the throat depth $K = 35$ through the discrete cap ladder; and the **line count** — single — from the mode analysis of the anomaly companion. The dashed link records that the same decay constant f_K that fixes the mass also sets the photon coupling.

with unit color-anomaly coefficient ($N = 1$) in the period convention adopted here, and consequently a domain-wall number $N_{\text{DW}} = 1$ in that convention (the descent with the anomaly companion's κ and $t = 5$ made explicit is open and could give $N_{\text{DW}} = 5$).

Proof. The Chern–Simons-improved field strength (6.13) of GG-A3-B16 contains the $\text{SU}(5)$ Chern–Simons form $\omega_3(A_5)$ with ledger coefficient $b = -1$. Under the boundary reduction of GG-A5-GDOS (T6) and the dualization (9.2), the inflow term (6.15) descends — by the same anomaly-inflow package (T9.7) that transports every ledger entry unchanged — to the boundary operator $(a_K/f_K) \text{Tr } F_5 \tilde{F}_5$ with unit weight in the units in which (6.15) is written with coefficient one; the anomaly companion's integrality datum $t = 5$ for the class $\hat{X} = (t/8\pi^2)X_4$ and its normalization constant κ are absorbed into the definition of f_K as the period-normalized decay constant of the 2π -periodic $B_{r\theta}$ zero mode, and the descent with both made explicit is registered as an open bookkeeping item in Part XI (the integer could be $t = 5$ in a different period convention, in which case $N_{\text{DW}} = 5$ with the single-patch clause making walls moot either way). The color group $\text{SU}(3)_c$ is embedded in $\text{SU}(5)$ as an unbroken block (GG-A2-GG6 T3; the Hosotani breaking leaves $\text{SU}(3)_c$ intact), and the trace over the fundamental restricts to the color block with index one. Equation (24.1) follows with $N = |b| = 1$, and $N_{\text{DW}} = N = 1$ by the standard counting of distinct vacua under $a_K \rightarrow a_K + 2\pi f_K$. ■

Physical picture. The Kaxion does not need to be *given* a gluon coupling; it inherits one from the very ledger entry that made the bulk anomaly-free. The same bookkeeping that forbids the theory from being inconsistent obliges the compact angle to talk to gluons. This is the precise sense in which the strong-CP solution is not an add-on here: the relaxation mechanism rides on equipment the theory was already required to carry.

24.3 Relaxation, residual, and the quality problem

Relaxation. With (24.1) in hand, the physical vacuum angle is the field combination $\bar{\theta}_{\text{eff}} = \bar{\theta} + \langle a_K \rangle / f_K$, and QCD itself supplies the potential: the vacuum energy of QCD as a function of the angle is $V(\bar{\theta}_{\text{eff}}) = \chi_{\text{QCD}} [1 - \cos \bar{\theta}_{\text{eff}}]$ to leading order, and the Vafa–Witten theorem (Vafa & Witten, 1984) — proved for vector-like gauge theories, which QCD at the relaxed point is — certifies that the minimum of the exact vacuum energy sits at $\bar{\theta}_{\text{eff}} = 0$; the relaxation that reaches it is the standard anomaly-generated potential, exhibited directly by the lattice, and needs no theorem at all. The Kaxion rolls there; CP is restored dynamically. This is the relaxation of the Peccei–Quinn mechanism (Peccei & Quinn, 1977; Weinberg, 1978; Wilczek, 1978), executed by a field the geometry forced — with the standard floor that any relaxed angle carries: the weak CP violation of the CKM phase displaces the minimum by $|\bar{\theta}_{\text{eff}}| \sim 10^{-17}$, far below the neutron bound. In the presence of the Kaxion the codebook's own $\bar{\theta}$ is absorbed into $\langle a_K \rangle$, so the residual the world sees is this relaxed minimum; GG-A9-SM's static bound below is the residual of the codebook without the Kaxion.

Residual. GG-A9-SM supplies the static companion bound: the geometric suppression of any residual angle is

$$|\bar{\theta}| \lesssim \sin\left(\frac{2\pi}{L}\right) \frac{e^{-K}}{2^J} = 0.951 \times \frac{6.31 \times 10^{-16}}{128} \simeq 4.7 \times 10^{-18}, \quad (24.2)$$

over seven orders of magnitude below the neutron-EDM bound $|\bar{\theta}| \lesssim 10^{-10}$, at GG-A9-SM’s grade (derived-conditional). With the Kaxion present the physical residual is the relaxed minimum, $|\bar{\theta}_{\text{eff}}| \sim 10^{-17}$, and GG-A9-SM’s number is the bare angle’s bound; the two are one mechanism read at two stages, and neither is printed as a prediction below the CKM floor.

Quality. A standard objection to Peccei–Quinn solutions is the *quality problem*: a global $U(1)_{\text{PQ}}$ is not respected by quantum gravity, and Planck-suppressed operators of low dimension generically displace the minimum by far more than 10^{-10} . The Kaxion is structurally protected, with one qualification. Its shift symmetry is not a global symmetry of some scalar potential; it is the gauge invariance of the two-form B_2 (Sections 9.1–9.2), and gauge redundancies cannot be violated by gravitational operators. The breakings are the topological ones the ledger itself dictates — Eq. (24.1) and, through the same ledger, the couplings to $\text{tr } R^2$ and to the Abelian entries; the gravitational entry is a topological density, not a Planck-suppressed potential, and on the intended class — an interval times a circle, with no closed two-cycle for a worldsheet instanton — is not expected to displace the minimum at the 10^{-10} level, though this is not computed; “immune” would overstate a protection that rests on the ledger’s completeness. This is the same protection enjoyed by the model-independent string axion, here with the entire structure forced rather than chosen.

24.4 Consequence: the susceptibility owns the mass

Because (24.1) carries unit coefficient, the Kaxion’s potential at temperatures below the QCD scale is exactly the QCD θ -vacuum energy with $\bar{\theta}_{\text{eff}} \rightarrow a_K/f_K$, and the curvature at the minimum gives the exact relation

$$\boxed{m_K^2 f_K^2 = \chi_{\text{QCD}}(0)} \quad \chi_{\text{QCD}}^{1/4}(0) = 75.5 \pm 0.5 \text{ MeV} \quad (24.3)$$

from the lattice (Borsanyi et al., 2016), i.e. $\sqrt{\chi_{\text{QCD}}} = (5.70 \pm 0.08) \times 10^{-3} \text{ GeV}^2$ ($\pm 1.3\%$). Equation (24.3) holds for *any* QCD-coupled axion with $N = 1$, independently of every ultraviolet detail; the precision chiral-perturbation-theory form is $m_K = 5.70(7) \mu\text{eV} (10^{12} \text{ GeV}/f_K)$ (Grilli di Cortona et al., 2016). One number remains free in (24.3): the decay constant. Closing it is the next section’s single task.

24.5 Where the susceptibility number comes from

The lattice value (24.3) is not a black box; chiral perturbation theory derives it from four measured constants, and the two-line computation is worth displaying because every per-cent of the Kaxion mass error traces here. At leading order the QCD topological susceptibility is

$$\chi_{\text{QCD}} = \frac{m_u m_d}{(m_u + m_d)^2} m_\pi^2 f_\pi^2, \quad (24.4)$$

the quark-mass prefactor expressing that topology can be rotated away entirely if any quark is massless (the massless-up-quark loophole, closed by lattice spectroscopy). With $m_u = 2.16(11) \text{ MeV}$,

$m_d = 4.67(9)$ MeV ($r \equiv m_u/m_d = 0.462$, so $r/(1+r)^2 = 0.216$), $m_\pi = 134.98$ MeV, $f_\pi = 92.1$ MeV:

$$\chi_{\text{QCD}}^{1/4} = \left[0.216 \times (134.98)^2 \times (92.1)^2 \right]^{1/4} \text{ MeV} = 76.0 \text{ MeV}, \quad (24.5)$$

within 0.7% of the direct lattice determination 75.5(5) MeV (Borsanyi et al., 2016) — the difference being the next-to-leading-order chiral corrections that the lattice resums. The propagated uncertainty is dominated by the quark-mass ratio ($\pm 3\%$ on r gives $\pm 0.3\%$ on $\chi^{1/4}$, $\pm 0.6\%$ on $\sqrt{\chi}$); the $\pm 1.3\%$ that (24.3) carries on $\sqrt{\chi}$ is the lattice’s own 75.5(5) MeV, and either way the susceptibility is the *least* of the Kaxion’s error budget. The reader should take from (24.4) the structural point: the Kaxion’s mass scale is built from m_π , f_π , and the quark masses — four numbers from the Standard Model’s best-measured sector — divided by one number from the series’ closure. Nothing else enters.

25 Calibrating the bulk normalization through the π -lock

25.1 Ground rules for this derivation

Nothing in this section is taken on authority; even the engineering formulae of GG-A10-Cosmos are re-derived from scratch. The inputs permitted below are exactly four: (i) the π -locked target $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ (GG-A10-Cosmos, derived-conditional under its closed-packet conjecture, Conjecture VIII of its codebook (Lemma VIII; the conditional theorem T-Cosmos-17); Part III); (ii) the benchmark cap loading $\Xi_{\text{cap}} = 43$ (Section 26.3), the closure below *calibrating* the bulk normalization at a generic misalignment angle and returning that angle; (iii) the lattice susceptibility, zero-temperature value and temperature exponent (Borsanyi et al., 2016); and (iv) standard Friedmann kinematics. Every constant is evaluated, every step is checkable, and the final number carries its full error budget.

25.2 The pipeline, assembled end to end

Part III derived the physics (bank-count-redeem, Eqs. (15.3)–(15.6)); here we assemble the numbers. The chain is

$$\Omega_K h^2 = m_K(0) \times \underbrace{\frac{\frac{1}{2} m_K(T_{\text{osc}}) f_K^2 \theta_i^2}{\frac{2\pi^2}{45} g_{*s}(T_{\text{osc}}) T_{\text{osc}}^3}}_{n/s \text{ at onset}} \times \underbrace{\frac{s_0}{\rho_{\text{crit}}/h^2}}_{\text{today's books}}, \quad (25.1)$$

with the cosmographic constants $s_0 = 2891.2 \text{ cm}^{-3}$ and $\rho_{\text{crit}}/h^2 = 1.0537 \times 10^{-5} \text{ GeV cm}^{-3}$, whose ratio is $2.744 \times 10^8 \text{ GeV}^{-1}$. The onset temperature solves $3H(T_{\text{osc}}) = m_K(T_{\text{osc}})$ with the lattice mass turn-on $m_K(T) = m_K(0)(T_{\text{QCD}}/T)^b$, $b = 4.08$, $T_{\text{QCD}} = 150$ MeV, and H from (12.3) with $g_* = g_{*s} = 73$ held fixed (the lattice-informed tables put g_* between 62 and 84 across 0.5–1 GeV; the integration with the tables moves the normalization by -2% , the row in the budget below). Carrying (25.1) through with the instantaneous-onset rule $m_K(T_{\text{osc}}) = 3H(T_{\text{osc}})$ — the line-by-line audit is Appendix E.1 — yields the estimate $0.090 \theta_i^2 (f_K/10^{12} \text{ GeV})^{1.165}$. That estimate undercounts: the field stays frozen while $m_K(T) \propto T^{-4.08}$ keeps rising, so the number of quanta is fixed later,

at a larger mass, than the instantaneous rule assumes. Integrating the equation of motion (12.2) directly through the onset with the same inputs (radiation domination, $g_* = g_{*s} = 73$, the power-law turn-on, the harmonic potential) gives the master normalization

$$\boxed{\Omega_K h^2 = 0.135 \theta_i^2 \left(\frac{f_K}{10^{12} \text{ GeV}} \right)^{1.165}} \quad (\text{harmonic regime, } \theta_i \lesssim 1), \quad (25.2)$$

the exponent being the derived $(b+3)/(b+2)$ of Eq. (15.6), which the integration reproduces (0.0898 at $f_K = 7.07 \times 10^{11} \text{ GeV}$, against $0.1345 \times 0.707^{1.165} = 0.0898$). Repeating the integration with tabulated $g_*(T)$, $g_{*s}(T)$ and entropy conservation moves the normalization by -2% (to 0.132; holding g_* at the extremes of the 0.5–1 GeV range, 62 or 84, moves it by up to $\pm 7\%$), and the switch-point at which the adiabatic count is read (between $T = 0.25$ and 0.20 GeV) moves it by less than 0.1% . The instantaneous estimate is retained in Appendix E.1 as the analytic scaffold, not as the number.

Equation (25.2) is not, however, the calculation of record, and this revision separates the two. That calculation integrates (12.2) on the *tabulated* $g_\rho(T)$, $g_{*s}(T)$ and $\chi_{\text{QCD}}(T)$ of Borsanyi et al. (2016) — Tables S2 and S7, cubic splines in $\log_{10}(T/\text{MeV})$ for the equation of state and a shape-preserving interpolant in $\ln T$ for the susceptibility — with the *full cosine* potential, from $T_i = 3 \text{ GeV}$ to $T_f = 0.35 \text{ GeV}$, reading the yield from the exact elliptic action of the librating pendulum. It uses no power-law turn-on and no fixed g_* , and it is what Section 15.3 and the budget of Section 29 now carry. Against it, Eq. (25.2) is high in Ω_K by some 11% at the benchmark, because its two approximations run in opposite directions and do not cancel: the fixed g_* with the power-law turn-on is high, and the harmonic force is low. Earlier revisions of this paper carried the two through separately and applied the anharmonic correction to the fixed- g_* baseline; that is the origin of the closure angle they printed, and Section 15.3 states the corrected value with the arithmetic in the open.

25.3 The pipeline in closed form

Before trusting any number, the student should see the pipeline as one closed-form expression. Solving the onset condition with the turn-on mass gives

$$T_{\text{osc}} = \left[\frac{m_K(0) T_{\text{QCD}}^b M_{\text{Pl}}}{4.98 \sqrt{g_*}} \right]^{\frac{1}{b+2}}, \quad (25.3)$$

and inserting (25.3) into (25.1) with $m_K(T_{\text{osc}}) = 3H(T_{\text{osc}})$ collapses the entire computation to

$$\Omega_K h^2 = \underbrace{\frac{45}{4\pi^2} \frac{4.98\sqrt{g_*}}{g_{*s}}}_{\text{kinematics}} \underbrace{\frac{s_0/(\rho_{\text{crit}}/h^2)}{M_{\text{Pl}}}}_{\text{today's books}} \underbrace{\frac{\sqrt{\chi_{\text{QCD}}} f_K \theta_i^2}{T_{\text{osc}}}}_{\text{the physics}}, \quad (25.4)$$

in which the three brackets cleanly separate what kind of number each factor is: pure kinematic constants; cosmographic bookkeeping; and the physics — the susceptibility (Standard Model), the decay constant (the unknown), the selected angle squared (the runtime), and one inverse power of the onset temperature carrying the stiffening-spring memory. Every exponent of the error budget

(Table 12) can be read off (25.4) with (25.3) by inspection: $\Omega \propto f_K^{1+1/(b+2)} = f_K^{1.165}$, $\Omega \propto \theta_i^2$, $\Omega \propto (\sqrt{\chi})^{1-1/(b+2)}$, $\Omega \propto T_{\text{QCD}}^{-b/(b+2)}$, $\Omega \propto g_*^{1/2+1/(2(b+2))}/g_{*s}$. Evaluating (25.4) at the reference point $f_K = 10^{12}$ GeV, $\theta_i = 1$ gives the instantaneous-onset estimate 0.090; the integrated normalization (25.2) is larger by the factor 1.49 that the frozen-field transient supplies, and the line-by-line numerics of the estimate are Appendix E.1.

25.4 Cross-validation against the precision literature

Before using (25.2) we test it. Setting $\theta_i = 1$ and $\Omega_K h^2 = 0.12$ gives $f_K = 0.90 \times 10^{12}$ GeV, i.e. $m_K = 6.3 \mu\text{eV}$. The leading-order community formula, $\Omega h^2 \simeq 0.12 (f_K/9 \times 10^{11} \text{ GeV})^{1.165} \theta_i^2$, is the same statement; the precision computations with NLO susceptibility and full $g_*(T)$ transport (Gorghetto & Villadoro, 2019; Grilli di Cortona et al., 2016) sit within roughly $\pm 10\%$ of it in mass once their $\chi(T)$ inputs are aligned. We do not quote a single literature mass to two figures: the published numbers differ in their $\chi(T)$ tables and in whether the anharmonic factor is included, and the residual we book below is their mutual spread.

Two upgrades, both available in the published literature, then sharpen the budget without any new theory. First, the temperature dependence of the susceptibility need not be parametrized at all: the lattice computation of Borsanyi et al. (2016) *tabulates* $\chi(T)$ through the QCD crossover, and evaluating the onset condition on the table rather than on the power-law fit would remove the parametrization freedom that carries the largest single line of the budget; the integration of this paper uses the power law with the same paper’s exponent, and books that line as the b row (the table, with the same paper’s exponent, would retire that row). Second, the relativistic degrees of freedom $g_*(T)$, $g_{*s}(T)$ are known to the percent level across the relevant temperatures (Saikawa & Shirai, 2018), retiring the transport item.

What remains is the matching of our leading-order pipeline to the NLO literature, which we treat as a calibration with a booked residual rather than as an uncontrolled error. The booked 5% is specifically the *mutual* spread of the precision computations at aligned $\chi(T)$ inputs — not their total quoted bands, whose $\chi(T)$ and g_* contents are already carried in this table’s own rows (double-counting them here would overstate the budget). One caveat is recorded for the freeze pass: this decomposition of the literature’s error into “already-itemized” and “residual” parts is our reading of the published numbers, and a line-by-line re-audit against the original error tables should confirm it; if the residual proves closer to its conservative ceiling of 8%, the total budget on the closure angle relaxes from $\pm 4\%$ to $\pm 5\%$ — with every central value unchanged, and with the mass untouched either way.

Two further line items are handled explicitly, and they pull against each other. At the closure angle $\theta_i \simeq 1.18$ rad (Section 15.3) the cosine potential is mildly anharmonic, which *lowers* the angle the closure asks for by 4.0% against the harmonic force computed on the same tables; and integrating on the tabulated $g_*(T)$ and $\chi_{\text{QCD}}(T)$ rather than on a fixed $g_* = 73$ with the power-law turn-on *raises* it by about 5%. Earlier revisions applied the first to a baseline that had not received the second and printed 1.12 rad; the two together give 1.18. None of these items touches the mass: with f_K fixed at its calibrated benchmark, the mass is $\sqrt{\chi_{\text{QCD}}}/f_K$ and inherits only the $\pm 1.3\%$ of the lattice susceptibility; every thermodynamic item above moves the angle, not the line. This is

the honest meaning of “first principles” here: a transparent derivation, anchored on the field’s best published numerics, with every residual named and booked against the quantity it actually moves.

25.5 The normalization, cross-checked against the series

The same test also guards the one number the series must carry consistently. A plausible-looking engineering formula of the form $\Omega_K h^2 \simeq 0.12 (\theta_i / (1/2\pi))^2 (f_K / 1.3 \times 10^{13} \text{ GeV})^{1.165}$ — which would imply $\Omega_K h^2 = 0.12$ at $f_K = 1.3 \times 10^{13} \text{ GeV}$ for $\theta_i = 1/(2\pi)$ — cannot be right: evaluating (25.2) at that point gives $\Omega_K h^2 = 0.068$, a factor 1.8 *below* target, so any such normalization would understate the decay constant by a factor $1.8^{1/1.165} \simeq 1.6$ and overstate the mass by the same factor. We record the check plainly, in the spirit of the falsification postulate: the series carries one normalization or it is not a series. GG-A10-Cosmos now carries the same normalization — the integrated $0.135 \theta_i^2 (f_K / 10^{12} \text{ GeV})^{1.165}$, the closure angle $\theta_i = 1.18 \pm 0.04 \text{ rad}$, the decay constant as one calibrated scale, the mass as QCD’s, and the single-patch history as this paper’s Conjecture C-Kaxion-d — with its own earlier statements withdrawn in its record, so the two papers state one doctrine and every number in this paper descends from the integrated pipeline. One disagreement between the two papers is live, and it is recorded in Section 13.5 rather than smoothed over: what the unpinned branch would have implied for the microwave sky.

25.6 The closure

The relation between Ω_K , f_K and θ_i has two unknowns once the bulk normalization is recognized as open (O14), and one of them must be fixed by fiat. The paper states here, once, *which one*. The calibration convention is the *scale*:

$$f_{\text{bulk}} \equiv 3.04 \times 10^{13} \text{ GeV} \quad \text{at} \quad \Xi_{\text{cap}} = 43, \quad \text{hence} \quad f_K = \frac{f_{\text{bulk}}}{43} = 7.1 \times 10^{11} \text{ GeV}, \quad (25.5)$$

declared in advance and not moved afterwards (Section 26.3). The closure angle is then an *output*. Setting $\Omega_K h^2 = \Omega_{\text{DM}} h^2 = 0.1215 \pm 1.5\%$ (Table 7) and solving the calculation of record — the tabulated, full-cosine integration of Section 25 — for the angle:

$$\boxed{\theta_i = 1.18 \pm 0.04 \text{ rad} \quad (\text{tabulated } \chi_{\text{QCD}}(T), \text{ full cosine}); \quad 1.17 \text{ rad on the measured } 0.1200} \quad (25.6)$$

— the generic $\mathcal{O}(1)$ misalignment angle. Which quantity carries the convention is not a matter of taste, and the reason belongs in the open. Were the *angle* declared instead, the scale and with it the laboratory line would be a function of the thermal transport, and would move whenever somebody else’s lattice calculation improved: at the transport of earlier revisions the same declared angle placed the line at 1.78 GHz rather than 1.95, a shift of 173 MHz against a stated continuous error of $\pm 25 \text{ MHz}$. A falsifier cannot be pre-registered against a frequency that drifts that way, so the scale is declared and the angle is allowed to move. The honest reading is otherwise unchanged: the abundance was not fitted to a free f_K , but the frequency is not independent of the abundance either — the benchmark row is the row at which a generic angle saturates the inventory, and the closure returns that angle by construction. What is checkable is that the angle is generic (1.18 rad,

nowhere near the tuned edges), and what would strain it is a confirmed line whose absolute rate implied a tuned angle. The mass is untouched by every item in the budget above except the lattice susceptibility.

26 The mass and the laboratory line

26.1 The mass from the single decay constant

With the decay constant at its calibrated benchmark $f_K = 7.1 \times 10^{11}$ GeV (Section 26.3), the susceptibility relation (24.3) fixes the mass with no further freedom:

$$m_K = \frac{\sqrt{\chi_{\text{QCD}}}}{f_K} = \frac{(5.70 \pm 0.08) \times 10^{-3} \text{ GeV}^2}{7.1 \times 10^{11} \text{ GeV}} = 8.1 \mu\text{eV} (\pm 1.3\%) \quad (26.1)$$

corresponding to the laboratory frequency

$$\nu_K = \frac{m_K c^2}{h} = 1.95 \text{ GHz} (\pm 25 \text{ MHz}), \quad (26.2)$$

using $1 \mu\text{eV} \leftrightarrow 241.80 \text{ MHz}$. This is the one physical line: there is no second frequency, and the bulk normalization scale of Appendix C.1 is an intermediate of the diagonalization, not the decay constant of a propagating field. In the campaign's terms (§10.5), the frequency ν_K of (26.2) is the *interval's* number — the address, carrying the discrete cap window — and the circle's contribution is the count: one light angle, one line. The search strategy of this paper is, in one clause, *search the window by the interval, and confirm the count the circle fixes*.

26.2 The decay constant and its residual window

The decay constant is written as a bulk normalization divided by an integer cap loading Ξ_{cap} (Section 9.3, Appendix C.1). Writing f_{bulk} for the bulk normalization scale,

$$f_K = \frac{f_{\text{bulk}}}{\Xi_{\text{cap}}}, \quad f_{\text{bulk}} \simeq 3.04 \times 10^{13} \text{ GeV}, \quad \Xi_{\text{cap}} \stackrel{\text{bench.}}{=} K + N + L = 43, \quad (26.3)$$

so the benchmark loading gives $f_K \simeq 7.1 \times 10^{11}$ GeV and the numbers above. The third figure of f_{bulk} is carried because Table 11 requires it: 3.04×10^{13} is the only three-figure value that reproduces every row, and rounding it to 3×10^{13} would move the benchmark mass to $8.2 \mu\text{eV}$. The cap computation that fixes Ξ_{cap} is not yet complete: this is the *compact-angle mass-uniqueness conjecture* (stated, with its assumptions enumerated, self-containedly in §26.3), and we fence it as named. The provenance of f_{bulk} must be stated beside it: 3.04×10^{13} GeV is *calibrated* — fixed through the relic closure of Part III, not derived from the geometry; the computation that would derive it (the orientation convention of the anti-self-dual sector and the cap, registered as the anomaly companion's Open Problem O14) is open, and no closed form is printed as distinguished. Its residual is one integer, spanning a *discrete* window ($\Xi_{\text{cap}} \in \{35, \dots, 47\}$) that maps to $\nu_K \simeq 1.59\text{--}2.13$ GHz (Table 11). Section 26.4 gives the ansatz that motivates the benchmark; pending a computation

that would replace it, a definitive-search instrument should treat the window, not the point, as the target (Part VIII). The bulk scale f_{bulk} is an intermediate normalization, not the decay constant of any propagating field — there is one physical f_K , hence one line.

Table 11: The discrete Ξ_{cap} window: each admissible cap loading $\Xi_{\text{cap}} = K + \Sigma$ fixes the single decay constant $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$ and hence the observables. The $\pm 1.3\%$ lattice error on $\sqrt{\chi_{\text{QCD}}}$ applies to every row; the bulk normalization is calibrated at the benchmark row, so the other rows are ratios to it, not independent predictions. Every admissible value is K plus a subset sum of $\{N, L, J\}$ (Section 26.4); the benchmark $\Xi_{\text{cap}} = 43$ is marked.

Ξ_{cap}	composition	m_K [μeV]	ν_K [GHz]
35	K	6.6	1.59
38	$K + N$	7.1	1.72
40	$K + L$	7.5	1.81
42	$K + J$	7.9	1.90
43	$K + N + L$ (benchmark)	8.1	1.95
45	$K + N + J$	8.4	2.04
47	$K + L + J$	8.8	2.13

26.3 The compact-angle mass-uniqueness conjecture, stated and derived

The cap index Ξ_{cap} of Eq. (26.3) is the single quantity on which the *absolute* mass depends, and, together with the calibrated bulk normalization, the ingredient of the mass derivation that this paper carries as a named conjecture rather than a theorem. Because so much leans on it — and because the reader is entitled to see it without reaching for the cosmology paper — we state it here self-containedly, reconstructing it from the GG-6 geometry directly: a precise definition, a precise statement, the one derived clause separated from the ansatz, and then the physical picture at length. *One fact should be held in view throughout: Ξ_{cap} parametrizes where the single line sits — and hence the single mass and frequency — relative to a calibrated benchmark; it moves only where the line sits, never the DFSZ address or the single-line count by which the prediction is identified. The conjecture below controls the address of the Kaxion within its window, not its fingerprints.*

The object, defined.

Definition 26.1 (Cap renormalization index). Let a_{GS} be the tensor remnant — the single light compact-angle pseudoscalar, in its bulk Green–Schwarz normalization — and let a_K be its canonically normalized form (the rotation (9.4) that separates it from the lifted direction plays no part in the normalization). Write f_{bulk} for the bulk Green–Schwarz normalization scale (a geometric intermediate of the diagonalization, not a propagating field) and f_K for the decay constant of the physical eigenstate a_K , the field a haloscope couples to. The *cap renormalization index* is the ratio of the two normalizations, parametrized by ansatz as a positive integer Ξ_{cap} ,

$$f_K = \frac{f_{\text{bulk}}}{\Xi_{\text{cap}}}, \quad m_K = \frac{\sqrt{\chi_{\text{QCD}}}}{f_K}.$$

Ξ_{cap} is thus a parametrization of the ratio between two normalizations of one field. Whether that ratio is an integer at all, and which one, is the content of the ansatz below; what the geometry delivers at face value is a factor of order e^K between the tensor's norm and the Wilson line's, and the absolute normalization is the anomaly companion's Open Problem O14. The mass-uniqueness question is whether the geometry fixes Ξ_{cap} ; this paper does not show that it does.

The statement.

Conjecture MU (compact-angle mass-uniqueness). *The GG-6 geometry determines Ξ_{cap} uniquely, at the value*

$$\Xi_{\text{cap}} = K + N + L = 35 + 3 + 5 = 43,$$

so that the single physical decay constant is $f_K = f_{\text{bulk}}/43 \simeq 7.1 \times 10^{11} \text{ GeV}$, the mass $m_K = 8.1 \mu\text{eV}$, and the Kaxion possesses a single laboratory line at $\nu_K = 1.95 \text{ GHz}$.

The conjecture has one derived clause and two clauses that are an ansatz, and the honesty of the whole paper depends on saying exactly where the line falls. We do so now.

What is derived, and what is ansatz.

Ansatz 26.2 (Cap-loading ladder). On the canonical metric of GG-A2-GG6 with throat depth K and the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$:

- (i) (derived-conditional, with O14 open and restated) the two boundary norms differ by e^{2K} , and neither face-value reading reproduces the calibrated f_K (Appendix C.1); the physical decay constant is therefore not either bulk norm, the bulk scale f_{bulk} of Eq. (26.3) is the bookkeeping quantity $\Xi_{\text{cap}} f_K$, and it is not a propagating line;
- (ii) (ansatz) the ratio of the two normalizations is a positive integer of the additive form $\Xi_{\text{cap}} = K + \Sigma$, with Σ one of the seven cap loadings of Eq. (26.4);
- (iii) (hence, under (ii)) $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$, i.e. $\nu_K \in [1.59, 2.13] \text{ GHz}$ (Table 11), the benchmark row calibrated as in §15.3.

Justification, and its limits. Step 1 (the exponential asymmetry $\Rightarrow$ (i)). Reducing the bulk tensor kinetic term on the canonical metric $ds_6^2 = e^{-2r} \eta dx^2 + dr^2 + d\theta^2$, $r \in [0, K]$ (Appendix C.1), the tensor sector and the Wilson line deliver two boundary kinetic norms: the infrared-dominated $Z_{\text{GS}} = 2\pi \int_0^K e^{2r} dr = \pi(e^{2K} - 1) \simeq \pi e^{2K}$ (C.1), and the ultraviolet-dominated, order-one $Z_\theta = 2\pi \int_0^K e^{-2r} dr = \pi(1 - e^{-2K}) \simeq \pi$ (C.2). The two norms are two *profiles*, and this revision states plainly which field carries each. Earlier revisions attributed the pair to the two chiralities of one tensor, with the orientation convention deciding between them. That attribution does not survive an explicit check, and the check is short enough to print. In the closed source-free scalar truncation of this same background — $H = h_3 + dr \wedge d\theta \wedge a_1$ with all legs of h_3 in the four-dimensional block — closure forces $d_4 h_3 = 0$ and $\partial_r h_3 = 0$, and the six-dimensional Hodge weight on such a three-form is $\sqrt{-g_6} (g_4^{-1})^3 = (e^{-4r} R_\theta)(e^{6r}) = R_\theta e^{2r}$, so that $a_1 = \sigma R_\theta e^{+2r} *_4 h_3$ with $\sigma = \pm 1$: the exponent is

e^{+2r} for *both* signs. Chirality flips the partner’s sign and leaves its radial norm untouched. The tensor’s own norm is therefore Z_{GS} , and the order-one Z_θ is what this section already said further down — the Wilson line’s integral, belonging to a different field. What the anomaly companion registers as open (Open Problem O14) is accordingly restated: not which of two chirality conventions the remnant obeys, but whether a boundary source, a coupled mode sector or the constrained action a genuinely self-dual tensor requires supplies an exponential separation inside the tensor sector at all. Until one of them is exhibited, clause (i) stands *conditional* on it, and this paper grades it so. The norms differ by $e^{2K} = e^{70} \simeq 2.5 \times 10^{30}$; canonical normalization divides the field by the square root of its norm, so the canonical normalization of the light field is *far* from the bulk one (the rotation (9.4) plays no part in this: a rotation preserves canonical normalization), and the configuration in which the bulk normalization would itself be canonical — $\Xi_{\text{cap}} = 1$ — misses by many orders. Hence (i) holds. The order-one integral (C.2) is the Wilson line’s; the tensor and the Wilson line are different fields sharing one integral, and only the tensor’s norm is at issue here. This step does *not* set the value: the factor the geometry delivers at face value is e^K , and $\Xi_{\text{cap}} = 43$ is a benchmark that stands in for a computation not yet formulated; the calibrated f_K corresponds to a descent depth near $K/2$ — which is interesting, and is not proof.

Step 2 (the ansatz behind (ii)). The ansatz reads the ratio of the two normalizations in the throat’s own integer units: the throat depth $K = 35$ — an e-fold count of the warp in curvature units, fixed in GG-A2-GG6, and not a flux count of anything — plus an integer Σ that the ultraviolet terminus of the interval (the “cap”) is taken to contribute, with Σ a subset sum of the ledger integers $\{N, L, J\}$ in which at most a pair of the three appears — $M = 1$ excluded and the triple $N + L + J$ excluded by stipulation, so that the window’s two edges are themselves part of the ansatz (a line at 2.27 GHz, $\Xi_{\text{cap}} = 50$, would retire the “pair” rule rather than the Kaxion). None of this is derived: the admissible class of GG-A3-B16 is a product $\mathcal{M}_4 \times I_r \times S_\theta^1$ with collars and carries no orbifold cap, no quantized flux equal to K is exhibited, and no computation in this paper or its parents produces an integer of the form $K + \Sigma$. The ansatz is motivated by how the same integers function elsewhere in the series, and it is what the discrete window rests on.

Step 2a (one route to the additive form, closed). The obvious mechanism for making a ratio of normalizations an integer of the throat’s own units is additive kinetic loading, and it is closed. With the primitive dualization identity $F^2 = 1/(4\pi^2 Z)$ of Appendix C.1, an additive contribution $Z \propto \Xi$ sends the decay constant as an inverse *square root* of the loading: thirty-five throat units plus eight cap units give $f/f_{35} = \sqrt{35/43} = 0.902$, not $1/43$, and reproducing $1/43$ that way would demand a norm ratio $43^2 = 1849$, hence 1848 added bulk units rather than eight. The ansatz of (ii) is therefore not a statement about kinetic loading — Ξ_{cap} parametrizes a ratio of two normalization *scales* — and the computation Conjecture MU awaits must reach that ratio by another route. We record the closure of this one against our own interest, because it is the route a reader would try first.

Step 3 (enumeration $\Rightarrow$ (iii)). With $(N, L, J) = (3, 5, 7)$ the seven loadings give $\Sigma \in \{0, 3, 5, 7, 8, 10, 12\}$, hence $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$ and, through Eqs. (26.3) and the frequency conversion, $\nu_K \in [1.59, 2.13]$ GHz — the rows of Table 11. Because the bulk

normalization is calibrated at the benchmark row (§15.3), the other rows are ratios to it: the window is the calibrated point times $\Xi_{\text{cap}}/43$.

Item (i) of Ansatz 26.2 is the part a reader can check line by line from the canonical metric, and the condition it now carries is checkable in the same way. Items (ii)–(iii) are the ansatz, and the discreteness of the window is inherited from it, not derived: the prediction lives in a finite window of seven frequencies *if* the ansatz holds, and F-Kaxion-1c/1d are pre-registered against that window with this grade stated.

What remains conjectural: one integer. Everything in Conjecture MU beyond Ansatz 26.2 reduces to two things: the ansatz itself, and, within it, which of the seven loadings Σ the cap carries. Neither is decided by a computation this paper can name on the anomaly companion’s admissible class, which carries no orbifold cap: the “cap” here is the ultraviolet terminus of the interval, and a defining computation for its loading — the manifold, the class, and the integrand — has yet to be formulated. The flux-bookkeeping argument that marks the benchmark as $\Sigma = N + L$ (the ultraviolet family and rank integers at the cap, the infrared dyadic seam J at the boundary) is the subject of §26.4; it is a motivation, not a derivation. Until a defining computation exists, the honest statement is the *window* of Ansatz 26.2, with $\Xi_{\text{cap}} = 43$ its marked, calibrated row.

Physical picture I: the throat sets the normalization. The angle of §9.3 is normalized in one place and read in another, and the conversion between them *is* the throat. The field is normalized deep in the bulk, where the scale profile e^{-2r} has run through its full K e-folds of redshift; the reading a detector takes sits at the throat’s end. Converting the rate of one into the rate of the other multiplies the frequency by the throat depth measured in the geometry’s integer units, K , plus whatever the cap adds. The mass-uniqueness conjecture is the assertion that this conversion factor is not a tunable dial but an integer the theory already owns — the throat depth K , sharpened by a small cap loading. The pictures in this and the next two paragraphs are the intuition behind the ansatz, not a derivation of it: the geometry delivers two face-value chiral norms a factor e^{2K} apart, neither of which is the calibrated f_K ; reading the conversion as 43 integer units is the ansatz.

Physical picture II: one line, not two. There is a single physical decay constant and therefore a single line. The bulk normalization scale f_{bulk} is *not* an alternative prediction at a different frequency: it is the normalization the Green–Schwarz direction carries in the bulk, related to the canonical f_K by the integer Ξ_{cap} of the ansatz, as a wavefunction renormalization relates a bare field to its physical normalization. Ξ_{cap} is, under the ansatz, an integer the geometry owns rather than a tunable parameter; what has yet to exist is the computation that would show the geometry owns it, and fix its value within the discrete window. Crucially, the abundance is computed with the physical f_K (Section 15.3): the relic and the laboratory line use the *same* decay constant, so the mass that a haloscope reads is the mass that set the cosmological oscillation.

Physical picture III: why an integer is a gift, not a hedge. Under the ansatz the residual uncertainty in the absolute mass is not a continuous error bar but a list of seven candidate frequencies, and that discreteness is what makes the conjecture disposable rather than permanent.

It can be retired two independent ways. *By theory*: a defining computation for the cap loading, once formulated, returns Σ and collapses the window to a point — or returns something not of the form $K + \Sigma$ and retires the ansatz. *By experiment*: at the $\pm 1.3\%$ the lattice leaves on the line at fixed f_K (± 25 MHz against rung spacings of 45–136 MHz), the discovery frequency read against Table 11 resolves every pair of rungs but 42/43 (separated by 1.8σ): a line on a rung measures Ξ_{cap} outright at that resolution, and a line confirmed between rungs retires the ladder ansatz while leaving the Kaxion, its address and its count untouched. That is the sense in which the absolute mass, though presently a calibrated benchmark inside an ansatz-graded window, is nonetheless sharply testable.

26.4 The cap index: from benchmark toward theorem

The window of Table 11 is not an error bar; it is one undetermined integer, and the table itself displays the structure that constrains it: *every admissible value of Ξ_{cap} is K plus a subset sum of the remaining ledger integers $\{N, L, J\}$,*

$$\Xi_{\text{cap}} = K + \Sigma, \quad \Sigma \in \{0, N, L, J, N+L, N+J, L+J\}, \quad (26.4)$$

with the benchmark at $\Sigma = N+L$, i.e. $\Xi_{\text{cap}} = 35 + 3 + 5 = 43$. This subsection records the physical argument that selects the benchmark, at the level of flux bookkeeping, together with the computation that would promote it to a theorem.

The bookkeeping argument. Under the ansatz, Ξ_{cap} measures how much kinetic normalization the *cap* — the geometry that closes the throat at its ultraviolet end — adds to the throat’s own K units. The cap is not free to carry arbitrary quanta: it is the terminus of the flux lines that thread the throat, and the question is which lines terminate *there* rather than at the infrared boundary. The ledger integers divide cleanly on this question. The family-projector flux (N units) and the gauge-rank flux (L units) encode the $\text{SU}(5) \rightarrow \text{SM}$ breaking data — ultraviolet structure, fixed where the gauge symmetry is whole, i.e. at the cap. The dyadic seam depth (J units), by contrast, is infrared data: it is the depth of the boundary seam that sets $\alpha^{-1} = 2^J$ and the lepton dyadics of GG-A9-SM, and it terminates at the boundary, not at the cap. The cap therefore carries the (N, L) pair and not J :

$$\Sigma = N + L \implies \Xi_{\text{cap}} = K + N + L = 43. \quad (26.5)$$

On this argument the window of Table 11 collapses to its boldface row: one line at $\nu_K = 1.95$ GHz, standing alone.

Status, honestly fenced. Equation (26.5) is a *flux-bookkeeping ansatz*, not a theorem and not yet a well-posed computation. What is established: the subset-sum structure (26.4) (arithmetic of the ledger, given the ansatz); the UV/IR division of the ledger integers (N, L ultraviolet, J infrared — this is how the same integers function everywhere else in the series). What is *not* done: a defining computation for the cap loading on the anomaly companion’s admissible class, which is a product with collars and carries no orbifold cap; the manifold, the class and the integrand of such a computation have not been written down, and until they are, “ $\Sigma = N+L$ ” is a motivated mark, not a

result. Formulating that computation is the sharpest single open item of this paper, and we register it as such in Part XI. Until it is done, the paper’s experimental guidance is unchanged — scan the window — but the window now has an interior structure: a marked row, an argument for the mark, and a named computation that decides it. GG-A10-Cosmos’s compact-angle-mass-uniqueness conjecture, which GG-A10-Cosmos states at exactly this integer, closes the moment the integral does.

27 The photon coupling, factor by factor

27.1 The exact matching formula

For *any* axion with the gluonic portal (24.1), the low-energy photon coupling is fixed by anomaly matching plus chiral dynamics, with no model freedom beyond one ratio:

$$g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K} \left[\frac{E}{N} - \frac{2}{3} \frac{4m_d + m_u}{m_d + m_u} \right] = \frac{\alpha}{2\pi f_K} \left[\frac{E}{N} - 1.92(4) \right]. \quad (27.1)$$

The structure of (27.1) deserves one transparent paragraph, because every factor in it is forced. The first term is ultraviolet: E/N is the ratio of the electromagnetic to the color anomaly of whatever fermions generate the portals — pure group theory. The second term is infrared and *unavoidable*: below the QCD scale, the gluon coupling (24.1) is removed by an axial rotation of the light quarks, which deposits an axion component into the neutral mesons; the physical axion state therefore mixes with π^0 and η , whose two-photon couplings feed back the model-independent $-1.92(4)$ (leading order $-\frac{2}{3}(4m_d+m_u)/(m_d+m_u) = -2.03$ with $m_u/m_d = 0.462$; next-to-leading-order chiral corrections bring -1.92 (Grilli di Cortona et al., 2016)). Any QCD-coupled axion — and Theorem 24.1 made the Kaxion one — pays this chiral toll. The fine-structure constant in (27.1) is the photon’s, evaluated at the photon’s scale: $\alpha = \alpha(0) = 1/137.036$, since a haloscope converts micro-eV quanta, twelve orders of magnitude below any running threshold.

27.2 $E/N = 8/3$: the ledger fixes the ratio

Theorem 27.1 (T-Kaxion-3: the anomaly ratio). *For the Kaxion, $E/N = 8/3$ exactly.*

Proof. By Theorem 24.1 the color anomaly descends solely from the $\omega_3(A_5)$ entry of (6.13). The electromagnetic anomaly descends from the same entry and from no other: the photon is a generator inside SU(5) (GG-A2-GG6 T3), while the remaining Abelian ledger entries couple to F_B and F_{Q_i} — the baryonic and informational gauge fields — to which the photon is blind at tree level (the one-loop B – Y kinetic mixing GG-A3-B16 concedes enters E/N at second order in the mixing and is negligible). Both anomalies are therefore traces over the *same* complete SU(5) structure, and for complete unified multiplets the ratio is fixed by the embedding alone: $E/N = \text{Tr } Q_{\text{em}}^2 / \text{Tr } T_{\text{color}}^2 = 8/3$, the classic grand-unification value (Kim & Carosi, 2010). No GG-6-specific input enters beyond the statement — itself a theorem (GG-A2-GG6 T3) — that the photon sits inside SU(5) and the anomaly inflow is the complete-multiplet trace. ■

The group theory deserves to be seen once with bare hands, because it is a computation any student can check against a table of Standard-Model charges. Take the $\bar{\mathbf{5}}$ of SU(5): its states are the

anti-down triplet ($Q = +\frac{1}{3}$, three colors) and the lepton doublet ($Q = 0, -1$). The electromagnetic anomaly weight is $E(\bar{\mathbf{5}}) = \sum Q^2 = 3 \times \frac{1}{9} + 0 + 1 = \frac{4}{3}$; the color weight is the Dynkin index of the single color triplet, $N(\bar{\mathbf{5}}) = \frac{1}{2}$. Ratio: $\frac{4/3}{1/2} = \frac{8}{3}$. Now the **10**: quark doublet ($Q = \frac{2}{3}, -\frac{1}{3}$, three colors each), anti-up triplet ($Q = -\frac{2}{3}$), positron ($Q = +1$): $E(\mathbf{10}) = 3(\frac{4}{9} + \frac{1}{9}) + 3 \cdot \frac{4}{9} + 1 = 4$; color content one doublet pair plus one triplet, $N(\mathbf{10}) = \frac{3}{2}$. Ratio: $\frac{4}{3/2} = \frac{8}{3}$ again. *Every* complete SU(5) multiplet returns 8/3 separately — which is why the theorem needs no inventory of the bulk fermion content, only the statement that the content fills complete multiplets (forced by GG-A2-GG6 T3).

The infrared half of the bracket is equally checkable. Removing the gluon coupling (24.1) by the axial rotation $q \rightarrow e^{i\gamma_5 a_K Q_a/2f_K} q$ of the light quarks (with $\text{tr } Q_a = 1$ distributed to minimize mass mixing) shifts the photon coupling by $-\frac{2}{3}(4m_d + m_u)/(m_u + m_d)$ at leading order — the π^0 and η poles doing the arithmetic. Numerically, with $r = m_u/m_d = 0.462$:

$$-\frac{2}{3} \cdot \frac{4+r}{1+r} = -\frac{2}{3} \times \frac{4.462}{1.462} = -2.03 \xrightarrow{\text{NLO}} -1.92(4), \quad (27.2)$$

the next-to-leading-order shift and its error being the chiral expansion at work (Grilli di Cortona et al., 2016). Assembling: $E/N - 1.92 = 8/3 - 1.92 = 0.75(4)$. Both halves of the Kaxion’s photon coupling are thus open-book computations: ultraviolet group theory anyone can redo in five lines, and infrared meson mixing whose error bar is the Standard Model’s, not the series’.

A remark: 8/3 is also the DFSZ value, for the same group-theoretic reason (complete multiplets); the consequence is experimentally golden and is drawn out below. Anomaly coefficients are integers protected by topology — they do not run, admit no throat dilution, and pair with the infrared α , not the electroweak one.

27.3 Numbers, with errors

The bracket in (27.1) evaluates to $8/3 - 1.92(4) = 0.75 \pm 0.04$ ($\pm 5.4\%$; the error is entirely the Standard-Model chiral term — the 8/3 is exact). Hence

$$g_{K\gamma\gamma} = \frac{0.75(4)}{137.036 \times 2\pi f_K} = (1.2 \pm 0.07) \times 10^{-15} \text{ GeV}^{-1} \quad (27.3)$$

with the single decay constant at its calibrated benchmark $f_K = 7.1 \times 10^{11} \text{ GeV}$, the error at fixed f_K being the chiral bracket’s (5.4%) with the lattice susceptibility’s 1.3% in quadrature ($\pm 5.5\%$ total, plus the discrete Ξ_{cap} window on the absolute scale).

27.4 The Kaxion sits on the DFSZ line

Because $E/N = 8/3$, the Kaxion’s photon coupling at any mass equals the DFSZ axion’s at the same mass: in the standard $(m_a, g_{a\gamma\gamma})$ exclusion plane, *the Kaxion line is the DFSZ line*, entered at the single mass of Section 26, inside its discrete window. The experimental convenience is hard to overstate: every published and projected “DFSZ sensitivity” statement in the haloscope literature applies to the Kaxion verbatim, with no rescaling. What distinguishes a Kaxion from a DFSZ axion is then *not* the photon coupling but everything DFSZ does not have: the predicted absolute

masses (DFSZ predicts none), the fenced window with its single-line count, and the absence of any Peccei–Quinn scalar sector at f_K . Identification, not reach, is where the Kaxion is falsifiable — exactly the design philosophy of Parts VI–VIII.

27.5 The baryonic entry, and why it opens no laboratory channel

The same descent that fixed the photon channel transports the ledger’s baryonic entry. The Abelian entries couple a_K to $F_B \tilde{F}_B$ with the trace-ratio weight (6.17): under the anomaly companion’s normalization the baryonic entry stands at

$$\frac{\text{Tr}(B_0^2)}{\text{Tr}(Y_0^2)} = \frac{L - N}{L} = \frac{2}{5} \text{ of the hypercharge entry,} \quad (27.4)$$

parameter-free and frame-invariant (a photon-relative amplitude would require the electromagnetic embedding’s trace normalization to be fixed alongside it, which this paper does not do). The entry opens no laboratory channel: the boundary condition that terminates $U(1)_B$ removes its zero mode, the lightest baryonic vector is the Z'_B at M_{lock} , and a μeV condensate has nothing baryonic to drive at laboratory frequencies. The nucleon-spin couplings that survive at the chiral level are quoted in Part VIII on the KSVZ-class sheet with their own fences.

28 Abundance recheck and lifetime

28.1 Closing the loop

Self-consistency demands that the derived pair (m_K, f_K) reproduce the abundance that closed it. Running the calculation of record forward with the benchmark $f_K = f_{\text{bulk}}/43 = 7.07 \times 10^{11} \text{ GeV}$ and the closure angle $\theta_i = 1.18 \text{ rad}$: the onset condition gives $T_{\text{osc}} = 1.04 \text{ GeV}$, the onset mass $m_K(T_{\text{osc}}) = 4.7 \times 10^{-4} m_K(0)$, and the redeemed density $\Omega_K h^2 = 0.1215$ — the π -locked target. (At the three-figure $f_K = 7.1 \times 10^{11} \text{ GeV}$ printed elsewhere it returns 0.122; the difference is the rounding of f_K , not of the pipeline.) The loop closes by construction, and we say so: the benchmark was calibrated to close it, so this is a check of arithmetic, not of physics.

28.2 Lifetime: stability derived, not assumed

The Kaxion’s only kinematically open decay is to two photons, each of energy $m_K/2$ (a pair of $\sim 0.98 \text{ GHz}$ quanta — microwave photons). The rate follows from the same textbook steps as $\pi^0 \rightarrow \gamma\gamma$: the operator $-\frac{g_{K\gamma\gamma}}{4} a F \tilde{F}$ gives the two-photon amplitude $\mathcal{M} = g_{K\gamma\gamma} \epsilon^{\mu\nu\rho\sigma} \epsilon_{1\mu}^* k_{1\nu} \epsilon_{2\rho}^* k_{2\sigma}$; squaring, summing polarizations, and integrating the two-body phase space with the identical-particle factor $\frac{1}{2}$ yields $\Gamma = g_{K\gamma\gamma}^2 m_K^3 / 64\pi$.

The width of $a \rightarrow \gamma\gamma$ for the pseudoscalar coupling (27.1) is the textbook

$$\Gamma_{K \rightarrow \gamma\gamma} = \frac{g_{K\gamma\gamma}^2 m_K^3}{64\pi}, \quad \tau_K = \Gamma^{-1}. \quad (28.1)$$

Numerically, with (26.1) and (27.3):

$$\tau_K = \frac{64\pi}{g_{K\gamma\gamma}^2 m_K^3} \simeq 1.7 \times 10^{50} \text{ s} \simeq 4 \times 10^{32} t_U, \quad (28.2)$$

with $t_U = 4.35 \times 10^{17} \text{ s}$. The error ($\tau \propto g^{-2} m^{-3}$; $\pm 12\%$) is irrelevant against thirty-two orders of margin. Decays to electrons are kinematically forbidden by eleven orders of magnitude ($m_K \ll 2m_e$); decays to neutrinos are helicity- and coupling-suppressed beyond relevance; stimulated decay in astrophysical photon baths shortens nothing observably at these couplings. Stability over cosmological time — assumed silently by every dark-matter fit — is for the Kaxion an output, with thirty-two orders of headroom: Corollary 8.3, now with numbers.

29 Error budget and the theorem-versus-conjecture ledger

29.1 The itemized budget

All errors are multiplicative and combined in quadrature in log-space; discrete (non-Gaussian) uncertainties are kept separate and quoted as windows.

Two rows carried by earlier revisions are absent below, and their absence is the substance of this revision’s change rather than an omission. The power-law susceptibility exponent ($b = 4.08 \pm 0.2$, 4.8%, formerly the largest row) and the anchor ($T_{\text{QCD}} = 150 \pm 10 \text{ MeV}$, 2.3%) parametrized a mass turn-on that the calculation of record no longer uses. Integrating on the tabulated $\chi_{\text{QCD}}(T)$ retires both — exactly as the comment column of the earlier table predicted it would — and carries the table’s own quoted errors instead. That is why the band below is $\pm 4\%$ and not $\pm 6\%$.

Table 12: Error budget for the four headline predictions, in two blocks. Under the calibration convention of Section 15.3 the mass and frequency at fixed f_K inherit only the lattice susceptibility; every thermodynamic item of the misalignment pipeline enters the closure angle. “GG” sources are series-internal; “SM” sources are Standard-Model inputs. Exponents, measured on the integration of record rather than assumed from the harmonic scaling ($\partial \ln \Omega_K / \partial \ln \theta_i = 2.17$, $\partial \ln \Omega_K / \partial \ln f_K = 1.163$): $\theta_i \propto \omega^{0.46} f_K^{-0.54}$; $m_K = \sqrt{\chi}/f_K$; $g_{K\gamma\gamma} \propto [E/N - 1.92]/f_K$.

Source	Type	Effect on θ_i	Comment
π -locked target $\rightarrow \omega$	GG	0.7%	target derived-conditional (GG-A10-Cosmos, with a declared epoch that enters at the percent level); $h = 0.674$ is the Planck-fit CMB-branch convention shared with GG-A10-Cosmos (the corridor’s own H_0 spread is the discrete row below)
Lattice $\sqrt{\chi_{\text{QCD}}} (\pm 1.3\%)$	SM	0.5%	enters the closure through the onset
$\chi_{\text{QCD}}(T)$ pointwise envelope	SM	2.5%	the tabulated errors of Borsanyi et al. (2016) moved coherently: $\mp 5.2\%/+5.5\%$ in Ω
EOS interpolation choice	SM	0.7%	shape-preserving against cubic splines on Table S2: $+1.5\%$ in Ω
$g_*(T)$, $g_{*s}(T)$ transport residual	SM	1%	Borsanyi et al. (2016) Table S2 against Saikawa & Shirai (2018)
NLO matching residual	SM	2.3%	mutual spread of the precision computations (Gorghetto & Villadoro, 2019; Grilli di Cortona et al., 2016), $\sim 5\%$ in Ω
Total (Gaussian), θ_i		$\pm 4\%$	$\theta_i = 1.18 \pm 0.04$ rad at the benchmark row
Source	Type	Effect on m_K, ν_K	Comment
Lattice $\sqrt{\chi_{\text{QCD}}} (\pm 1.3\%)$	SM	1.3%	the only continuous error on the line at fixed f_K : $\pm 0.1 \mu\text{eV}$, ± 25 MHz
f_{bulk} calibration	GG	convention	the declared input $f_{\text{bulk}} \equiv 3.04 \times 10^{13}$ GeV at $\Xi_{\text{cap}} = 43$ defines the benchmark row (25.5); the closure angle is its output and carries the band above. The convention is pre-declared and will not be moved after a null. It is stated on the scale rather than on the angle so that the address does not follow the thermal transport: declaring $\theta_i(\text{bench})$ instead would have moved the line to 1.78 GHz when the transport was corrected in this revision. No Gaussian error, and no derivation (O14)
Chiral bracket 1.92(4) (adds to $g_{K\gamma\gamma}$ only)	SM	5.4%	$\Rightarrow g_{K\gamma\gamma}$ total $\pm 5.5\%$ at fixed f_K
Single-patch history	GG	assumed	named conditional on the inflation-era pinning of the angle (Conjecture C-Kaxion-d, §13.5); the post-inflationary variant (θ_i randomized) moves the line to $\gtrsim 6$ GHz, outside the window
$\Xi_{\text{cap}} = 43$ (cap loading; fixes f_K)	GG	discrete	ladder $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\} \Rightarrow \nu_K \in [1.59, 2.13]$ GHz; the ladder is an ansatz (§26.4)
H_0 convention (Planck-fit vs. corridor)	GG	discrete	the series’ derived $H_0 = 66.6\text{--}67.4$ shifts ω by up to -2.4% ($\theta_i \rightarrow 1.17$ rad at the cohort’s canonical 66.6, on $\theta_i \propto \omega^{0.46}$) — a convention choice, not a Gaussian error

29.2 Why not a few percent? (And when it will be.)

The reader of GG-A9-SM is entitled to ask what the Kaxion’s mass actually carries, when the same program delivered twenty-eight Standard-Model parameters at the per-mille-to-percent level. The answer has two parts, and the second is the one that matters.

The Standard-Model successes are *algebraic locks*: $\alpha^{-1} = 2^J$, $\sin^2 \theta_W = L^2/4N^3$, $v = M_{\text{lock}}(2/L)^N$ — finite expressions in the ledger integers, evaluated once. The Kaxion mass at fixed f_K is algebraic too, $m_K = \sqrt{\chi_{\text{QCD}}}/f_K$, and carries only the lattice’s $\pm 1.3\%$; the QCD thermodynamics of Section 25 — $\chi(T)$, $g_*(T)$, the onset transport — enters the closure *angle*, not the line. The mass’s real uncertainty is of a different kind and is not a percentage: the absolute normalization of f_K is calibrated, not derived (O14), and the geometry’s discrete residual is the ladder $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$, a window of $-19\%/+9\%$ around the benchmark ($\pm 14\%$ about its center). The geometry supplies the ladder and the count exactly; it does not yet supply the rung.

The corollary is a statement about where improvement can come from. The lattice campaign on $\chi(T)$ exists (Borsanyi et al., 2016) and the transport tables exist (Saikawa & Shirai, 2018); this revision integrates on them, which has retired the two largest rows of the earlier budget and brought the angle’s band from $\pm 6\%$ to $\pm 4\%$ with no new theory, and a tightened next-to-leading-order matching would take it toward $\pm 2\%$. Nothing on the Standard-Model side can narrow the *line*: at $\pm 1.3\%$ the rungs of the ladder are already resolved from one another in principle (their spacing is 45–136 MHz against ± 25 MHz; the 42/43 pair only at 1.8σ), so a confirmed line inside the window would read off Ξ_{cap} at that resolution. What would narrow the line is on the geometry’s side — the computation behind O14 and a defining computation for the cap loading (Section 26.4) — and until it exists the window, not the benchmark, is the prediction.

29.3 The master ledger

Table 13: Theorem-versus-conjecture ledger for every load-bearing claim of Parts II–V. Every load-bearing claim of the middle Parts appears once, with what it rests on, whether it is a theorem or a fenced conjecture, and how much margin it carries. The third column is the one to read first, because it is where the paper says which of its own claims are not yet proved. A framework that printed only its theorems would be easier to read and less useful.

Claim	Parent	Status	Error
Kaxion exists (T-Kaxion-1)	GG-A2-GG6 T1–T6, GG-A3-BI6, GG-A5-GDOS, GG-A6-GRIP	theorem	—
Gluonic portal, $N=1$, $N_{\text{DW}}=1$ (T-Kaxion-2)	GG-A3-BI6 ledger + GG-A5-GDOS descent	theorem ($N=1$ in the adopted period con- vention; N_{DW} open, 1 or 5)	—
$\bar{\theta} \rightarrow 0$; residual (24.2)	GG-A9-SM (bare angle) + T-Kaxion-2 relaxation	derived-cond. (GG-A9-SM)	physical residual $\sim 10^{-17}$ (CKM floor of the relaxation); GG-A9-SM’s 4.7×10^{-18} is the bare angle’s bound
$m_K f_K = \sqrt{\chi_{\text{QCD}}}$	T-Kaxion-2 + lattice	theorem + SM input	$\pm 1.3\%$
Line count: single (no partner)	mode analysis (GG-A3-BI6); Part VI	derived- cond. (five registered exposures, §9.4)	structural
Baryonic entry: no channel below M_{lock}	termination (GG-A3-BI6); Part VI	theorem	structural
$E/N = 8/3$ (T-Kaxion-3)	GG-A2-GG6 T3 + GG-A3-BI6	theorem	exact
Abundance target $\Omega_{\text{DM}} = 1/\pi - 1/2\pi^2$	GG-A10-Cosmos	derived-cond. (GG-A10- Cosmos Conj. VIII / T-Cosmos-17; C-Kaxion-c)	$\pm 1.5\%$ (as ω)
$\theta_i = 1.18$ rad (initial angle)	abundance closure	output (generic)	$\pm 4\%$; the angle at which the declared benchmark scale saturates Ω_{DM}
Misalignment completeness f_K (single decay constant)	GG-A10-Cosmos, named calibrated f_{bulk} (O14); cap loading $\Xi_{\text{cap}}=43$	conjecture calibrated/ bench.	sole saturator of Ω_{DM} ladder
m_K, ν_K (one physical line)	Eqs. (26.1)–(26.2)	calibrated (via f_K)	$\Xi_{\text{cap}} \in \{35, \dots, 47\}$ $\pm 1.3\% + \text{window}$
$g_{K\gamma\gamma}$ (DFSZ line)	Eq. (27.3)	calibrated (via f_K)	$\pm 5.5\%$
Lifetime $\tau \gg t_U$	Eq. (28.1)	derived	32 orders of headroom

29.4 What Part V has established

Starting from the constitution and the series theorems alone — taking nothing on authority, and re-deriving the one misalignment normalization the series must carry (Section 25.5) along the way — we have established: the gluonic portal that makes the strong-CP assignment legitimate (T-Kaxion-2); the single decay constant at its calibrated benchmark, $f_K = 7.1 \times 10^{11}$ GeV (cap loading $\Xi_{\text{cap}} = 43$, bulk normalization calibrated through the relic closure), and with it the one physical line — mass $8.1 \mu\text{eV}$, frequency 1.95 GHz, each $\pm 1.3\%$ at fixed f_K — inside a discrete window of loadings pending a defining computation for the cap (Section 26.4); the misalignment angle at which that benchmark saturates the observed abundance, the generic $\theta_i \simeq 1.1$ rad; the photon coupling on the exact DFSZ line, $g_{K\gamma\gamma} = (\alpha/2\pi f_K) \times 0.75(4)$; the ledger-locked baryonic entry, with no channel below M_{lock} ; and a lifetime thirty-two orders beyond the age of the universe. Every number carries its error against the quantity it moves; the one thing no percentage covers — the absolute normalization of f_K — is named as open; and the entire sheet is killed or confirmed by the experiments of Parts VII–VIII.

In plain words, this Part did the thing dark-matter theories are never able to do: it wrote the answer key before the exam. A mass, a frequency, a brightness, a lifetime — each the end of an arithmetic chain that begins at five integers and one calibrated scale, and passes through nothing adjustable after that. From here the paper changes character: the geometry has said everything it has to say, and every remaining page is about machines. Part VI sharpens the two signatures the machines will interrogate first.

30 Part V in plain English: the four numbers, and where each one comes from

This is the Part the paper exists for. Four quantities are derived: the mass, the photon coupling, the abundance, and the lifetime.

The mass comes from a relation the QCD axion literature has used for decades — the mass times the decay constant equals the square root of the topological susceptibility — with the difference that here the decay constant is not free after one calibration: its benchmark is set once, through the observed abundance, and the geometry’s integers fix the ladder around it. The photon coupling comes from the charge assignment, and the assignment is not chosen either; it is what the embedding permits, which is why the Kaxion lands on the DFSZ line rather than near it. The abundance follows from the cosmology of the previous Part. The lifetime follows from the coupling.

The honest accounting is in the error budget rather than in the central values. The framework’s integers carry no uncertainty; its one calibrated scale, its ladder and its conditionals are the GG-side rows of the budget, quoted as windows rather than bands — so every Gaussian band quoted here comes from the standard-model side: the lattice value of the susceptibility, the QCD transition temperature, the count of relativistic degrees of freedom. The exponents in the caption say how each of those propagates. A reader who wants to know how firm the frequency is should read the exponents, not the central value.

What the Part does not claim is that these four numbers are exact. It claims that they are not adjustable, which is a different and more useful property.

Part VI

Unique fingerprints

The single line, the fenced window, and the count test

31 The single line and the count

31.1 The claim

The spectral claim of this paper is a count: the compact-angle sector delivers *exactly one* light QCD-coupled pseudoscalar, and therefore exactly one dark-matter line — inside the fenced window of Part V, with no partner tone at any spacing. The count is a derived-conditional consequence of the assembled structures, for the two parent-backed reasons of §9.3 — the circle’s holonomies are Hosotani-lifted at the circle’s curvature scale, and the boundary anomaly is carried once, by the single tensor’s inflow — conditional on the five registered exposures (nine modes) the parents carry with open masses being heavy or not QCD-coupled (§9.4). One angle of the compact-angle sector stays light, and one line rings; whether anything else does is exactly what the count test asks.

31.2 The count test

A confirmed dark-matter line, at any frequency inside the window, examined for a second resolved line elsewhere in the window, costs a reanalysis of data the scan has already taken — no new hardware, no new integration time. The pre-registered statement: the Kaxion of this paper stands *alone*; a second light QCD-coupled line inside the window, established at the same grade as the first (two persistent excesses with the halo line shape of §33, each surviving the same systematic vetoes), would falsify the count and, with it, the conditional on the parents’ registered exposures that the count rests on (§9.4) — and it would say which exposure was light. The test is two-sided in the honest direction: it cannot confirm the framework — many models produce one line — but it can kill the count cheaply — the Kaxion, its address and the window stand, and a registered mode is named light — and every scan that crosses the window performs it for free.

32 The baryonic entry, stated once

The BI-6 ledger’s baryonic entry is the $-\frac{2}{5}$ of X_4 , the trace ratio $\text{Tr}(B_0^2)/\text{Tr}(Y_0^2) = (L - N)/L$ under the anomaly companion’s normalization. It is stated once, because its laboratory consequence is a negative one: *there is no baryonic channel below M_{lock}* . The boundary condition that terminates $U(1)_B$ removes its zero mode; the lightest baryonic vector is the Z'_B at 3.85 TeV (Arisaka, 2026, P1-KK) — placed on a vanishing seam kinetic term, $\rho_K = 0$; a seam term of the size GG-A3-BI6

registers as O15 would move it to ~ 22 TeV, which only strengthens the absence — and a μeV condensate has nothing baryonic to drive at laboratory frequencies. The ledger entry is real, and its laboratory-scale consequence is its absence: the Kaxion’s measurable portals are the photon and gluon couplings of Part V; the baryonic sector’s own exposure is the octet–graviton pattern and the massive-tower inflow bookkeeping (O13) that GG-P1-KK and GG-A3-B16 register, the Dirichlet vector itself being unproduced by light quarks at tree level. Nucleon-*spin* couplings survive at the chiral level, on the KSVZ-class sheet quoted with its fence in Part VIII — and that sheet places them out of reach at this f_K .

33 Finite- Q morphology and falsifiers

33.1 What the line looks like in a real receiver

Two widths govern the observed morphology. The *condensate linewidth* is set by the galactic velocity dispersion, $\delta\nu/\nu \simeq v^2 \sim 10^{-6}$ ($Q_a \equiv \nu/\delta\nu \simeq 10^6$): about 2.0 kHz at the laboratory line. The *receiver bandwidth* at realistic loaded quality factors, $1/Q_L \sim 10^{-5}$, comfortably contains it. A single line, kilohertz-narrow, at one frequency inside a fenced window: morphologically, the search is the classic haloscope search, and every published scan protocol applies without modification.

33.2 The line shape and the search template

For completeness, the spectral shape the search must fit. A virialized, isotropic halo gives the tone the standard asymmetric profile

$$f(\nu) \propto \sqrt{\nu - \nu_0} \exp\left[-\frac{3(\nu - \nu_0)}{\nu_0 \langle v^2 \rangle}\right] \quad (\text{galactic frame}), \quad \nu \geq \nu_0, \quad (33.1)$$

rising from the rest-mass edge ν_0 with width $\sim \nu_0 \langle v^2 \rangle \simeq 10^{-6} \nu_0$ — the sharp low edge is the rest mass, the tail is kinetic energy. The laboratory rides the halo wind, so the template a search fits is the lab-frame form of Part IX — edge at ν_0 , maximum $\simeq 0.6$ kHz above it, full width at half maximum $\simeq 1.8$ kHz — of which this is the galactic-frame parent; the template is a single-tone shape either way; the count test of §31.2 runs on the same spectra, as a second-line search with the same profile.

33.3 The pre-declared falsifier

Collecting the Part’s results under the falsification postulate (MF) of GG-A1-Constitution:

F-Kaxion-1. The Kaxion of this paper is excluded if any one of the following is established: (c) a completed scan of the closed window $\nu_K \in [1.59, 2.13]$ GHz with no persistent line, at a depth defined as follows: a 90%-confidence exclusion $g_{90} \leq 0.8 g_{\text{DFSZ}}$ at every frequency in the window, evaluated at the nominal local density $\rho_{\text{loc}} = 0.45 \text{ GeV cm}^{-3}$ with the standard-halo line shape of Part IX (the factor 0.8 is a 1.6σ margin on the 12% quadrature sum of the $\pm 5.5\%$ coupling error and the $\pm 22\%$ halo-density prior of Part IX propagated as $\sqrt{\rho}$, so that the sheet’s DFSZ line is excluded rather than grazed; “DFSZ depth” means this criterion wherever the paper uses the phrase, and a

scan reaching only g_{DFSZ} itself at nominal density is registered as a partial null) — the window kill, pre-registered against the seven loadings $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$ so that a null across the window counts regardless of the ladder’s later status. What (c) kills is the calibrated benchmark with its ladder and the pre-inflationary reading that placed it; the calibration convention is declared in advance — $f_{\text{bulk}} \equiv 3.04 \times 10^{13}$ GeV at $\Xi_{\text{cap}} = 43$, with $\theta_i(\text{bench}) = 1.18$ rad as its output — and will not be moved after a null. (d) a confirmed axion-like dark-matter line at a frequency *outside* the window — the location kill, which kills the same objects as (c); a line at the randomized-angle mass ($\gtrsim 6$ GHz) would instead kill the inflation-era pinning C-Kaxion-d, and a line at 2.27 GHz the “pair” rule of the ladder. (e) a confirmed line inside the window whose coupling sits off the DFSZ locus at $> 5\sigma$ — the address kill, which falsifies $E/N = 8/3$ and with it the theorem-level sheet. And the count test: a *second* light QCD-coupled line inside the window, established at the same grade as the first, falsifies the count (T-Kaxion-5) and identifies which registered mode is light; the Kaxion, its address and its window stand. (The identifiers F-Kaxion-1a and 1b are unassigned.)

33.4 What a discovery sequence looks like

The falsifier has a constructive mirror: the step-by-step protocol by which a candidate line would be promoted to a Kaxion discovery. *Step 1 — the line.* A persistent excess with the lab-frame template of Part IX (the boosted form of (33.1)), surviving cavity retuning, rescans, and the standard systematic vetoes. *Step 2 — the address.* The frequency inside the fenced window, on the DFSZ line at $E/N = 8/3$ — the location Part V fixes, with its calibrated normalization graded as such. *Step 3 — the count.* The count test of §31.2: the line stands alone across the window. *Step 4 — the modulations.* The annual and daily fingerprints of Part IX, which no terrestrial artifact shares. A candidate that passes all four is a Kaxion at the strength this paper can confer; what would remain is the community’s replication.

One fingerprint, then, and no knob left free anywhere in it: a single kilohertz-narrow line, inside a discrete window of cap loadings around a calibrated benchmark, at the DFSZ location, alone. The window is the ladder read out as a frequency band, its placement a calibration graded as such in Part V and its ladder an ansatz graded as such; the address is the ledger’s theorem. Every clause is stated at the grade its parent can carry.

34 Part VI in plain English: one line, and the cheap test that it is alone

A frequency alone is not a fingerprint. Any number of candidates could be arranged to sit at the same place, and a single line in a single scan would be evidence of something rather than of this.

What this Part offers instead is a fence, an address, and a count. The fence: the line must fall inside a narrow band of frequencies read from the geometry’s integers by an ansatz, around a pitch calibrated once and declared in advance. The address: within that band, the coupling must sit on one specific model line that experimenters already draw on their exclusion plots. The count: the line must be alone — the geometry supplies exactly one light field of this kind, so a second line in

the band would kill the count as this paper states it — and name which of the parents' registered modes is light — as surely as an empty band would.

What the Part offers is less decorative than a spectrum and harder to evade: one line, in one band, alone, checkable partly with data every scan already takes.

Part VII Detection

Principle, current experiments, and projections

35 The haloscope principle and the QGI readout

35.1 Resonant conversion

Sikivie’s haloscope (Sikivie, 1983) converts the portal of Part V into power. In a static field $\mathbf{B}_0$, the electromagnetic coupling turns the oscillating condensate $a_K(t) = \theta_0 f_K \cos(m_K t)$ into an effective current density oscillating at exactly $\nu_K = m_K/2\pi$, which drives any electromagnetic mode it overlaps. On resonance, a cavity of volume V , form factor C , and loaded quality factor Q_L extracts

$$P_{\text{sig}} = g_{K\gamma\gamma}^2 \frac{\rho_{\text{loc}}}{m_K} B_0^2 V C \min(Q_L, Q_a), \quad (35.1)$$

the structure worth memorizing: *one* factor of density (not two — the field is classical, Part III), one inverse mass, and the instrument’s entire art in $B_0^2 V C Q$. Numerically, for the laboratory eigenstate in a representative design cavity ($B_0 = 8$ T, $V = 100$ L, $C = 0.5$, $Q_L = 10^5$, $\rho_{\text{loc}} = 0.45$ GeV/cm³):

$$P_{\text{sig}} \simeq 2.4 \times 10^{-22} \text{ W} \times \left(\frac{g_{K\gamma\gamma}}{1.2 \times 10^{-15} \text{ GeV}^{-1}} \right)^2, \quad (35.2)$$

about two hundred microwave photons per second — pitifully small and entirely detectable, which is the standing miracle of the field. (The natural-to-SI unit conversion behind this number is worked once, step by step, in Appendix D.1, so that every power in the paper can be traced to it; the survey rows quote critical antenna coupling absorbed in Q_L , while the over-coupled KAXIO budget of Part X carries its antenna factor $\beta/(1 + \beta)$ explicitly.) Against it stands noise at the system temperature: the Dicke radiometer equation gives signal-to-noise $\text{SNR} = (P_{\text{sig}}/k_B T_{\text{sys}}) \sqrt{t/\Delta\nu_{\text{eff}}}$ with the matched filter’s effective width $\Delta\nu_{\text{eff}} = (\int g^2 d\delta)^{-1} \simeq 2.9$ kHz for the lab-frame template (Part IX; the condensate linewidth $\nu/Q_a \simeq 2.0$ kHz is the scale, not the filter). At the standard quantum limit $k_B T_{\text{sys}} \rightarrow h\nu$ (94 mK equivalent at 1.95 GHz), Eq. (35.2) yields 5σ in minutes per tuning step; the experimental cost is never one frequency but the *scan* — which is why a theory that names its frequencies in advance (Part V) changes the economics of the entire enterprise.

35.2 The haloscope as a measurement event of the series’ own theory

Before the engineering, one structural remark that places the experiment inside the framework it tests. The quantum-mechanics paper of the series (GG-A8-QM) derives measurement as a complete runtime event: switching on an apparatus reorganizes the joint system–apparatus landscape

(an endogenous, scheduled GRIP-Out — the laboratory instance of GG-A7-Qi4’s exploration-to-commitment switch); the registration that follows is GRIP-In descent into one pointer basin; and the stable record is the halt. A haloscope run is precisely this anatomy with magnets and amplifiers: energizing the cavity and its readout chain opens record basins that did not exist before (the GRIP-Out prologue); the condensate-driven cavity field relaxing into a recorded photon count is the GRIP-In; and each logged spectrum is a halt — permanent for the same reason every record in the series is permanent. The experiment that hunts the purest GRIP-In specimen in nature is itself a GRIP event chain; detector and detected run on one grammar.

35.3 The QGI readout

The series’ preferred formulation makes the fingerprint native. The cavity-plus-condensate system is a driven harmonic oscillator — the Quantum Geometric Interface (QGI) reading of the boundary law (GG-A5-GDOS T6, Langevin dress code): each receiver mode x_i obeys

$$\ddot{x}_i + \frac{\omega_i}{Q_L} \dot{x}_i + \omega_i^2 x_i = g_i \mathcal{A} \cos \omega_K t + \xi_i(t), \quad (35.3)$$

with one forcing frequency $\omega_K = 2\pi\nu_K$ inside the fenced window, g_i the photon coupling of the port, and ξ_i the thermal-quantum noise. Every identification test of this paper is a property of (35.3) that survives ignorance of $\mathcal{A}$: the location of the forcing inside the fence, and the absence of a second forcing anywhere the fence allows. The detector concept of Part VIII is a hardware realization of (35.3) read across the window.

35.4 The response, derived end to end

Equation (35.1) is usually quoted; a student should once see it earned. Four steps take us from the portal Lagrangian to the photon rate.

Step 1: the effective current. Varying the portal term $-\frac{g_{K\gamma\gamma}}{4} a_K F \tilde{F}$ together with the Maxwell action modifies Ampère’s law in a static external field $\mathbf{B}_0$:

$$\nabla \times \mathbf{B} - \partial_t \mathbf{E} = \mathbf{J}_a, \quad \mathbf{J}_a = g_{K\gamma\gamma} \mathbf{B}_0 \partial_t a_K, \quad (35.4)$$

an oscillating current density parallel to the magnet’s field, with amplitude $|\mathbf{J}_a| = g_{K\gamma\gamma} B_0 \sqrt{2\rho_{loc}}$ (using $\langle \dot{a}_K^2 \rangle = \rho_{loc}$). Everything a haloscope does is harvest the work done by (35.4).

Step 2: the mode equation. Expand the cavity field in eigenmodes, $\mathbf{E}(\mathbf{x}, t) = \sum_n e_n(t) \mathbf{E}_n(\mathbf{x})$. Projecting Maxwell’s equations onto mode n gives precisely the QGI oscillator (35.3), with the drive strength fixed by the overlap of the current with the mode — the form factor

$$C_n = \frac{|\int dV \mathbf{B}_0 \cdot \mathbf{E}_n|^2}{B_0^2 V \int dV |\mathbf{E}_n|^2}, \quad (35.5)$$

a pure geometry number ($C_{010} \simeq 0.69$ for the fundamental mode of an ideal cylinder, ~ 0.4 – 0.5 engineered).

Step 3: steady state on resonance. A damped oscillator driven at its natural frequency rings up until the drive’s work balances the dissipation: amplitude $\propto Q_L \times \text{drive}$, stored energy $U \propto Q_L^2$, extracted power $P = \omega U / Q_L \propto Q_L$. Carrying the constants through (and noting the ring-up saturates at the *condensate’s* coherence if the cavity is sharper than the line, hence $\min(Q_L, Q_a)$):

$$P_{\text{sig}} = g_{K\gamma\gamma}^2 \frac{\rho_{\text{loc}}}{m_K} B_0^2 V C \min(Q_L, Q_a), \quad (35.6)$$

which is (35.1). The single most useful reading: the $1/m_K$ is the coherence dividend — lighter condensates oscillate more slowly, letting the cavity integrate longer per quantum.

Step 4: photons and time. At the Kaxion benchmark, (35.6) with the representative parameters of Eq. (35.2) gives $P = 2.4 \times 10^{-22}$ W, i.e. a rate

$$\dot{N}_\gamma = \frac{P}{h\nu} = \frac{2.4 \times 10^{-22}}{(6.63 \times 10^{-34})(1.95 \times 10^9)} \simeq 1.8 \times 10^2 \text{ photons/s.}$$

Against quantum-limited noise ($k_B T_{\text{sys}} \rightarrow h\nu$), the matched-filter time to 5σ on the line is $t = (5 k_B T_{\text{sys}} / P)^2 \Delta\nu_{\text{eff}} \simeq 2.1$ s with the template’s effective width $\Delta\nu_{\text{eff}} = 2.9$ kHz (Part IX) — an idealized floor. Sweeping the full Kaxion window (544 MHz, 2.7×10^4 cavity bandwidths at $Q_L = 10^5$) then costs of order a day of idealized live time; real campaigns pay form-factor, duty-cycle, and tuning overheads of 10–100, landing at weeks to months — precisely the scale at which the current DFSZ campaigns operate. The conclusion deserves emphasis: *at DFSZ coupling, with a named window, the Kaxion search is a campaign, not an era.*

And the count test. The same transfer function quantifies the single-mode blindness asserted in Part VI: a receiver tuned to one frequency suppresses a hypothetical second line at a representative in-window detuning ($\Delta\nu/\nu = 2.5\%$, say) by $|\chi(\nu')/\chi(\nu)|^2 \simeq [2Q_L \Delta\nu/\nu]^{-2} = (2 \times 10^5 \times 0.025)^{-2} = 4 \times 10^{-8}$ — seventy-four decibels. For a scan that crosses the window the count test is a reanalysis of data already taken; for a camp it is a deliberate retune or a second mode, never a by-product of a single-mode camp; the transfer-function algebra behind this number is Appendix D.2.

36 Current and near-future experiments

36.1 The decisive simplification: the Kaxion is on the DFSZ line

Theorem 27.1 collapses the experimental question to one sentence: *the Kaxion’s photon coupling at any mass equals the DFSZ axion’s, so every DFSZ sensitivity statement in the haloscope literature applies to the Kaxion verbatim.* DFSZ is the community’s canonical hard target; three decades of instrument development have just, in the last few years, begun to reach it. The question “can anyone see the Kaxion?” therefore has a sharp answer: *yes — any experiment that demonstrates*

DFSZ reach inside a Kaxion window either discovers it or kills that window (with the count-test rider of Section 33: scans crossing the window should also re-bin for a second line).

36.2 The unique window: closing the two discrete questions

Every continuous quantity in the Kaxion sheet is fixed once the benchmark is calibrated (Part V). Two *discrete* questions could, in principle, have multiplied the prediction — the cosmological history of the compact angle, and the frame whose frequency a detector reads — and a less constrained framework would have been forced to quote a menu. The series closes the second question by theorem and carries the first as a named conditional (C-Kaxion-d), and the result is a single laboratory window:

$$\boxed{\nu_K \in [1.59, 2.13] \text{ GHz} \quad (\text{benchmark } 1.95 \text{ GHz}), \quad m_K = 6.6\text{--}8.8 \text{ } \mu\text{eV} \quad (\text{benchmark } 8.1)} \quad (36.1)$$

— one line, standing alone, in this window and nowhere else: *the Kaxion window*. This subsection derives the uniqueness by rejecting the two alternatives that would have moved the line — one by theorem, one conditionally.

36.2.1 The first rejected alternative: the randomized-angle history

A post-inflationary (randomized-angle) history — causal patches with independent values of θ , defect debris, the circle-averaged $\langle \theta^2 \rangle = \pi^2/3$ in place of the selected $\theta_i \simeq 1.1$ rad — makes more dark matter per unit decay constant and therefore closes the inventory at a *smaller* f_K and a heavier line: integrating the full cosine over a uniform angle, the closure moves to $f_K \lesssim 2.3 \times 10^{11}$ GeV and the mass to $\gtrsim 25 \mu\text{eV}$ ($\nu \gtrsim 6$ GHz) from misalignment alone, above the window — the familiar post-inflationary QCD-axion mass scale, which misalignment alone puts near $28 \mu\text{eV}$ (Borsanyi et al., 2016); the defect contribution that a randomized circle-valued field must carry raises it further and is not computed here (a bulk-tension defect network would itself be excluded). The hypothesis has a physical prerequisite: the angle must be *randomized*, and there are exactly two known mechanisms — thermal restoration ($T \gtrsim f_K$) or de Sitter fluctuation ($H_{\text{inf}}/2\pi \gtrsim f_K$). The Kaxion forecloses the first structurally and numerically; the second is the conditional.

The structural exclusion of thermal restoration: no Peccei–Quinn scalar to restore.

For a PQ axion the angle is the phase of a complex scalar $\Phi = \rho e^{i\theta}$: above $T \sim f_a$ the radial mode restores ($\langle \rho \rangle = 0$), the phase loses meaning, and on cooling each Hubble patch picks an independent θ , with cosmic strings (windings of Φ) as the inevitable debris. *Every step of that story requires the radial mode*. The Kaxion has none: by Theorem 8.2 its directions are the dual of a forced two-form and the holonomies of a rigid circle — gauge-topological angles whose “decay constant” is a geometric normalization (Appendix C.1), not the vacuum value of any order parameter. There is no field whose thermal restoration could erase the angle at any temperature below the bulk scale, and no thermal epoch in which separate patches could acquire separate dials by that route. This is the structural property model-independent (string-type) axions share — no radial mode to restore — and for the Kaxion it is a corollary of the existence theorem: no scalar was ever added. It says

nothing about de Sitter randomization, which needs no radial mode and is the separate question of C-Kaxion-d.

Physical picture. A PQ axion’s angle lives on a dial whose face can be melted and recast — heat the radial field and the markings vanish; cool it and every neighborhood paints its own. The Kaxion’s dial is machined into the geometry. The hands can be turned (the misalignment of Part III); the face can never be melted — but inflation could have spun the hands, which is exactly what C-Kaxion-d says it did not. One universe, one face, and, under the conditional, one initial angle.

The quantitative side: what the numbers do and do not exclude. Thermal restoration is excluded by the numbers and by the structure together: the reheating temperature after primary inflation is $T_{\text{rh}} \simeq 6 \times 10^{10}$ GeV (GG-A10-Cosmos), a factor $\simeq 12$ below f_K , and there is in any case no radial mode to restore. De Sitter randomization is *not* excluded by the numbers: the primary inflation of GG-A10-Cosmos has $H_1 \simeq 1.5 \times 10^{13}$ GeV (from $r = 0.0038$ and $A_s = 2.1 \times 10^{-9}$, as GG-A10-Cosmos now prints), so a light angle would receive $H_1/(2\pi f_K) \simeq 3.4$ rad per e-fold and be randomized — which is why the single-patch history rides on the inflation-era pinning, Conjecture C-Kaxion-d of Section 13.5. The lock crossing at which the angle is deposited is prompt and low-curvature; writing H_2 for its curvature scale,

$$H_2 \simeq \frac{M_{\text{lock}}^2}{\sqrt{3} \bar{M}_{\text{Pl}}} = \frac{(3.85 \times 10^3)^2}{1.73 \times 2.44 \times 10^{18}} \text{ GeV} \simeq 3.5 \times 10^{-12} \text{ GeV}, \quad (36.2)$$

so that nothing at the release epoch itself can randomize the angle,

$$\frac{H_2}{2\pi f_K} \simeq \frac{5.6 \times 10^{-13}}{7 \times 10^{11}} \simeq 8 \times 10^{-25}, \quad (36.3)$$

and the lock-crossing temperature — the temperature at which the radiation Hubble rate equals the lock basin’s curvature H_2 (not a reheating temperature; there is no second inflation), and the T_{lock} of Part III — is

$$T_{\text{lock}} \simeq \left(\frac{90}{\pi^2 g_*} \right)^{1/4} \sqrt{H_2 \bar{M}_{\text{Pl}}} \simeq 1.6 \text{ TeV}, \quad (36.4)$$

far below any scale at which the angle could be disturbed. (GG-A10-Cosmos’s timeline labels the same prompt crossing by the bath temperature at which it begins, $T = M_{\text{lock}} \simeq 3.85$ TeV, $t \simeq 1.6 \times 10^{-14}$ s; the two labels bracket one event.) What these margins establish is that the release epoch is quiet. What they do not establish is the single-patch history itself: that requires the angle to have been pinned through primary inflation (C-Kaxion-d), and the history is carried as a named conditional — assumed in the ledger, killable by a minicluster population or an isocurvature detection — not as forced. Under the conditional, the misalignment angle enters the abundance exactly as Part V used it, and the randomized-angle history is a rejected hypothesis; without it, the randomized history is the one the numbers select, and it puts the line at $\gtrsim 6$ GHz, outside the window. (One corollary for Part IV: the minicluster phenomenology of the randomized history is

off under the conditional, and is fenced there as the contingency that a failure of C-Kaxion-d would make live.)

36.2.2 Why the line lies in the gigahertz window

The remaining question is where in the gigahertz band the line lies — and in particular why it sits there rather than at the much lower frequency the bulk normalization scale alone would suggest. The physical decay constant is that of the canonically normalized light eigenstate, $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$; the only way the line could sit near the bulk scale is the limit $\Xi_{\text{cap}} = 1$, and the series closes that loophole twice over:

(i) *The norms are exponentially asymmetric.* The component reduction of Appendix C.1 gives the tensor’s kinetic norm as πe^{2K} (infrared-weighted) against the order-one π of the Wilson line’s integral (ultraviolet-weighted) — a mismatch of $e^{2K} = e^{70} \simeq 2.5 \times 10^{30}$. So asymmetric a system does not mix its directions — the hierarchy decouples them, with any mixing angle suppressed by e^{-K} — and the canonical normalization of the light field is far from the bulk normalization for that reason: “ $\Xi_{\text{cap}} = 1$ ” would require the infrared-enhanced and order-one norms to coincide, a statement off by thirty orders of magnitude in the very integrals that define the two normalizations.

(ii) *Which norm, and at what value, is open.* Which chirality the propagating remnant carries is the orientation convention of O14, and the value of the ratio between the bulk and canonical normalizations is what the ansatz of §26.3 parametrizes as Ξ_{cap} — benchmark $K+N+L = 43$ among the seven loadings $\Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$, the benchmark calibrated through the closure. The line sits in the gigahertz window because the calibration puts it there; the geometry’s contribution is that it cannot sit at the bulk scale.

Physical picture. A bare field and its physical normalization can be written in any convention — but a microphone in the room hears the canonical frequency, never the bookkeeping one. The physical line is the eigenfrequency of the canonically normalized field — not at the bulk normalization scale, and, at the calibrated benchmark, in the gigahertz window.

36.2.3 Uniqueness, and the scorecard

Both alternatives rejected, the prediction is singular: one window, Eq. (36.1), whose residual width is the one unfinished geometric item of the program — the absolute normalization (O14) and the defining computation for the cap loading — bracketed, under the ladder ansatz, by a discrete OGN-5 window rather than by a continuum. The uniqueness is itself falsifiable twice over: exclusion of the window kills the calibrated placement and its ladder (F-Kaxion-1c), and so does a confirmed compact-angle-like line *anywhere else* (F-Kaxion-1d), a line at the randomized-angle mass killing the pinning conditional instead — while a second line inside the window kills the count and names a registered mode. This is the feature no adjustable competitor can reproduce:

The PQ axion’s menu can never collapse, because its forks are free model choices — no calculation, even in principle, decides KSVZ against DFSZ or pre against post. One of the Kaxion’s two questions had a calculable answer and the calculation is done; the other is a named conjecture with a named killer. The haloscope program of Part VIII is therefore not scanning a parameter space; it is

Table 14: Why no adjustable competitor can reproduce the signature. For each model the table counts the continuous dials it carries, the discrete forks it must choose among, and whether the pair together fix a fenced window with a marked rung. The Kaxion’s row is the only one whose answer to the last column is yes, and that is the whole content of the claim: a candidate with a continuous dial predicts a region, and a region cannot be missed in the way a fenced window can. The falsifiers of this section are stated against the window.

Model	Continuous dials	Discrete forks	Fenced window?
WIMP	m_χ, σ (decades each)	—	no: a plane, not a window
PQ axion	f_a (10 orders), θ_i	pre/post; E/N	no: forks are free choices
Generic ALP	m, g unconstrained	—	no
Kaxion	one (calibrated)	two: closed, cond.	yes, cond.: 1.59–2.13 GHz (bench. 1.95)

sweeping a pre-registered window with a marked rung, expecting one line standing alone and silence elsewhere.

36.3 The experimental landscape

All comparisons below are made against the unique window (36.1); for orientation in the broader experimental program see the community reviews (Adams et al., 2023; Irastorza & Redondo, 2018). Table 15 surveys the field row by row — band, demonstrated depth, and standing relative to the window — and repays a careful read: the rows above the window show DFSZ depth already achieved, the rows at high mass show KSVZ-only depth, and exactly three rows (CAPP, ADMX-EFR, HAYSTAC) matter enough to receive the deep dives that follow.

Table 15: The experimental landscape against the Kaxion window. Sensitivities are quoted relative to the DFSZ line (= the Kaxion line); extents are from the cited publications. Nothing in this table is this paper’s work. Every row is what a collaboration has published or announced, quoted at the depth they themselves quote.

Experiment	Mass coverage	Depth	Relation to the Kaxion windows
<i>Published inside the window</i>			
IBS-CAPP (8TB) (Lee et al., 2020)	6.62–6.82 μeV (1.60–1.65 GHz)	$\simeq 4\times\text{KSVZ}$	$\simeq 10.2\times\text{DFSZ}$; the only modern search inside the window, covering 50 of its 544 MHz.
UF (Hagmann et al., 1990)	5.6–7.0, 8.1 μeV	7.6– $\sim 30\times\text{DFSZ}$	pilot; its gap at 1.70–1.84 GHz holds the $\Xi_{\text{cap}} = 38, 40$ rungs.
RBF (Wuensch et al., 1989)	4.5–10.2 μeV	$\sim 10^2\times\text{DFSZ}$	pilot experiment; all of the window.
<i>Announced to cover the window</i>			
ADMX four-cavity array (Yang et al., 2020)	5–9 μeV (1.21–2.18 GHz)	(1.21–none stated)	its band contains all seven addresses; a design and prototype study stating neither a sensitivity nor a start.
<i>The DFSZ front, and the rest of the dial</i>			
ADMX (ADMX Collaboration, 2025)	2.7–4.2 μeV	DFSZ	DFSZ demonstrated; below the window — the technology benchmark.
ADMX 1.1–1.3 GHz (ADMX Collaboration, 2025b)	4.54–5.41 μeV	better than KSVZ	adjacent to the window’s lower edge; DFSZ not reached in that run.
CAPP-12TB (CAPP Collaboration, 2024)	4.51–4.59 μeV	DFSZ	DFSZ at a second site; window-relevant technology (1–2 GHz class).
IBS-CAPP (MAX) (Kim et al., 2024)	4.24–4.91 μeV (1.025–1.185 GHz)	at or near DFSZ	the deepest search above 1 GHz; below the window.
HAYSTAC (HAYSTAC Collaboration, 2025)	17.0–19.5 μeV	$2.9\times\text{KSVZ}$	above the window in mass; squeezed-state readout pioneers sub-SQL noise.
QUAX (Crescini et al., 2020)	$\simeq 43 \mu\text{eV}$	$\sim\text{KSVZ}$	above the window; high-frequency pathfinder.
ORGAN	63–67 μeV	$\sim 10^2\times\text{KSVZ}$	far above the window in mass and well above the QCD band in depth.
ADMX SLIC (Crisosto et al., 2020)	0.176–0.203 μeV	$\sim 10^5\times\text{DFSZ}$	historic sub- μeV pathfinder, far below the window.
DMRadio- m^3 (DMRadio Collaboration, 2022)	$\lesssim 0.83 \mu\text{eV}$ (lumped)	DFSZ goal	sub- μeV program, below the window.
ADMX-EFR (ADMX-EFR Collaboration, 2024)	8.3–16.5 μeV (2–4 GHz)	DFSZ goal	first DFSZ program overlapping the window’s upper edge.
MADMAX (MADMAX Collaboration, 2024); ALPHA (Lawson et al., 2019)	40–400; 100 μeV	$\sim 20\times\text{KSVZ}$ –DFSZ	above the window; dielectric/plasma concepts.
CAST (CAST Collaboration, 2024)	solar, all masses	$g < 6\times 10^{-11}$	contextual ceiling, nearly five orders above the Kaxion line.

36.4 IBS-CAPP: the best-situated instrument

Status. The Center for Axion and Precision Physics (IBS-CAPP, Korea) operates the most advanced axion-haloscope facility now running. Its flagship CAPP-MAX — a 12 T magnet, a 37-liter ultralight copper cavity, a dilution refrigerator holding the cold mass below 40 mK, flux-driven




Figure 8: The Kaxion against the haloscope landscape in the $(m_K, |g_{K\gamma\gamma}|)$ plane. Red regions: the *published* exclusion envelopes of the ADMX, CAPP (including CAPP-MAX), HAYSTAC, TASEH, CAST-CAPP, GrAHal, QUAX and ORGAN campaigns and the historical RBF/UF scans, digitized from the community AxionLimits compilation (O’Hare, 2020) (MIT license, concept DOI 10.5281/zenodo.3932429) with lower-envelope decimation; vertical strokes are single-frequency results. The Kaxion line coincides with the DFSZ line (Theorem 27.1); the single predicted window (36.1) — $\nu_K \in [1.59, 2.13]$ GHz, benchmark 1.95 GHz — is the vertical needle-band, with the discrete Ξ_{cap} bracket drawn on the line and the $\pm 1.3\%$ (m), $\pm 5.5\%$ (g) benchmark errors of Part V at fixed f_K . The black star marks the 1.036 GHz candidate faced and excluded in early 2026 (§37.4) — outside the window and brighter than the DFSZ line, its confirmation would have falsified the Kaxion (F-Kaxion-1d). Three published searches reach inside the window and are labeled: IBS-CAPP’s 8 T run over 1.60–1.65 GHz at $\simeq 4 \times \text{KSVZ}$, and the historical UF and RBF scans an order of magnitude above that. *Announced programs* are drawn dashed and only as far as each states (§36.4–36.5): ADMX-EFR on the DFSZ line it targets across 2–4 GHz, and ADMX’s four-cavity array — a stated mass range, $5-9 \mu$ eV, with no stated sensitivity — as a coverage bracket above the plane. No DFSZ-depth search has been published anywhere in the window: the deepest inside it falls a factor of ten short. Territory adjacent, exclusion no.

Josephson parametric amplifiers, and a total system noise $\lesssim 200$ mK — delivered the first axion search *at or near DFSZ sensitivity over more than 100 MHz above 1 GHz, with a record DFSZ scan speed of $\simeq 3$ MHz/day* (Kim et al., 2024). Decisively for the Kaxion, CAPP has already operated *inside* the window: its 8 T / 165 mm solenoid swept 1.60–1.65 GHz and reached the QCD band at $\simeq 4\times$ the KSVZ coupling (Lee et al., 2020), which at these masses is 10.2 times the DFSZ line. The band and the depth have not yet been combined: the DFSZ demonstration is on the 12 T / 320 mm magnet near 1 GHz, and the in-window running was a factor of ten above the line the Kaxion sits on. The published roadmap is to cover 1–8 GHz at DFSZ within about five years and to extend toward 25 GHz thereafter, supported by high-temperature-superconductor cavities that hold a high Q in-field and by 18–25 T HTS magnet development. The scan front is moving toward the window from below: multi-chip Josephson-parametric-amplifier modules now deliver near-quantum-limited gain over 1.2–1.5 GHz, and in early 2026 the group completed a textbook candidate-validation campaign — following up, and excluding, a 5.1σ -local excess near 1.036 GHz (Ahn et al., 2026) — demonstrating at discovery-level sensitivity precisely the vetting discipline the Kaxion program requires (Section 37).

Match to the window. An 8 T solenoid has already held a cavity at the window’s lower edge, and the 12 T / 320 mm magnet that carries CAPP’s DFSZ demonstration is a retune away from the window’s body. What the window still wants is *depth*: the deepest search ever performed inside it falls a factor of ten short of the DFSZ line, over 50 of the window’s 544 MHz, and no DFSZ-depth search has been published anywhere in the window. CAPP is the only group that has both operated in this band and demonstrated DFSZ depth above 1 GHz — in two different runs, on two different magnets.

Does it need to upgrade? For a *generic* DFSZ detection at 1.95 GHz: essentially no new detector — a CAPP-MAX-class cavity tuned into 1.8–2.1 GHz, plus continued running. (The signal power falls and the scan slows with cavity volume as the frequency rises, so 2 GHz is more expensive per megahertz than 1 GHz, but it is squarely within the existing architecture and roadmap.) For *Kaxion identification* — the fence, the address, and the count — it needs only analysis and coverage: complete the window at DFSZ depth, and run the count test of Part VI on the accumulated spectra (a second instrumented mode, Part VIII, runs it live, and is a modest retrofit CAPP’s cryogenics and quantum-limited readout already support).

Concrete proposal (increasing cost). (i) A *count-test re-binning* of any CAPP scan that crosses [1.59, 2.13] GHz — search the in-window spectra for a second line beside any candidate — at *zero* hardware cost. (ii) Prioritize the benchmark rung at 1.95 GHz — a retune of a CAPP-MAX-class cavity — and from it the highest-value 544 MHz in the Kaxion program. (iii) Retrofit a second instrumented mode for the live count test (Part VIII). Steps (i)–(ii) are additive to CAPP’s funded program at negligible marginal cost; (iii) is the identification upgrade a partner group is positioned to co-develop.

Timescale to $> 5\sigma$. At the demonstrated DFSZ scan speed (slower at 2 GHz than the $\simeq 3$ MHz/day record near 1 GHz), a dedicated DFSZ-depth sweep of the 544 MHz window is a ~ 5 -year campaign for a single cavity (the demonstrated rate scaled as ν^{-4} to 2 GHz gives ~ 0.3 MHz/day), and a 1–2-year one only for a multi-cavity array of the EFR class; CAPP’s stated

1–8 GHz roadmap contains the window, though no published schedule says when it is reached. Because the Kaxion sits on the DFSZ line, a $> 5\sigma$ excess at the benchmark is then a matter of integration time, not of a sensitivity breakthrough. Identification (the fence, the address, and the count together) follows with the analysis and coverage, late-2020s to early-2030s.

36.5 ADMX-EFR: the backstop at the upper edge

Status. The parent instrument, ADMX, began a new run in January 2026, pushing toward DFSZ at below-nominal halo density through lower system noise, and is developing a four-cavity coherently-combined system for the $5\text{--}9\ \mu\text{eV}$ range (Yang et al., 2020). Its own published run report points the other way: ADMX writes that its coming runs will probe *downward* toward 1 GHz (ADMX Collaboration, 2025b), so the main instrument is not approaching this window, and the four-cavity array is the only ADMX route into it. The ADMX Extended Frequency Range (ADMX-EFR) is the next-generation US Department of Energy Office of Science (HEP) project — the Generation-2 dark-matter line — whose dedicated magnet is in hand at Fermilab’s Dark Wave Laboratory, with the experiment in the design-and-prototyping phase targeting operations in the late 2020s. It is an array of **eighteen** cavities with coherent power combining on a 9.4 T MRI magnet, designed to reach **DFSZ sensitivity across 2–4 GHz** ($8.3\text{--}16.5\ \mu\text{eV}$) in a three-year campaign, with a scan speed $\sim 5\times$ current ADMX from the larger field, higher- Q (superconducting-film or dielectric) cavities, and reduced amplifier noise via squeezed states or single-photon counting (ADMX-EFR Collaboration, 2024).

Match to the window. ADMX-EFR’s design floor is 2 GHz, so it cleanly covers only the *top* of the Kaxion window, 2.0–2.13 GHz ($8.3\text{--}8.8\ \mu\text{eV}$); the benchmark 1.95 GHz/ $8.1\ \mu\text{eV}$ sits just below its floor, and the body of the window (1.59–2.0 GHz) is out of its design band. ADMX-EFR is thus the *backstop* that would catch the window’s upper edge, complementary to CAPP covering the body.

Does it need to upgrade? To cover only the overlap (2.0–2.13 GHz): no change — it is in the design band. To reach the *benchmark* 1.95 GHz, ADMX-EFR would extend its lowest cavity $\sim 50\text{--}100$ MHz below 2 GHz, a cavity-size adjustment within the 9.4 T bore (the bottom element of the array runs slightly larger). Identification again needs only coverage and the count-test analysis; the EFR array’s coherent combining and photon-counting readout are well suited to the faint DFSZ signal.

Concrete proposal. (i) Prioritize the *bottom* of the EFR band (2.0–2.13 GHz) in early operation — it is the Kaxion overlap and costs nothing extra. (ii) Carry a *count-aware analysis*: any in-window second line beside a candidate falls inside the EFR band, so the count test is available within EFR’s own data. (iii) Engineer the lowest cavity to extend down to 1.95 GHz, bringing the benchmark into reach. Steps (i)–(ii) cost only analysis and run-plan priority.

Timescale to $> 5\sigma$. Once operations begin — the start awaits a funding decision — the 2.0–2.13 GHz overlap is covered at DFSZ in the first $\sim 1\text{--}2$ years of EFR running; the benchmark requires the lower-edge extension on the same timescale. As with CAPP, DFSZ reach equals Kaxion reach.

36.6 HAYSTAC: the high-mass frontier, and a technology partner

Status. HAYSTAC (Yale/Berkeley/Colorado) pioneers sub-quantum-limited readout: its squeezed-state receiver evades the standard quantum limit and accelerates the scan by $2\text{--}3\times$. Its 2025 Phase-II results scanned $16.96\text{--}19.46\ \mu\text{eV}$ ($\sim 4.1\text{--}4.7\ \text{GHz}$) to $\sim 2.75\text{--}2.96\times\text{KSVZ}$ (HAYSTAC Collaboration, 2025).

Relation to the window. HAYSTAC operates *above* the Kaxion window in mass and at KSVZ — not DFSZ — depth, so it is **not** a discovery instrument for the benchmark. Two precise remarks. First, HAYSTAC’s mass range is exactly where the abundance-frame (bookkeeping) reading would once have placed a higher line; the canonically-normalized eigenstate of Part V is the single physical line at $8.1\ \mu\text{eV}$, below HAYSTAC’s band — so HAYSTAC’s null at $17\text{--}19.5\ \mu\text{eV}$ does not constrain the window. Second, HAYSTAC’s squeezed-state and photon-counting technology is precisely what the Kaxion identification needs; it is the natural *technology* collaborator for the count-test receiver’s readout even though its magnet sits at the wrong frequency.

36.7 The near-term targeting plan

The three deep-dives collapse to one strategy, fixed by the decisive simplification of §36: *the Kaxion is on the DFSZ line, so any DFSZ reach inside $[1.59, 2.13]\ \text{GHz}$ discovers it or kills the window*. The plan follows the instruments’ own roadmaps and asks of them, in increasing cost:

1. **Count-test analysis (zero hardware).** Reanalyze every existing and planned scan crossing the window for a second line anywhere inside it, asking of each candidate whether it stands alone. This converts ordinary scan data into a Kaxion identification test at no instrumental cost, at both CAPP and ADMX-EFR.
2. **Window priority.** Direct a retuned CAPP cavity at the benchmark rung ($1.95\ \text{GHz}$) and ADMX-EFR’s early operation at its band bottom ($2.0\text{--}2.13\ \text{GHz}$). ADMX’s announced four-cavity array, designed for $5\text{--}9\ \mu\text{eV}$ (Yang *et al.*, Springer Proc. Phys. **245**, 53 (2020)), would cover all seven addresses, with EFR reaching the topmost two from above — between them the whole window, on a schedule neither has published.
3. **Identification retrofit.** A second instrumented mode for the live count test (Part VIII) — the modest upgrade that turns a discovery into an identification.

The projected DFSZ-depth coverage is the CAPP and ADMX-EFR bands of Figure 8: on the collaborations’ stated (and somewhat optimistically extrapolated) roadmaps, both cross the Kaxion line inside the window between ~ 2030 and ~ 2032 — CAPP’s date on the assumption of a multi-cavity array of the EFR class; a single CAPP-MAX-class cavity at $\sim 0.3\ \text{MHz/day}$, entering in 2027–2029, completes the window in 2032–2034 (§36.4). The primary approach is IBS-CAPP (which has operated in this band and demonstrated the depth, in two different runs); the US backstop is ADMX-EFR (the upper edge), whose start awaits a funding decision; HAYSTAC is a readout-technology partner. In every case the dedicated contribution — the count-test analysis and the multi-mode retrofit — is *additive* to a funded program rather than a competing build.

36.8 The bottom line

Three statements summarize Figure 8, each with its receipt. *First*, the window is not excluded: published at-or-near-DFSZ coverage ends near $5.4\,\mu\text{eV}$ (ADMX’s 1.1–1.3 GHz run; CAPP-MAX), below the window, and inside the window three searches have been published — IBS-CAPP’s 8 T run at $10.2\times$ the Kaxion line over 50 of the window’s 544 MHz, and the historical RBF and UF pilots at some 60–140 and 30–35 times it — so the floor inside the window stands an order of magnitude above the prediction at best: territory adjacent, exclusion no. *Second*, the window is bracketed: the announced DFSZ program at 1–4 GHz — CAPP-class instruments from below, ADMX-EFR from above — converges on exactly $[1.59, 2.13]$ GHz, so the Kaxion will be found, or killed, by instruments already funded or proposed. *Third*, the identification hardware is ready: a multi-mode retrofit turns any cavity in the band into a count-test instrument at negligible cost. The strategic conclusion for a dedicated program (Part VIII): strike the window with a count-aware instrument riding the 2030 DFSZ program.

36.9 The discovery calendar: F-Kaxion-1c as dates

Kill criteria without dates are rhetoric, so we lay each instrument’s route into the window beside its stated roadmap and read off what those roadmaps would imply. One caution belongs at the top of the table rather than under it: what follows is *this paper’s extrapolation from published statements*, not a schedule anyone has announced. No program has published a schedule that covers this window at DFSZ depth. One roadmap contains it — IBS-CAPP’s 1–8 GHz program — as a program-level range that neither names this band nor attaches a campaign to it; of the two designs aimed at the window itself, the array designed for it states neither a sensitivity nor a start, and the program that states DFSZ depth reaches only the window’s top and waits on a funding decision.

Table 16: The routes into the window. Stated roadmaps as of July 2026 (§36.4–36.5); the camping row is the pre-registered option of §37, available to any listed instrument at any time. Dates move, and the table is a snapshot rather than a commitment by anyone named in it. A single CAPP-MAX-class cavity at ~ 0.3 MHz/day completes the window about five years after entry (§36.4); the 2030–2032 completion read below assumes an array of the EFR class on the lower half, and reads 2032–2034 for a single cavity.

Instrument	Route into $[1.59, 2.13]$ GHz	DFSZ depth	Window entry
CAPP-MAX (JPA upgrades)	up from 1.5 GHz	demonstrated	2027–29; depth $\gtrsim 2030$
ADMX four-cavity array	$5\text{--}9\,\mu\text{eV}$ by design	none stated	not scheduled
ADMX-EFR	2–4 GHz (upper edge)	design	$\sim 2029\text{--}2032$
any row, camped on 1.95 GHz	pre-registered	yes, in days	any time

Reading the table as dates: the scan fronts reach the window’s lower edge around 2027–2029 (CAPP from 1.5 GHz), the EFR campaign closes the upper edge around 2030–2032, and a full DFSZ-depth, count-aware sweep of the window is therefore plausibly *complete by $\sim 2030\text{--}2032$* on funded and stated programs if an array of the EFR class takes the lower half, and by 2032–2034 with a single cavity there — with the camping option standing open at every moment in between.

That is the kill date this paper accepts in advance, on the explicit understanding that the date is our arithmetic on other people’s roadmaps and not a commitment by anyone named in the table: if, by the early 2030s, the window has been swept at DFSZ depth with a count-aware analysis and no line has appeared, F-Kaxion-1c has fired: the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d are retired, and the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinned branch or the relic is not this sector’s — not a kill of the framework, and this paper does not claim one. A falsifier whose firing waits on a sweep nobody has scheduled is a falsifier in principle, and we would rather say so than let the table imply otherwise. Part VIII closes by drawing this same calendar as the cohort’s kill-test timeline (Figure 10), with the one test already faced marked at its date.

37 The pre-registered single-frequency campaign: camping on the frequency

Everything in this Part so far accepts the field’s default posture: the haloscope *scans*, because a generic axion could sit anywhere over decades of frequency. The Kaxion changes the posture. The theory names its frequency in advance — the benchmark 1.95 GHz, inside the closed window (36.1) — so the search can be run the way a confirmation experiment is run: park the resonator on the predicted line and integrate. We call this *camping on the frequency*, and this section quantifies what it buys. The answer is twofold: the look-elsewhere penalty vanishes, and the scan — the only expensive part of a DFSZ campaign — is deleted. Both factors are arithmetic, not hardware; they are available to every instrument in the band on the day it chooses to use them.

37.1 The trials factor collapses to the window

A scanning search pays for its breadth in statistics. Over the condensate linewidth $\Delta\nu_a \simeq 2.0$ kHz, a blind hunt across 1–10 GHz interrogates $N \simeq 4.6 \times 10^6$ independent frequency bins at a fixed 2.0 kHz width (counting at the ν -dependent linewidth ν/Q_a gives $N \approx 2\text{--}5 \times 10^6$; the requirement below is $7.3\text{--}7.4\sigma$ either way), and the significance of any single excess must be discounted by that trials factor — the look-elsewhere effect. The arithmetic is unforgiving: a local 5σ excess in such a scan carries a global false-alarm probability of order unity, and discovery-grade *global* 5σ requires local 7.4σ — a factor 2.2 more integration time everywhere, paid on every bin of the scan. Pre-registering the Kaxion window (2.8×10^5 bins) demands local 7.0σ — a modest gain over the blind scan, and the honest one, because the window is the prediction. The benchmark rung is narrower — ± 25 MHz at fixed f_K , 2.6×10^4 bins, local 6.7σ — but it is a calibrated row, not a derived point, and the paper does not claim the one-bin collapse that a derived frequency would earn: local significance is not global significance here, and the difference is stated on the search card.

The field has just supplied the textbook illustration. In early 2026 the IBS-CAPP group reported the recovery of a previously unanalyzable data segment in which an excess at 1.036315 GHz met every candidate-selection criterion — local significance 5.1σ , spectral profile consistent with a virialized halo line, signal strength tracking the cavity response (Ahn et al., 2026). After the look-elsewhere discount the same excess was worth only 3.5σ globally: an anomaly, not an answer. Had that

frequency been named in advance, the identical dataset would have carried discovery-grade *global* significance — and would then have faced the same vetting that in fact killed it. That is the entire content of pre-registration, demonstrated on live data: *the difference between an anomaly and an answer is whether the frequency was written down first*. This paper writes the Kaxion’s window down first, and its benchmark with it.

37.2 Signal-to-noise at the benchmark, instrument by instrument

What does camping cost in time? Table 17 evaluates the paper’s registered signal and noise machinery (Eqs. (35.1) and the Dicke radiometer equation; Appendix D.1) at the benchmark 1.95 GHz and DFSZ coupling $g_{K\gamma\gamma} = 1.2 \times 10^{-15} \text{ GeV}^{-1}$, for three instrument classes. The volume entries are honest: a single CAPP-MAX-class cavity retuned from its published ~ 1.1 GHz operating point to 1.95 GHz shrinks as $V \propto \nu^{-3}$, from 37 to ~ 6 liters — the high-frequency volume penalty that multi-cavity arrays exist to repair, and repair well: the eighteen-cavity EFR-class array restores ~ 200 liters at the window.

Table 17: Camping on the benchmark. The registered Sikivie power (35.1) and matched-filter time to 5σ on the lab-frame template (effective width $\Delta\nu_{\text{eff}} \simeq 2.9$ kHz, Part IX), at 1.95 GHz, DFSZ coupling, $\rho_{\text{loc}} = 0.45 \text{ GeV/cm}^3$, C and $Q_L = 10^5$ as listed, critical antenna coupling absorbed in Q_L (App. D.1; signal photon rates 3×10^1 – 4×10^2 per second across the rows). Idealized live time; the 10–100 $\times$ duty-cycle and calibration overheads apply to *scanning* campaigns, while a camped single-bin run pays only its calibration share (the KAXIO run plan of Part X budgets duty explicitly). The last row shows a deliberately modest receiver for contrast: even a 4-kelvin semiconductor amplifier, camped on the named bin, reaches 5σ within a day.

Instrument class	B_0	V	C	T_{sys}	P_{sig}	$t_{5\sigma}/\text{bin}$
Representative design cavity (§35)	8 T	100 L	0.5	94 mK (SQL)	$2.4 \times 10^{-22} \text{ W}$	2.1 s
CAPP-MAX-class single cavity at 1.95 GHz	12 T	6.4 L	0.6	200 mK	$4.1 \times 10^{-23} \text{ W}$	5.4 min
EFR-class 18-cavity array	9.4 T	200 L	0.4	150 mK	$5.3 \times 10^{-22} \text{ W}$	1.1 s
CAPP-MAX-class cavity, 4 K HEMT receiver	12 T	6.4 L	0.6	4 K	$4.1 \times 10^{-23} \text{ W}$	36 h




Figure 9: SNR versus integration time, camped on the benchmark. Dicke radiometer growth $\text{SNR} \propto \sqrt{t}$ at 1.95 GHz and DFSZ coupling for the instrument classes of Table 17 (idealized live time; matched filter on the lab-frame template, effective width $\Delta\nu_{\text{eff}} = 2.9$ kHz, Part IX). Every curve crosses discovery-grade 5σ within seconds to minutes — even the deliberately modest 4-kelvin receiver crosses within a day. At DFSZ coupling, with the frequency named in advance, integration time was never the cost of the Kaxion search; the scan was. Camping deletes the scan.

Three lessons. *First*, at the named bin every quantum-limited instrument in the band reaches 5σ in seconds to minutes of live time; with the scanning-campaign overheads of 10–100, minutes to hours. *Second*, the honest cost sits in the window, not in the benchmark bin: the prediction is the 544 MHz window with a marked benchmark, and a CAPP-MAX-class cavity covers the window at DFSZ depth, at one 5σ -local dwell per cavity bandwidth, in ~ 100 idealized days, which the scanning overheads of 10–100 turn into the multi-year single-cavity campaign of §36.4; an EFR-class array covers the entire window in about eight idealized hours, days to weeks of calendar time. The benchmark rung itself is ± 25 MHz wide at fixed f_K and is where a campaign should begin; no lattice improvement narrows the window, because the window is the geometry’s residual, not QCD’s. (These band-coverage times are 5σ -local dwell; claiming a discovery anywhere in the window at *global* 5σ costs local 7.0σ , a factor $\times 2$ in dwell, §37.1.) *Third*, the count test costs no additional exposure on a scan: every listed instrument’s tuning range spans a useful stretch of the window, so for any instrument that crosses the window the identification test of Part VI is a reanalysis of data already taken. A camp must ask for it (a retune or a second mode, §49).

37.3 The campaign protocol

The pre-registered campaign, stated as a protocol a collaboration could adopt on Monday morning:

1. **Pre-registration.** The window [1.59, 2.13] GHz, the benchmark 1.95 GHz, the single-tone template, and the count test are fixed in advance by this paper and its falsifier registry (Appendix E.2); the trials factor of §37.1 is thereby $\leq 2.6 \times 10^4$ bins for the benchmark rung

(local 6.7σ for global 5σ) and $\leq 2.8 \times 10^5$ for the window (local 7.0σ) — both declared before the data.

2. **Benchmark first, then the window.** Camp on 1.95 GHz and its ± 25 MHz lattice width — the highest-value 50 MHz in the field — then sweep the full window $[1.59, 2.13]$ GHz at DFSZ depth, marking the rungs of the discrete Ξ_{cap} ladder of Part V as it passes them. The window is finite; the campaign is enumerable.
3. **Count-aware from day one.** Analyze every spectrum for a second in-window line beside any candidate. A candidate anywhere in the band is confirmed or killed within days: the count test and the persistence tests (field-off, cross-instrument, re-scan) settle it, precisely the validation discipline of Ahn et al. (2026).
4. **The count test on standby.** A second instrumented mode (Part VIII) is a software-and-synthesis add-on; running it alongside the camp turns a confirmed line into a verified count — identification, not detection.

The Kaxion Search Card — everything a haloscope collaboration needs, on one page

Target. benchmark $\nu_K = 1.95$ GHz (± 25 MHz at fixed f_K ; $m_K = 8.1 \mu\text{eV}$); closed window $\nu_K \in [1.59, 2.13]$ GHz, the prediction proper; discrete ladder $\Xi_{\text{cap}} \in \{35, 38, 40, 42, \mathbf{43}, 45, 47\}$ (benchmark bold; per-rung frequencies in Table 31). A confirmed compact-angle line at any frequency *outside* the window falsifies the calibrated placement and its ladder (F-Kaxion-1d); a line at the randomized-angle mass ($\gtrsim 6$ GHz) falsifies the pinning conditional C-Kaxion-d instead of the framework.

Strength. Exactly on the DFSZ line: $g_{K\gamma\gamma} = (\alpha/2\pi f_K) [8/3 - 1.92] = (1.2 \pm 0.07) \times 10^{-15} \text{ GeV}^{-1}$; expected cavity power $P_{\text{sig}} = g_{K\gamma\gamma}^2 (\rho_{\text{loc}}/m_K) B_0^2 V C \min(Q_L, Q_a)$ — worked budgets at three instrument classes in Table 17.

Line shape. Lab-frame standard-halo template (Part IX): sharp low edge *at* ν_K , peak $\simeq 0.6$ kHz above it, FWHM $\simeq 1.8$ kHz, 77% of the power within 1.95 kHz, matched-filter effective width 2.9 kHz; energy centroid 1.38 kHz above the edge, breathing ± 74 Hz annually and ± 1.3 Hz (Daejeon) / ± 1.2 Hz (Fermilab) with the *sidereal* day (Eqs. (48.1)–(48.2)).

The identification tests. (i) The count: the line stands alone across the window — any second in-window line established at the same grade as the first (a persistent excess with the halo line shape, surviving the same vetoes) kills the count; a single-mode receiver suppresses an in-window second line by ~ 74 dB, so the test must be asked for (reanalysis, retune, or a second mode). (ii) The address: on the DFSZ line at $E/N = 8/3$, inside $[1.59, 2.13]$ GHz.

Monday morning, in increasing cost.

- reanalyze any archived scan crossing the window for a second line inside it, and ask whether the candidate stands alone (zero hardware);
- camp the benchmark: at DFSZ depth the matched line reaches 5σ in seconds-to-minutes of live time (Table 17); the full window is hours (array-class) to ~ 100 days (single cavity) of idealized dwell;
- retune to any second candidate frequency in the window — the count test settles it within a day;
- complete the window at depth — the fence and the count close together.

The four registered kill conditions (the count, the address, and clauses (c) and (d) of F-Kaxion-1; App. E.2): $\square$ Count test: a second in-window line established at the same grade as the first $\Rightarrow$ the count refuted; the window swept at that grade with none $\Rightarrow$ the count confirmed. $\square$ Address: a confirmed line off the DFSZ locus at $> 5\sigma \Rightarrow$ dead. $\square$ F-Kaxion-1c: window swept at DFSZ depth, silent $\Rightarrow$ the calibrated placement, its ladder and C-Kaxion-d retired, not the framework; the Kaxion owed at $\gtrsim 6$ GHz or the relic not this sector's. $\square$ F-Kaxion-1d: confirmed line outside the window $\Rightarrow$ the same objects retired; at $\gtrsim 6$ GHz, C-Kaxion-d.

Reproduce everything. The discovery simulation (spectra, modulations, the count-test acceptance; fixed seeds) is a single standard-library Python script, available from the author on request (§52); replace its instrument block with your own and re-run.

37.4 The first kill test has already been faced

Falsifier F-Kaxion-1d is the sharpest clause in the registry: a *confirmed* compact-angle-like line at any frequency outside $[1.59, 2.13]$ GHz kills the calibrated placement and its ladder outright, wherever it appears and however faint (at the randomized-angle mass it would kill the pinning

conjecture instead). Its reach was demonstrated within weeks of this paper’s drafting. The candidate of Ahn et al. (2026) sat at 1.036 GHz — outside the window — and at a strength near $1.3\times$ the KSVZ coupling, roughly an order of magnitude brighter in power than the DFSZ line the Kaxion sits on: wrong frequency *and* wrong brightness. Had it been confirmed, the Kaxion would have been dead that week, with no adjustment available — the window is closed and the depth is locked. It was not confirmed: the excess failed persistence and cross-instrument checks and was excluded at 90% confidence, with the follow-up reaching DFSZ-adjacent depth. The theory faced a kill test it could not have dodged, and survived.

The general point outlasts the episode. Because F-Kaxion-1d bites at *every* frequency, every haloscope on Earth — whatever band it scans, whatever model it targets — is running a Kaxion kill test whenever it runs at confirmation-grade depth. The candidate-validation machinery the field has built for its own protection (persistence, field-off, cross-instrument, template consistency) is exactly the machinery that executes this paper’s falsifiers. The Kaxion does not merely invite the field’s scrutiny; it is structurally exposed to all of it, all the time. That is what it means for a prediction to be unique: one window, known today, and everywhere else on the dial a standing offer of refutation.

Gather what this Part has established. The Kaxion’s window is not a distant target: it lies inside the band that two programs — IBS-CAPP, running, and ADMX-EFR, designed and awaiting a funding decision — would be built to reach at exactly the DFSZ depth the prediction demands, and the arithmetic of a named-frequency search puts a decisive exposure of the benchmark bin at live-times of minutes rather than years. The discovery calendar turns that into dates: a sweep of [1.59, 2.13] GHz completed at DFSZ depth with no line retires the Kaxion by the early 2030s, and a confirmed line anywhere else retires it sooner — a clause the field has already exercised once, near 1.036 GHz, and the theory survived.

Yet every instrument in that landscape shares one limitation: each reads a *single* frequency at a time. None can interrogate the window’s count — the line standing alone, the identification test Part VI pre-registers — within one run. *Confirming* the Kaxion, as opposed to merely detecting an axion-like line, therefore calls for apparatus built to the prediction’s own shape: a receiver that covers the fenced window completely and checks the count as it goes. Part VIII designs that instrument, and shows that it asks for no physics the laboratory does not already have.

38 Part VII in plain English: how you would actually look, and how close the field already is

The detection principle is fifty years old and has not needed improving. Put a resonant cavity in a strong magnetic field; a passing pseudoscalar converts to a photon at the cavity frequency; tune the cavity and listen. The difficulty has never been the principle. It is that the signal is tiny and the interesting frequency is unknown, so the scan is long.

The Kaxion changes the second half of that sentence and not the first. The frequency is stated in advance, which turns an open-ended scan into a bounded one, and a bounded scan can be costed, scheduled, and finished.

This Part sets the prediction against what the field has actually achieved. Some experiments have demonstrated the required depth at frequencies below the window. Others are working in the window at shallower depth. Nobody has yet done both at once in the right place, which is the honest summary and also the opportunity: the gap is not a technology gap, it is a scheduling gap.

The projections quoted are the collaborations' own, not this paper's, and they are quoted at the depth those collaborations quote. That is deliberate. A paper that re-derived other people's sensitivities to make its own case look better would deserve the skepticism it got.

Part VIII

Future detection

A dedicated Kaxion instrument for the fenced window and the count test

39 The multi-mode count-test receiver

Principle. A single-mode haloscope answers “is there power at ν ?”; this paper’s count (derived-conditional, §9.4) adds a second question worth instrumenting: “is the line *alone* across the window?” Concept A is a receiver with two (or more) simultaneously instrumented modes — a multi-mode cavity, or cavities phase-locked to one synthesizer — placed at configurable spacings across the fenced window and scanned together. One mode carries the discovery scan; the others run the count test of Part VI in real time, at zero marginal live time. Coincidence logic still earns its keep, in the inverted sense the single-line doctrine demands: a genuine Kaxion excites exactly one mode, while wideband technical artifacts (gain ripples, interference combs) light several at once — so demanding *exclusivity* rather than coincidence suppresses the classes of background a single-mode scan finds hardest to veto.

Design quantities. Center 1.95 GHz, secondary modes configurable across the [1.59, 2.13] GHz window, condensate linewidth 2.0 kHz; single-photon-counting readout (transmon-qubit class) as the upgrade path beyond linear amplification — at 1.95 GHz the thermal occupation at 40 mK is $n_{\text{th}} \simeq 0.1$, so the gain over a quantum-limited amplifier is at most the factor $(1 + 2n_{\text{th}})^{-1}$ in noise, worthwhile but not decisive (Objection 43); every baseline in this paper is computed for the amplifier.

40 The spin couplings, and why no spin port is staged

At the chiral level the Kaxion carries ordinary nucleon- and electron-spin couplings, and the sheet must be printed with its verdict. The Kaxion has no heavy-quark completion; what it shares with the KSVZ class is the property that matters for the sheet — no tree-level quark coupling, $C_q^0 = 0$, the gluonic portal alone — so the chiral-level couplings are those tabulated for that class: $C_p = -0.47 \pm 0.03$ and $C_n = -0.02 \pm 0.03$ (Grilli di Cortona et al., 2016), with the electron coupling loop-suppressed at $C_e \sim 3 \times 10^{-4}$. The sheet depends on $C_q^0 = 0$ and on nothing else about a completion, and that fence is carried wherever it is quoted; the photon coupling, by contrast, sits on the DFSZ line by the anomaly ratio $E/N = 8/3$ alone and carries no such fence. At $f_K = 7.1 \times 10^{11}$ GeV every spin port — NMR-class nucleon readout (CASPER-style, CASPER Collaboration, 2021) or the ferromagnetic-magnon electron channel (QUAX-class, Crescini et al.,

2020) — is out of reach on any useful timescale, and the Larmor arithmetic settles the practical one before sensitivity does: at $\nu_K = 1.95$ GHz the proton port would need $B_p = 45.8$ T, the electron port $B_e = (70 \pm 7)$ mT. No spin port is therefore staged in this paper; the Kaxion’s measurable portals are the photon and gluon couplings of Part V, and the collider carries the baryonic sector’s exposure (Arisaka, 2026, P1-KK). Should a future completion move the sheet, a spin port would ride on the photon detection — a retrofit to an instrument that has already found the line — and would confirm a second channel of a known signal rather than search blind.

41 Search strategy and roadmap

Table 18 condenses the strategy of Parts VII–VIII into three stages — what runs now at zero marginal cost (the count-test re-binning), what decides discovery-or-closure around 2030 (the DFSZ window campaigns), and what the identification instruments of this Part add in the 2030s — naming for each stage its deliverable and the falsifier it retires. The prose below walks the same three stages.

Table 18: The search program. Stage costs are relative; depths are on the DFSZ=Kaxion line. Three stages, each with what it does, what it delivers, and what it costs relative to the others. Stage I needs no new hardware at all — it is a reanalysis of data that already exists, made possible by knowing what to look for — which is why it is listed first and why its cost is the lowest. Depths are quoted on the DFSZ line, which is the Kaxion line, so a stage that reaches DFSZ reaches the prediction.

Stage	Instrument	Target	Outcome
I (now–)	Count-test reanalysis of existing 1–2 GHz data (Part VI); multi-mode pathfinder retrofit	the window at KSVZ depth; the count test run on archival data	method validation; first count-aware limits
II (2030 era)	Concept-A multi-mode in the window, photon counting; ride the announced DFSZ program	full [1.59, 2.13] GHz at DFSZ depth	<i>discovery, or closure of the window</i>
III (2030s)	Deep rescan and modulation program on any candidate (Part IX’s fingerprints)	the single line, confirmed and characterized	identification: the address, the count, and the modulations together

Impact. A Stage-II detection would be, with respect to specificity, stronger than the Higgs discovery template: the Higgs was predicted without its mass, while the Kaxion arrives with its window (a discrete ladder around a calibrated benchmark, Section 29.2), its coupling (DFSZ-exact), and its count (single, no partner) all written down in advance. A Stage-III identification — the line confirmed alone across the window, at the DFSZ address — would simultaneously identify the dark matter and read the cap loading of a six-dimensional geometry off a frequency. Conversely, the full program executed null fires F-Kaxion-1c: it retires the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d — the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinned branch, or the relic is not this sector’s — exactly as a falsifiable placement should die, and without a kill of the framework, which this paper does not claim.

41.1 The kill-test timeline

The instruments of this Part, the scan fronts of Part VII, and the falsifiers of Part VI meet on one calendar. Figure 10 draws it in the cohort’s kill-test-timeline convention: the test already faced and survived (early 2026), the standing camping option, the scanning fronts approaching the window from below, the upper-edge campaign, and — in red — the kill date on which F-Kaxion-1c stands executed if no line has appeared. Every card names its venue and its falsifier; the wager is dated, and the dates are near.

This Part began with a question every unique prediction must eventually face — *what instrument would you build if you believed the number?* — and ends having answered it once, honestly: a multi-mode receiver that covers the fenced window at DFSZ depth and runs the count test as it goes, riding hardware the field already knows how to build, placed on a calendar that ends in the early 2030s. In plain words: nothing in the identification program waits on an invention; it waits on a decision. What such a decision would buy — day by day, panel by panel — is exactly what the next Part computes.

42 Part VIII in plain English: what an instrument built for this prediction would look like

Existing haloscopes were designed to search a range. If the prediction of this paper is right, the right instrument is designed to interrogate a point — and the two are not the same machine.

The difference is the fence and the count. A machine that scans looks for one line and moves on. A machine built for this prediction has to cover one narrow band completely, at full depth, and to check as it goes that any line it finds is alone in the band — which favors several modes live at once over one mode swept past everything. That requirement drives most of the design: the mode structure, the receiver chain, the exclusivity logic.

The Part sets out what would be required rather than what would be convenient. Each requirement is traced to the Part that forces it and to the falsifier it executes, and the reason for that discipline is to keep the proposal honest. A requirement with no falsifier behind it is a preference, and preferences are how instruments become expensive.

Nothing here is beyond current practice. The magnet exists, the amplifiers exist, the cavities exist. What does not yet exist is a machine that puts them together with the fenced window and the count in mind, and the argument of this Part is that building one is a modest step rather than a new program.




Figure 10: Falsification / kill-test timeline for the Kaxion. The discovery calendar of Table 16 drawn on a forward time axis (2026–2033, broad data-window eras rather than precise launch dates), in the cohort’s kill-test-timeline convention. Violet: the F-Kaxion-1d test already faced and survived (the 1.036 GHz candidate, early 2026, §37.4). Orange: the standing pre-registered camping option (§37). Blue: the scanning fronts approaching the window from below (CAPP-MAX, then the ADMX four-cavity array). Green: ADMX-EFR closing the upper edge. Red: the kill date — a completed DFSZ-depth sweep of [1.59, 2.13] GHz with no line retires the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d (the Kaxion then owed at $\gtrsim 6$ GHz, or the relic not this sector’s; not a kill of the framework), by ~2030–2032 on stated roadmaps if an array of the EFR class sweeps the lower half, 2032–2034 for a single CAPP-MAX-class cavity (§36.4).

Part IX

Simulating the Discovery

The data of discovery week, computed in advance: inputs, equations, error budgets, and the verdicts

43 Why a discovery can be simulated at all — and what is simulated here

A theory with free parameters cannot tell an experimentalist what its discovery will look like; it can only promise to explain the discovery afterwards. The Kaxion has no adjustable parameter left once its one scale is calibrated (Part V), and that changes what this Part can do: every property of the signal — its frequency, its strength, its spectral shape, its seasonal breathing, its count — is already fixed by Parts V–VII (and graded, theorem by conjecture, in the master ledger of Table 13). The data of discovery week can therefore be *computed before the week arrives*, and this Part computes it, to the standard the claim deserves: every input parameter is tabulated with its provenance and its assigned uncertainty (§44), every equation is stated and worked (§45–46), every figure carries both its statistical and its systematic errors, and the full error budget is collected in one table (§50). The generator is a single deterministic script, available from the author on request (§52), so that a collaboration can reproduce every panel exactly — and, more usefully, re-run it with its own instrument parameters.

The scope should be stated as precisely as the inputs. *Simulated*: the received power spectrum of the condensate line in a camped cavity (amplitude, shape, statistics); its annual and daily modulation; and the single-mode acceptance that quantifies why the count test of Part VI must be a deliberate measurement. *Not simulated, and why*: radio-frequency interference and its vetoes (site-specific; handled operationally by the environmental-antenna protocol the field already runs, §37.3); cryogenic interruptions and duty cycle (bookkeeping, folded in as the 10–100× overhead factor quoted throughout Part VII); receiver gain drift (removed in analysis by baseline filtering, entering the budget only as the residual in Table 26); and correlated noise between mirrored JPA bins (a documented few-percent effect in current receivers, second-order for a camped — as opposed to scanned — measurement). Each omission is an engineering perturbation on a signal whose defining properties — location, shape, count — it cannot mimic.

44 The simulation pipeline and its inputs

Figure 11 is the pipeline. Three input blocks feed two physics modules (signal and noise), a spectrum generator, and an analysis chain; the systematic budget propagates onto every output.




Figure 11: The simulation pipeline. Three input blocks (theory sheet, frozen by Parts II–V; halo model; instrument model — Tables 19–21) feed the signal and noise modules, whose equations are stated in §45–46; the spectrum generator draws fixed-seed Gaussian noise on a 100 Hz grid; the analysis chain applies baseline removal, template matching, and the candidate protocol of §37.3 with the signal-injection efficiency ϵ_a . The systematic budget (Table 26, red dashed) propagates onto every output as the envelopes drawn in Figures 12–14. The generator is a single self-contained script (§52).

Table 19: Simulation inputs I — the theory sheet (frozen; Parts II–V). Central values with the Standard-Model-side errors of Part V; the geometric integers carry no uncertainty. These inputs are common to every figure of this Part. These inputs are frozen before any figure in this Part is drawn, and none is adjusted afterwards.

Quantity	Value	Error	Provenance
benchmark frequency ν_K	1.95 GHz	$\pm 1.3\%$ at fixed f_K	Part V, Eq. (26.1)
mass m_K	$8.1 \mu\text{eV}$	$\pm 1.3\%$ at fixed f_K	Part V
photon coupling $g_{K\gamma\gamma}$	$1.2 \times 10^{-15} \text{ GeV}^{-1}$	$\pm 5.5\%$	Part V (DFSZ, $E/N = 8/3$)
line count	single	count (conditional)	Part VI
window	[1.59, 2.13] GHz	discrete Ξ_{cap} ladder (ansatz)	Part V

The three blocks are deliberately separated because their epistemic status differs: the *theory sheet* (Table 19) is frozen by Parts II–V and carries only the Standard-Model-side errors of Part V; the *halo model* (Table 20) is astrophysics, shared with every haloscope in the field and adopted here in exactly the form the collaborations use (Ahn et al., 2026; ADMX Collaboration, 2025b; Kim et al., 2024); the *instrument model* (Table 21) is engineering, benchmarked on the published CAPP-MAX system because it is the closest existing realization of the camp of §37. A reader who disputes an entry can trace exactly which figures it moves, and by how much, through Table 26.

Table 20: Simulation inputs II — the halo model (SHM). The isothermal-sphere parameters of this simulation (conventions stated per row), with the systematic uncertainties assigned in this simulation and their justification. These are *model-level* inputs: they calibrate the astrophysics, not the geometry, and never touch the location or the count.

Quantity	Value	Assigned syst.	Justification
local density ρ_{loc}	0.45 GeV/cm ³	$\pm 22\%$	rotation-curve and stellar-kinematics determinations span 0.35–0.55; value and spread as adopted by the collaborations
dispersion parameter v_0	220 km/s	$\pm 8\%$	Maxwellian $f(v) \propto v^2 e^{-v^2/v_0^2}$, most-probable-speed convention ($v_0 = v_c$; rms 270 km/s); sets the width via Eq. (45.2)
solar speed $v_\odot$	232.8 km/s	± 3 km/s	galactocentric solar motion
orbital speed v_{orb}	29.79 km/s	—	ephemeris (exact)
orbit projection $\cos \gamma$	0.49	$\pm 5\%$	angle between orbital plane and solar apex
rotation speed $v_{\text{rot}} \cos \lambda$	0.374 km/s	site	Daejeon latitude 36.4° (example site)

45 The signal model: equations and worked numbers

Signal power. The camped cavity extracts the registered Sikivie power of Part VII,

$$P_{\text{sig}} = g_{K\gamma\gamma}^2 \frac{\rho_{\text{loc}}}{m_K} B_0^2 V C \min(Q_L, Q_a), \quad (45.1)$$

identical to Eq. (35.1); nothing is re-derived here. What this Part adds is the worked budget: Table 22 evaluates (45.1) factor by factor with each factor’s uncertainty carried alongside, so the reader can see where the number — and its error — comes from.

Line shape. In the standard halo model the condensate velocities are Maxwellian in the galactic frame, $f(v) dv \propto v^2 e^{-v^2/v_0^2} dv$ with the most-probable speed $v_0 = 220$ km/s (rms speed $\sqrt{3/2} v_0 = 270$ km/s; the two conventions are a factor 1.5 apart in v_0^2 and must not be confused). Each quantum arrives at $\nu = \nu_K(1 + v^2/2c^2)$, so substituting $v = \sqrt{2c^2\delta/\nu_K}$ for the detuning $\delta = \nu - \nu_K$ gives, in the galactic frame, the shape

$$g(\delta) d\delta = \frac{2}{\sqrt{\pi}} \frac{\sqrt{\delta}}{\Delta^{3/2}} e^{-\delta/\Delta} d\delta, \quad \Delta \equiv \frac{\nu_K v_0^2}{2c^2} = 525 \text{ Hz}, \quad (45.2)$$

normalized to unit area for $\delta > 0$: a sharp low-frequency edge at ν_K (the condensate at rest), a maximum at $\delta = \Delta/2$, and a thermal tail upward, with full width at half maximum $1.79 \Delta = 942$ Hz and 94% of the power within the first 1.95 kHz. The laboratory rides the halo wind at $v_\odot \simeq 232$ km/s, and the boost does not shift this shape rigidly: the lab-frame distribution is $f_{\text{lab}}(v) \propto v [e^{-(v-v_\odot)^2/v_0^2} - e^{-(v+v_\odot)^2/v_0^2}]$, whose edge stays at ν_K (the slowest quanta are still the ones at rest in the laboratory) while the maximum moves to $\simeq 0.6$ kHz above ν_K and the width

Table 21: Simulation inputs III — the instrument model (CAPP-MAX-class camp at 1.95 GHz). Engineering parameters, their calibration methods, and the assigned systematics; the volume row carries the honest $V \propto \nu^{-3}$ retuning penalty from the published 37 L at ~ 1.1 GHz. The three sibling classes of Table 17 differ only in the first five rows.

Quantity	Value	Assigned syst.	Method / justification
magnetic field B_0	12 T	$\pm 1\%$ ($\pm 2\%$ on B_0^2)	NMR probe map of the solenoid bore
cavity volume V	6.4 L	$\pm 2\%$	machining + cryogenic contraction; $37 \text{ L} \times (1.09/1.95)^3 = 6.5 \text{ L}$, carried as 6.4 L inside the stated $\pm 2\%$
form factor C_{010}	0.60	$\pm 5\%$	bead-pull measurement cross-checked against FEM eigenmode computation
loaded quality factor Q_L	10^5	$\pm 5\%$	vector-network-analyzer S -parameters at each tuning
system noise T_{sys}	200 mK	$\pm 6.4\%$	Y-factor + shot-noise calibration of the JPA chain; dominant instrument systematic
analysis efficiency ϵ_a	0.93	$\pm 1\%$	software signal injection through the full analysis chain
spectral bin $\Delta\nu_{\text{bin}}$	100 Hz	—	analysis choice; $\sim 1/18$ of the lab-frame FWHM
condensate bandwidth $\Delta\nu_a$	1.95 kHz	from ν_0	registered Part-VI value (77% of the lab-frame line power; FWHM $\simeq 1.8$ kHz lab frame, 942 Hz galactic frame)

Table 22: The worked power budget. Equation (45.1) evaluated term by term for the instrument of Table 21. Theory-side errors (first row) are quoted separately from halo and instrument systematics because they move the *prediction*, not the measurement. The product reproduces the registered Table 17 entry.

Factor	Value	Uncertainty	Class
$g_{K\gamma\gamma}^2$	$(1.2 \times 10^{-15} \text{ GeV}^{-1})^2$	$\pm 11\%$ ($2 \times 5.5\%$)	theory (SM-side)
ρ_{loc}/m_K	$0.45 \text{ GeV}/\text{cm}^3 / 8.1 \mu\text{eV}$	$\pm 22\%$	halo
B_0^2	144 T^2	$\pm 2\%$	instrument
V	6.4 L	$\pm 2\%$	instrument
C	0.60	$\pm 5\%$	instrument
$\min(Q_L, Q_a)$	10^5	$\pm 5\%$	instrument
P_{sig}	$4.1 \times 10^{-23} \text{ W}$	$\pm 23\%$ (halo+instr.)	$\Rightarrow 3.2 \times 10^1 \text{ photons/s}$

grows to a full width at half maximum of $\simeq 1.8$ kHz, with 77% of the power within the first 1.95 kHz and 93% within 3 kHz. The simulation uses the lab-frame form; Eq. (45.2) is its galactic-frame parent and the scale Δ is the halo thermometer the fit reads.

Table 23: Time to 5σ , camped on the benchmark, with the systematic range. Central values from the matched-filter expression of §46 with $\epsilon_a = 0.93$ applied ($t \propto \epsilon_a^{-2}$; the idealized Table 17 quotes $\epsilon_a = 1$); the range propagates the total $\pm 24\%$ SNR systematic of Table 26, dominated by the halo density. Idealized live time; the duty and calibration conventions of Table 17 apply (scanning-campaign overheads 10–100 $\times$; camped runs far less).

Instrument class	$t_{5\sigma}$ (central)	syst. range
Representative design cavity (SQL 94 mK)	2.5 s	1.6 s–4.3 s
CAPP-MAX-class single cavity (200 mK)	6.2 min	4.0 min–11 min
EFR-class 18-cavity array (150 mK)	1.3 s	0.84 s–2.2 s
CAPP-MAX-class cavity, 4 K HEMT	41 h	27 h–71 h

Per-bin expectation. On the analysis grid of bin width $\Delta\nu_{\text{bin}}$, the expected signal-to-noise in bin k after live time t is

$$\mu_k = \epsilon_a \frac{P_{\text{sig}} w_k}{k_B T_{\text{sys}}} \sqrt{\frac{t}{\Delta\nu_{\text{bin}}}}, \quad w_k = \int_{\text{bin } k} g(\delta) d\delta, \quad (45.3)$$

with w_k the fraction of line power in the bin (peak bin of the lab-frame template: $w = 0.050$; the galactic-frame form would give 0.092). This is the Dicke radiometer statement of Part VII applied bin-wise, with the analysis efficiency ϵ_a applied explicitly.

46 The noise model and the time to detection

Noise. Over an integration $t \Delta\nu_{\text{bin}} \gg 1$ (here $\geq 10^3$ even in the shortest panel), the averaged power in each bin is χ^2 -distributed with $2t \Delta\nu_{\text{bin}}$ degrees of freedom, hence Gaussian to excellent approximation, with standard deviation

$$\sigma_P = \frac{k_B T_{\text{sys}} \Delta\nu_{\text{bin}}}{\sqrt{t \Delta\nu_{\text{bin}}}} \implies y_k = \mu_k + n_k, \quad n_k \sim \mathcal{N}(0, 1) \text{ i.i.d.} \quad (46.1)$$

in signal-to-noise units. The i.i.d. assumption sets aside two documented effects — mirrored-bin correlations in JPA receivers (a $\sim 4\%$ SNR effect when optimally combined) and baseline-removal residuals (sub-percent after standard filtering) — both absorbed as the analysis-efficiency and gain-drift rows of Table 26.

Time to detection. Matched-filtering the full line with its own template gives the time to a 5σ statistical detection, $t_{5\sigma} = \Delta\nu_{\text{eff}} [5 k_B T_{\text{sys}} / (\epsilon_a P_{\text{sig}})]^2$, with $\Delta\nu_{\text{eff}} \equiv (\int g^2 d\delta)^{-1} = 2.9$ kHz for the lab-frame template (a boxcar over the first 1.95 kHz captures 77% of the power and is not the optimal filter), evaluated per instrument class in Table 23 together with its systematic range: since $t \propto \text{SNR}^{-2}$, the $\pm 24\%$ envelope on the expected SNR (Table 26) maps to a factor $\times 0.65$ to $\times 1.7$ on the time, which changes nothing qualitative: seconds stay seconds, minutes stay minutes, hours stay hours.

47 Results I: the camped spectra

Figure 12 shows the camp as data: the same benchmark bin at one hundred seconds, one hour, and one day of live time. Three curves are drawn in every panel, and the distinction between them is the point of this whole Part. The black trace is the synthetic data y_k of Eq. (46.1): its scatter *is* the statistical error — unit variance per bin by construction, shrinking relative to the signal as $\sqrt{t}$. The red curve is the expectation μ_k of Eq. (45.3): everything about it is computed, nothing fitted — amplitude from Table 22, shape from Eq. (45.2). The dashed orange envelope is the systematic band: $\pm 24\%$ on μ_k , the quadrature total of Table 26, dominated by the halo density — visibly a second-order caveat on an unmistakable signal by the third panel, where the peak-bin expectation reaches 20σ and the line *shape* (the halo thermometer, Eq. (45.2)) is resolved bin by bin.

Two systematic classes deliberately do *not* appear as bands in these panels. The theory-side window on ν_K moves the whole camp band, not the line within a panel — it is the window-scan cost already priced into §37.2. And the halo-density systematic, while it scales the amplitude, cannot delay discovery by more than the factor 1.7 of Table 23: at camped exposures the verdict is over-determined.

48 Results II: the modulation fingerprints

A real halo line is not stationary; the laboratory rides the Earth. The line’s edge stays at ν_K ; what moves is the line’s *energy centroid* above the edge — the observable a template fit returns — which oscillates annually as the orbital velocity adds to and subtracts from the Sun’s galactic motion,

$$\delta\nu_{\text{cent}}(t) = \frac{\nu_K}{c^2} \left[\frac{v_\odot^2}{2} + \frac{3v_0^2}{4} + v_\odot v_{\text{orb}} \cos \gamma \cos\left(\frac{2\pi(t - t_{\text{Jun}})}{1 \text{ yr}}\right) \right] = 1375 \text{ Hz} + 74 \text{ Hz} \times \cos(\cdots), \quad (48.1)$$

and daily as the rotation velocity sweeps the laboratory through the halo wind,

$$\delta\nu_{\text{day}}(t) = \frac{\nu_K v_\odot v_{\text{rot}} \cos \lambda_{\text{lab}} \cos \delta_{\text{apex}}}{c^2} \cos\left(\frac{2\pi t}{23.934 \text{ h}}\right) \approx 1.3 \text{ Hz} \text{ at Daejeon's latitude}, \quad (48.2)$$

where $\cos \delta_{\text{apex}} \simeq 0.67$ projects the rotation velocity onto the direction of solar motion (the apex sits at declination $\simeq 48^\circ$); Fermilab’s latitude gives 1.2 Hz. (The peak of the template sits 613 Hz above the edge and moves less, ± 63 Hz annually, because the boost reshapes the line rather than shifting it; the centroid carries the full velocity shift and is the fitted quantity throughout.) Both are small against the linewidth but within reach of centroid fitting, because the fit precision is set by the line’s total signal-to-noise, not by its width alone:

$$\sigma_{\text{cent}} \simeq \frac{\text{FWHM}}{\text{SNR}_{\text{line}}(t)}, \quad \text{SNR}_{\text{line}}(t) = \epsilon_a \frac{P_{\text{sig}}}{k_B T_{\text{sys}}} \sqrt{\frac{t_{\text{live}}}{\Delta\nu_{\text{eff}}}} \implies \sigma_{\text{cent}} \approx 35 \text{ Hz per fortnight} \quad (48.3)$$

at a deliberately conservative 1/30 duty cycle ($t_{\text{live}} \simeq 11$ h per fortnight, FWHM 1.8 kHz, $\Delta\nu_{\text{eff}} = 2.9$ kHz; at full duty the same fortnight gives ≈ 6 Hz). This is the estimator used in Part X as




Figure 12: Results I — the camped spectra, with statistical and systematic errors shown.

Synthetic power spectra ($\Delta\nu_{\text{bin}} = 100$ Hz, units of the per-bin thermal noise σ ; identical fixed-seed noise across panels, §52) for the CAPP-MAX-class camp of Table 21 at $t = 100$ s, 1 h, 24 h. Black: data $y_k = \mu_k + n_k$, whose unit-variance noise *is* the statistical error (Eq. (46.1)). Red: the expected signal μ_k of Eq. (45.3) — amplitude from the worked power budget of Table 22, shape from Eq. (45.2). Dashed orange: the $\pm 1\sigma_{\text{sys}}$ envelope, $\pm 24\%$ on μ_k from the total systematic budget of Table 26 (halo density dominated). The peak-bin expectations are $\mu \approx 0.7\sigma$, 4.1σ , and 20σ ; the frequency *position* carries the separate theory-side window of Part V, which moves the camp band, not the line within these panels.

well. These are the statistical error bars of Figure 13; the daily panel folds one year of full-duty camping into three-sidereal-hour bins (± 3.5 Hz per bin). The systematic envelope on the curve itself (dashed) collects the halo rows of Table 20: $\pm 10\%$ on the annual amplitude, $\pm 5\%$ on the mean offset. Table 24 summarizes. As throughout, these are model-level characteristics — they calibrate the halo, not the geometry — and their diagnostic power is negative space: a clock harmonic, an interference tone, a thermal artifact sits still. A line that breathes with the year and turns with the sidereal day lives in the halo. The 1.036 GHz validation campaign (§37.4) applied exactly this discipline.

49 Results III: single-mode acceptance, and why the count test is a deliberate act

The count test of Part VI asks whether a confirmed line stands alone across the window, and the first fact the simulation establishes is that a camped single-mode receiver cannot answer that

Table 24: The modulation observables. Amplitudes from Eqs. (48.1)–(48.2); the statistical column assumes fortnightly centroid fits (Eq. (48.3), 1/30 duty) accumulated for one year — 26 points at ± 35 Hz, sinusoid-fit precision $\sigma_A = \sigma\sqrt{2/N}$; the daily row uses the three-hour-bin analysis of a full-duty year, since fortnightly centroids average over the sidereal day and carry no daily information; systematics from Table 20.

Observable	Period	Amplitude	Stat. (1 yr)	Syst.
centroid above the edge	—	1375 Hz	± 6.9 Hz	$\pm 9\%$ ($v_\odot^2, v_0^2$)
annual modulation	365.25 d	74 Hz	± 9.7 Hz ($\sim 7.6\sigma$)	$\pm 10\%$ ($v_\odot v_{\text{orb}} \cos \gamma$)
daily modulation	23.934 h	1.3 Hz	± 1.8 Hz ($\sim 0.7\sigma/\text{yr}$)	site ($\cos \lambda_{\text{lab}}$), $\cos \delta_{\text{apex}}$




Figure 13: Results II — the modulation fingerprints, with errors. Left: the line’s energy centroid rides $\simeq 1375$ Hz above its edge at ν_K and breathes annually with amplitude $A_{\text{ann}} \simeq 74$ Hz (Eq. (48.1); solid blue). Dashed: the halo-side systematic envelope — $\pm 10\%$ on A_{ann} from $v_\odot = 232.8 \pm 3.0$ km/s and the orbit-projection geometry, $\pm 9\%$ on the mean offset ($\propto v_\odot^2$ and v_0^2) (Table 20). Points: simulated fortnightly centroid fits with their *statistical* error of ± 35 Hz (Eq. (48.3), at a conservative 1/30 duty cycle). Right: the daily (sidereal) signature at the latitude of Daejeon, amplitude $\simeq 1.3$ Hz (Eq. (48.2), including the apex projection $\cos \delta_{\text{apex}} \simeq 0.67$); site-dependent through $\cos \lambda_{\text{lab}}$. Points: one *year* of full-duty camping, folded into three-sidereal-hour bins, statistical error ± 3.5 Hz per bin — even a full-duty year leaves the daily signature below 1σ , and the 1/30-duty fortnightly program does not see it at all (Table 24). At single-cavity aperture the sidereal fingerprint is not detected — it is the observable the KAXIO aperture of Part X exists to buy. Both are halo-model characteristics, quoted at model level; their diagnostic power is that no laboratory artifact reproduces either period.

passively. A receiver reads any second line through its mode’s Lorentzian acceptance,

$$\left| \frac{\chi(\nu)}{\chi(\nu_c)} \right|^2 = \left[1 + \left(\frac{2Q_L(\nu - \nu_c)}{\nu_c} \right)^2 \right]^{-1}, \quad (49.1)$$

with ν_c the camped frequency, which at a representative in-window detuning — drawn here at $\Delta\nu/\nu = 2.5\%$, about 49 MHz at the benchmark, a tenth of the window’s width and typical of separations inside $[1.59, 2.13]$ GHz — evaluates to -74 dB for $Q_L = 10^5$: blindness, for every practical purpose total, drawn with its calibration systematic in Figure 14 and tabulated against Q_L

Table 25: Single-mode suppression of a second line at $\Delta\nu/\nu = 2.5\%$ detuning versus Q_L (Eq. (49.1); a representative in-window separation). At every realistic Q_L a second line elsewhere in the window is invisible without a deliberate second measurement; a multi-mode receiver regains it in full. The row to read is the one at the realistic loaded Q , not the best case.

Q_L	second-line suppression (dB)
50000	−68.0
100000	−74.0
200000	−80.0
10^6	−94.0




Figure 14: Results III — the count test, and why it must be asked for. A camped line at ν_c and a hypothetical second line ν' at a representative in-window separation (dashed red; the second line is a *hypothesis* the count test interrogates, drawn at -3 dB for illustration — the count predicts silence there). A single-mode receiver tuned to ν_c applies the Lorentzian acceptance of Eq. (49.1) (dark blue) and receives the second line suppressed by 74 dB — at -3 dB intrinsic, -77 dB, invisible in any realistic exposure. Light blue: the $Q_L \pm 5\%$ calibration systematic (Table 21), ∓ 0.4 dB on the suppression — immaterial to the conclusion (Table 25). The count test is therefore a deliberate act: retune, or instrument a second mode elsewhere in the window (Part VIII), and any second line would appear at camped SNR within a day. Statistical errors at the retuned mode are those of Figure 12.

in Table 25. The engineering meaning of the $\pm 5\%$ on Q_L is a ∓ 0.4 dB shift on a 74 dB suppression: the conclusion — *no single-mode instrument runs the count test by accident* — is systematics-proof by four orders of magnitude. The test is therefore a deliberate act (a re-binning of scan data that crossed the window, a retune, or the Part-VIII second mode), and at the camped SNRs of Table 23 a retune verdict on any candidate second frequency arrives within a day, carrying the same statistical and systematic structure as Figure 12.

Table 26: The systematic budget, in one table. Fractional effects on the expected per-bin SNR unless noted. Quadrature totals: amplitude-instrument $\pm 7.7\%$ (with the $T_{\text{sys}} \pm 6.4\%$ noise-floor row, combined instrument $\pm 10\%$); with the halo density, $\pm 24\%$ on SNR, i.e. a factor ≈ 1.5 – 1.7 either way on $t_{5\sigma}$. The two identification observables (last rows) are by construction immune to the amplitude-class systematics.

Source	Class	Size	Moves	Method
ρ_{loc}	halo	$\pm 22\%$	amplitude, $t_{5\sigma}$	surveys
v_0 (line width)	halo	$\pm 8\%$	shape, $\Delta\nu_a$	surveys
$v_\odot, \cos \gamma$	halo	$\pm 10\%$ (A_{ann})	modulation	astrometry
B_0^2	instrument	$\pm 2\%$	amplitude	NMR field map
V (retuned)	instrument	$\pm 2\%$	amplitude	metrology + contraction
C_{010}	instrument	$\pm 5\%$	amplitude	bead pull + FEM
Q_L	instrument	$\pm 5\%$	ampl.; ∓ 0.4 dB suppr.	VNA S -parameters
T_{sys}	instrument	$\pm 6.4\%$	noise floor	Y-factor + shot noise
ϵ_a , gain drift	analysis	$\pm 1\%$	amplitude	signal injection
line count (single)	—	immune	—	a yes/no, not a drift

50 The systematic-error budget

Table 26 collects every systematic of this Part in one place, with its class, its calibration method, and which observable it moves. Three structural points. *First*, the budget is halo-dominated: the $\pm 22\%$ on ρ_{loc} dwarfs the $\pm 7.7\%$ amplitude-instrument total ($\pm 10\%$ including the $T_{\text{sys}} \pm 6.4\%$ noise-floor row), and it is common mode across every haloscope on Earth — it rescales times, never verdicts. *Second*, the identification observables are engineered to cancel it: the count is a frequency comparison (no amplitude enters), and the window is a fence, not a rate; that is why the tests that carry the geometry sit in the shortest rows. *Third*, the theory-side errors of Table 19 are not “systematics” of the measurement at all — they are the width of the prediction, priced separately in Part V, and they shrink with lattice QCD, not with instrument time.

51 Discovery week — and what failure looks like

With the machinery of §45–50 on the table, the pre-declared sequence of Part VI reads as a schedule. *Day zero*: the excess of Figure 12 (middle panel — an hour of live time); the vetting protocol of §37.3 runs in hours, the camp having no scan schedule to protect. *Day one*: the count — a deliberate retune, the second mode, or the re-binning of whatever scan data crossed the window (Figure 14); the count test is read, either way, at camped SNR within the day. *Weeks*: the line shape sharpens bin by bin against Eq. (45.2), and the centroid fits begin tracing the annual curve of Figure 13. *Weeks to months*: the count test extends across the full window as the deep rescan proceeds — the single line standing alone, interrogated at every candidate separation. At the end of the week the logbook does not say “an axion-like particle has been detected”; it says: a line at the pre-registered frequency, with the halo’s shape, the sky’s modulation, standing alone across the fenced window — every number of it computed, with its errors, in the figures of this Part before any of it happened.

Honesty requires running the other tape. If the camp integrates down past DFSZ depth in silence, the same pipeline writes an exclusion instead: with no signal, the 90%-confidence upper limit on the coupling after live time t is

$$g_{90}(\nu) = g_{K\gamma\gamma}^{\text{DFSZ}}(\nu) \left[\frac{1.28}{\epsilon_a} \frac{k_B T_{\text{sys}}}{P_{\text{sig}}^{\text{DFSZ}}} \sqrt{\frac{\Delta\nu_{\text{eff}}}{t}} \right]^{1/2} \times (1 + \sigma_{\text{sys}})^{1/2}, \quad (51.1)$$

the standard construction with the systematic inflation the field applies (compare the 6.4% noise-calibration inflation of the 1.036 GHz follow-up, [Ahn et al., 2026](#)). One camped week at a single Ξ_{cap} rung excludes that rung at DFSZ depth; the ladder of Part V is finite; and the scan fronts of §36.9 sweep the seams between the rungs. The asymmetry deserves emphasis: *discovery* can happen in any week the camp runs, but *failure* is only declared by the full window — F-Kaxion-1c fires when [1.59, 2.13] GHz has been swept at DFSZ depth and no line has appeared, on the dates of Table 16. There is no third tape. The window cannot move, the depth cannot be renegotiated, and a confirmed line anywhere else on the dial retires the same objects (F-Kaxion-1d), or, at $\gtrsim 6$ GHz, the pinning conjecture. Either the logbook above gets written by the early 2030s, or the ledger entry of Appendix E.2 retires the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d — the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinned branch, or the relic is not this sector's. Neither is a kill of the framework, and this paper does not claim one.

52 Reproducibility and code release

Every figure and table of this Part is generated by a single, self-contained Python script (standard library only; no external packages), which evaluates Eqs. (45.1)–(51.1) from the input tables verbatim and draws the noise realizations from fixed seeds — seed 35 for the spectra, modulation, and count-test figures — so that every panel regenerates bit-identically. The script is available from the author on request, with the input tables as a separate, editable parameter block: a collaboration wishing to re-run the simulation for its own instrument replaces Table 21's eight numbers and re-executes. Nothing in this Part depends on proprietary tooling, unpublished data, or hidden constants; the reader who doubts a curve is invited — operationally, not rhetorically — to regenerate it.

53 The bottom line: why this discovery cannot be mistaken

Step back from the machinery and count what this Part has actually put on the table. Six independent, quantitative signatures of one particle, each computed in advance, each with its statistical reach and its systematic budget assigned: a *frequency* — one 2 kHz line in a pre-declared 544 MHz window, with its benchmark rung marked; an *amplitude* — the worked budget of Table 22, good to $\pm 24\%$; a *shape* — the sharp edge and thermal tail of Eq. (45.2), resolved within one camped day; *two clocks* — the year and the sidereal day, written into the line at 74 Hz and 1.3 Hz; and a *count* — the line standing alone across the window, interrogated deliberately and answered with silence. No other proposed constituent of the dark matter — no WIMP, no sterile neutrino, no generic axion — has ever arrived with a dossier like this, because every one of them keeps at least

one number free after the data. The Kaxion keeps one calibration, declared before the data, and no dial after it.

That completeness is what makes the identification proof against error. The six signatures are *orthogonal*: one is a location, one an intensity, one a spectral form, two are time dependences pinned to astronomy no laboratory can emulate, and one is a count the geometry fixes conditionally on its parents' registered modes. A false positive must therefore be a conspiracy: some artifact would have to sit in the right bin, at the right strength, with the right asymmetric shape, breathe with the Earth's orbit *and* turn with its day, and leave the rest of the window silent — each to within the printed error bars. No known artifact class counterfeits even two of these at once. (What this completeness has meant historically for the field's candidate signals — and what it buys programmatically — is developed where it belongs, in the case for the observatory that reads all six: §67.)

And the sword has two edges, which is what makes the first edge credible. Every panel that makes the discovery unmistakable makes the refutation unarguable: a line that fails the shape, sits still through the seasons, or refuses to stand alone does not become an interesting anomaly — it becomes the death certificate of this paper (F-Kaxion-1c and 1d priced in dates and depths; the count test priced in a re-binning). A theory that can be confirmed by its locked numbers can be killed by any one of them; the Kaxion accepts both verdicts from the same figures.

The deepest point is worth stating last, plainly. The reason this discovery can be simulated at all — the reason this Part could be written — is that the Kaxion is not a hypothesis fitted to the dark matter but a consequence forced by a geometry with one calibrated scale, and after that calibration there is no dial to turn. The five integers (1, 3, 5; 7, 35) and the one lattice input fixed the coupling's DFSZ locus, the shape of the window, and the count before any experimentalist was asked for anything; a line on a rung would read the cap loading Ξ_{cap} under the ladder ansatz, and a line off every rung would retire the ansatz. That is what it looks like when a theory stakes everything on numbers it can no longer change. The dark-matter problem has waited ninety years not for a better detector but for a prediction worth pointing one at; the preceding pages are that prediction, drawn to scale, errors included, waiting only for the data to be laid on top.

One thing the simulation has deliberately left undone. It drew the discovery week for a *single* camped cavity — one mode at a time; but the signatures were meant to be read *together*, so that no artifact can forge them one by one. That combined instrument is what the final Part builds. Part X takes the pipeline of this simulation, replaces its single cavity with a multi-mode array covering the fenced window, and lets it run for a year: the whole window at DFSZ depth, the count test live throughout, every surviving falsifier of this paper executed in one continuous experiment.

54 Part IX in plain English: what discovery week would look like, written down before it happens

This Part computes the data of a discovery in advance — the spectra, the modulation, the count — and it is worth being clear about why that is more than a display.

A prediction stated as a number can be matched after the fact by a sufficiently flexible analysis. A prediction stated as a full simulated dataset cannot, because it commits to the shapes as well as the positions: how the line moves with the day and with the year, how sharply it rises from its rest-mass edge, and that it stands alone across the band. All of those are consequences of the same fixed inputs, and all of them are printed here before any data exists.

The inputs are frozen at the head of the Part and are not adjusted for any figure in it. That is the property that makes the exercise worth doing. If a future dataset matches these figures, nobody has to take on trust that the match was not arranged; the arrangement is on the page, dated.

If a future dataset does not match them, the framework is wrong in a way that can be pointed at. That is the harder half of pre-registration and the reason for accepting it.

Part X

The KAXIO Proposal

A second-generation Kaxion observatory at Fermilab: the fenced window, one line, one year

55 Executive summary of the proposal

Parts VII–IX established, quantitatively, three facts: the Kaxion search is a camping problem, not a scanning problem (§37); the identification burden is carried by the fence, the address, and the count — the calibrated window, the DFSZ location, and the single-line count (derived-conditional, Part VI) — a combination no existing instrument is organized to read completely; and every observable of discovery week, with its errors, is computable in advance (Part IX). This Part converts those three facts into an instrument. We propose **KAXIO** — the **Kaxion Observatory** (*KAX-ee-oh*) — a major upgrade of ADMX-EFR at Fermilab that reorganizes its eighteen cavities into *two* nine-cavity sub-arrays, independently tunable across the fenced window — stacked on one rung for depth, split across rungs for coverage and the live count test — and runs steadily for one year. The instrument executes, in a single continuous run, every surviving test this paper has pre-registered: the DFSZ-depth line at the benchmark rung within the first week if the address is right, the count test — one line, no partner anywhere the fence allows — running whenever the sub-arrays differ and completing across the window within a quarter, the daily modulation emerging at $\sim 2\sigma$ within a month (a $\sim 6.5\sigma$ measurement by year end), the annual modulation across the year, and, on the null tape, a DFSZ-depth exclusion of the entire window (F-Kaxion-1c). The design reuses the two most expensive elements of EFR — the 9.4 T MRI magnet and the Dark Wave Laboratory cryogenic plant — and adds one genuinely new item: the independent tuning layer with its common-clock analysis (Table 34). Either outcome of the year is a landmark: the identified discovery of the dark matter, or the closure of the sharpest dark-matter prediction on record.

56 Design requirements

Most instruments derive their requirements from what technology makes convenient; KAXIO derives its six from what the theory sheet makes *mandatory*. Each row of Table 27 is traced to the Part of this paper that forces it and to the falsifier it executes; none is discretionary, and none can be relaxed without surrendering a pre-registered test. They fall into two blocks.

The discovery block (R1–R3). R1 — DFSZ depth at every tuning — is not a sensitivity goal but an identity: Part V places the Kaxion *on* the DFSZ line by theorem ($E/N = 8/3$), so an instrument shallower than DFSZ is not a less sensitive Kaxion search, it is no Kaxion search at all. R2 — window coverage of $\nu_0 \in [1.59, 2.13]$ GHz — is the closed window of the discrete Ξ_{cap} ladder: the theory permits no line outside it, so the instrument must be able to reach all of it, and

nothing beyond it is required. R3 — two sub-arrays independently tunable across the window — is the requirement that turns a detector into an identifier: it doubles the deep-camp exposure when both sit on one rung, halves the window-completion time when they split, and runs the count test of Part VI live whenever they differ. The count test is the Kaxion’s signature, and carrying both sub-arrays into one common-clock analysis means every second of split-setting data interrogates the window for a partner — the count test sits in the trigger, not in a follow-up proposal.

The modulation block (R4–R6). R4 — one steady year at $\geq 85\%$ duty — is the annual modulation’s non-negotiable price: the 74 Hz breathing of Eq. (48.1) has a period of exactly one year, and no instrument that runs in campaign fragments can trace it. R5 — absolute timing far below a minute over that year — guards the experiment’s sharpest artifact discriminant, the 3.9-minute-per-day gap between the sidereal and solar periods: GPS discipline at the 10^{-8} level makes the phase of the daily modulation a measurement rather than an assumption. R6 — centroid stability at ± 1 Hz per season — is the one requirement aimed at a systematic rather than a signal: it holds the instrument’s own seasonal drift an order of magnitude below the annual amplitude and below even the 1.2 Hz daily amplitude at Fermilab’s latitude, so that neither modulation can be manufactured or masked by the apparatus; it is implemented by the continuous calibration tone of §58 and audited by the sidereal/solar discrimination itself.

Table 27: KAXIO design requirements. Each requirement is traced to the Part that forces it and the falsifier it executes; none is discretionary. Each row names a requirement, the Part that forces it, and the falsifier it executes. The middle column is what makes the instrument an argument rather than a wish list: nothing is specified because it would be pleasant to have, and a requirement with no Part behind it would not belong here. A reader looking to cut cost should look for a row with no falsifier against it.

Requirement	Origin	Serves
R1 DFSZ depth at every tuning	Kaxion on the DFSZ line (Part V)	discovery; F-Kaxion-1c
R2 window coverage $\nu_0 \in [1.59, 2.13]$ GHz	the closed window (Part V)	F-Kaxion-1c
R3 two independently tuned sub-arrays	depth, coverage, count (Part VI)	count test; F-Kaxion-1c
R4 one-year continuous run, duty $\geq 85\%$	annual modulation (Part IX)	identification
R5 absolute timing $\ll 1$ min over a year	sidereal-day discrimination	identification
R6 centroid stability ± 1 Hz/season	modulation systematic floor	identification

57 Detector design

Figure 15 is the concept; Table 28 the specification. Two subsystems carry the physics.




Figure 15: KAXIO conceptual design. The eighteen ADMX-EFR cavities are reorganized into two nine-cavity sub-arrays inside the existing 9.4 T magnet, each family built as two cavity sets specified by bare frequency (upper 1.79 GHz, lower 1.55 GHz; exchanged once) and independently tunable across the fenced window — stacked on one rung for depth, split across rungs for coverage and the live count test (Table 28). The two sub-array streams enter a common-clock DAQ that runs the count test of Part VI live. Reused EFR infrastructure is drawn gray; new elements are the independent tuning layer and the common-clock DAQ (Table 34).

The two sub-arrays, independently tuned. The eighteen EFR cavities are rebuilt as two families of nine, each family carrying two cavity sets specified by their bare frequency (upper set 1.79 GHz, bare radius 6.4 cm; lower set 1.55 GHz, 7.4 cm; the geometry paragraph below), each tuned by its own rod and control loop across the rungs its set reaches. In the *deep-camp* setting both sub-arrays tune to one frequency — the benchmark first — and their combined volume doubles the exposure. In the *split* setting they take different rungs, halving the window-completion time and running the count test of Part VI live: any persistent excess in one sub-array is immediately interrogated for a second line anywhere the other sits. Each sub-array is coherently power-combined as in the EFR baseline. The cost of this configuration is re-machining and a control-software layer, not a new cryostat, and the control layer is simple: two independent loops, each reporting its tuning to the common clock.

Cavity geometry, stated. The numbers above imply the cavities, and the implications are engineering commitments rather than conveniences. The length is set by the EFR volume budget — about 200 L in eighteen cavities — rather than by a published EFR cavity drawing: at 116 L per nine cavities of the upper set (bare radius 6.4 cm) each cavity is 1.0 m long (aspect ratio 7.8 for the upper set, 6.8 for the lower); at those aspect ratios the TM_{010} mode crosses the TE_{11n} family roughly every 100 MHz across the tuning range, each crossing costs a dead band of a few megahertz that the other sub-array covers in the split setting, and a 1 m cavity can localize its mode along its length wherever the bore’s field or the rod’s straightness varies — so the crossings and the mode map are measured on the bench in the engineering year, not assumed away, and the design assumes the EFR field homogeneity over the cavity length, verified by a field map of the bore before commissioning (Table 33). The window’s 30% span exceeds the range of one rod-tuned TM_{010} cavity, so each sub-array carries two cavity sets, exchanged once at the 1.79 GHz seam, and each set is specified by its bare (rod-out) frequency, since a metal tuning rod moves TM_{010} *upward*: an upper set of bare radius 6.4 cm (1.79 GHz), rod-tuned up to 2.15 GHz (+20%) and so reaching the $\Xi_{\text{cap}} = 47$ rung at 2.13 GHz with margin, and a lower set of bare radius 7.4 cm (1.55 GHz), tuned up to 1.82 GHz (+17%) — rod ranges at the upper end of demonstrated practice, and a commitment the engineering year must meet. Both sets respect the packing limit of the 0.8 m bore — eighteen circles pack a circle only if their outer diameter is at most 16.4 cm, which the lower set meets at wall thickness ≤ 8 mm. The benchmark budget of Table 29 (95 L, $C = 0.40$) is the effective VC product at the benchmark rod position of the upper set, conservatively below the bare-cavity volume (116 L per nine); the set exchange is a scheduled line item of the duty budget (Table 33). Form factor and Q_L are carried at the EFR baseline for both sets; the antenna coupling is held over-coupled ($\beta \simeq 2$, $Q_L = Q_0/(1 + \beta)$) so that the loaded bandwidth $\nu/Q_L \simeq 20$ kHz stays ten times the linewidth.

The two-stream exclusivity layer. Both data streams — sub-array α and sub-array β — are digitized against one GPS-disciplined local oscillator, with a continuously injected calibration tone in every chain (R6) and the standard environmental veto antenna. The exclusivity demand is the background killer. The prediction is a candidate in *one* stream, at one frequency, with every other instrumented bin empty — so the trigger records everything and vetoes nothing: a single-stream persistent excess is the signature; the same excess appearing in both sub-arrays at *different* tunings marks a wideband artifact, since no halo signal drives two well-separated frequencies at once; and a second genuine line in the other sub-array would be the count-test kill of Part VI — a result to report, never to suppress. Interference, thermal, and mechanical artifact classes are exactly the ones that light multiple streams, which is what makes the split setting a background instrument as well as a coverage one.

58 The RF and readout chain

Figure 16 draws one chain end to end; two such streams run in parallel, each nine-fold coherently combined after its JPA stages as in the EFR baseline. Three design choices matter for the physics.

Table 28: KAXIO specification. Baseline ADMX-EFR values from the collaboration’s design (ADMX-EFR Collaboration, 2024); KAXIO entries are this proposal.

Subsystem	ADMX-EFR baseline	KAXIO
magnet	9.4 T MRI, 0.8 m bore	unchanged (in hand)
cavities	$18 \times \text{TM}_{010}$, one frequency	2×9 , independently tuned
cavity sets	one	two (bare 1.79/1.55 GHz)
volume (eff.)	~ 200 L	95 L + 95 L (benchmark)
form factor C_{010}	0.40	unchanged
Q_L	10^5 (SC-coated)	unchanged
T_{sys}	150 mK (JPA chain)	unchanged; photon-counting upgrade path
tuning	piezo, single	two independent piezo loops, reported to the common clock
DAQ	single-stream	2-stream common clock, GPS, cal tone
run plan	3-yr scan 2–4 GHz	1-yr camp + split-array sweep of R2

One clock: every mixer in every stream runs from a single GPS-disciplined local oscillator, so the two streams share phase provenance and the modulation analyses inherit absolute timing at the 10^{-8} level (R5). *One calibration tone:* injected at every antenna through a -60 dB coupler and tracked continuously in the FPGA filter banks, it measures each chain’s gain and centroid response at all times — the hardware behind the ± 1 Hz seasonal floor (R6) — and, because it is injected *after* the cavity, it cannot mimic a cavity-resonant halo line. *Real-time exclusivity:* the FPGA layer compares the two sub-array streams on a common clock before anything is flagged, so the trigger itself embodies the geometry being tested.




Figure 16: The RF/readout chain (one per data stream; each is nine-fold coherently combined after the JPA stage). Cold section at 10 mK: over-coupled antenna, triple circulator, flux-pumped JPA at +20 dB with $n_{\text{add}} \simeq 1.1$ quanta added, counting the vacuum as half a quantum ($T_{\text{sys}} = 150$ mK against $h\nu/k_B = 94$ mK; Table 29), replaced by a photon-counting front end in the upgrade path; 4 K HEMT and room-temperature stages follow. All mixers run from one GPS-disciplined local oscillator (R5); the calibration tone (red, dashed) is injected at every antenna through -60 dB couplers and tracked in the FPGA polyphase filter bank, pinning the centroid response continuously (R6); a second, weakly coupled port carries a swept probe in transmission through the cavity (S_{21}), tracking the mode to ± 10 Hz for the Lorentzian-tilt correction of Table 30. The common-clock engine aligns the two sub-array streams in real time.

Table 29: Sensitivity and noise at the benchmark ($\nu_0 = 1.95$ GHz, DFSZ coupling, $\rho_{\text{loc}} = 0.45$ GeV/cm³; registered machinery of Part IX; live time at 85% duty). Every row is computed. Halo-density systematic ($\pm 22\%$) applies to amplitudes as in Table 26 and cancels in the frequency-based tests — the count and the modulations.

Quantity	Value	Basis
P_{sig} per sub-array	1.66×10^{-22} W	Eq. (45.1); 9.4 T, 95 L (conservative: the bare cavity volume is 116 L, and the 18% margin covers the tuning rod and the mode's fill), $C = 0.40$, $Q_L = 10^5$, antenna $\beta \simeq 2$, i.e. $\beta/(1 + \beta) = 2/3$
photon rate per sub-array	1.3×10^2 s ⁻¹	$P_{\text{sig}}/h\nu$
noise density $k_B T_{\text{sys}}$	2.1×10^{-24} W/Hz	150 mK JPA chain, i.e. $n_{\text{add}} \simeq 1.1$ quanta added with the vacuum counted as half a quantum, $T_{\text{sys}} = (n_{\text{add}} + \frac{1}{2}) h\nu/k_B$ at $h\nu/k_B = 94$ mK (Y-factor $\pm 6.4\%$)
$P/k_B T_{\text{sys}}$	80 Hz	ratio of the above
$t_{5\sigma}$ per sub-array	12.8 s (live)	matched line, $\epsilon_a = 0.93$, $\Delta\nu_{\text{eff}} = 2.9$ kHz
week-1 peak-bin SNR	$\approx 2.7 \times 10^2$ per sub-array	stat. only ($\pm 24\%$ syst.); shape fully resolved
centroid precision (fortnight)	± 0.64 Hz stat.	Eq. (48.3) (FWHM 1.8 kHz), both sub-arrays stacked on one rung
daily-modulation reach	2σ in one month; $\sim 6.5\sigma$ over the year	Fig. 17
annual-modulation reach	syst.-limited at ± 1 Hz	Fig. 18; R6
window sweep on null	full R2 range, DFSZ depth	split-array sweep; executes F-Kaxion-1c

59 Sensitivity and noise

Table 29 carries the full budget, evaluated with the registered machinery of Part IX (Eqs. (45.1)–(46.1)) at the benchmark rung. The headline: each sub-array camps its tone with $P_{\text{sig}}/k_B T_{\text{sys}} = 80$ Hz at DFSZ coupling — a matched-line 5σ in 12.8 s of live time — so the electromagnetic discovery reach is never in question; the instrument's purpose is what comes *after* the excess, and its noise budget is engineered around the two locks and the two modulations rather than around raw depth.

60 Systematic errors

Table 30 extends the Part-IX budget (Table 26) with the rows KAXIO adds; the inherited rows apply unchanged. Two design choices dominate the table's shape. First, everything amplitude-like remains halo-dominated and common-mode — and the two frequency-domain verdicts — the window and the count — do not depend on it, which is why the instrument is built around them. Second,

Table 30: KAXIO-specific systematics, additional to the inherited budget of Table 26. The seasonal-drift row is the one channel through which an instrument effect could imitate physics; its mitigation is hardware (R6) and its residual is verified by the sidereal/solar-period discrimination.

Source	Size	Moves	Mitigation / method
inter-array gain	$\pm 3\%$	tone-power comparison	common cal tone; cross-swap runs
inter-array frequency ref.	10^{-8} (GPS LO)	count-test comparison	common disciplined LO
sub-array cross-talk	< -80 dB	fake coincidence	isolation + offset LO scheme
seasonal LO/thermal drift	± 1 Hz residual	annual centroid	continuous cal tone (R6); sidereal check
clock/timing	$< 10^{-8}$	modulation phase	GPS-disciplined LO (R5)
duty-cycle gaps	15%	exposure bookkeeping	persistent-mode magnet; redundancy; set exchange scheduled
cavity Lorentzian tilt	2 Hz/100 Hz	centroid	mode tracked to ± 10 Hz by a swept probe through a weakly coupled port (S_{21} , Fig. 16); tilt fitted from the tracked resonance, subtracted

the one systematic that could mimic physics — a *seasonal* instrument drift masquerading as annual modulation — is attacked in hardware (R6): a continuously injected calibration tone in every chain, referenced to the GPS-disciplined clock, measures the instrument’s own centroid response at all times; the residual after correction is held at ± 1 Hz, an order below the 74 Hz signal, and — decisively — the calibration tone does not know the sidereal day, while the signal does.

61 Simulated performance

All figures of this section are generated by the Part-IX pipeline (Fig. 11) with the instrument block replaced by Table 28 — the exercise the reproducibility section promised a proposing collaboration could do, here done for our own proposal, with the same fixed seeds and the same script, available from the author on request (§52). Nothing else changes: the theory sheet and the halo block are untouched, so every difference between the figures below and their Part-IX counterparts is attributable to the instrument alone, and the comparison is itself informative — it prices, panel by panel, exactly what the KAXIO aperture buys over a single-cavity camp.

The section is organized as the run itself would be. First the camped spectrum as it stands at the end of week one — the per-sub-array arithmetic of Table 29, 2.7×10^2 peak-bin SNR per camped week — then the daily modulation as it sharpens from week one to month one (Figure 17); then the year-long annual curve, the observable that forces the calendar requirement R4 (Figure 18); and finally the verdict clock (Figure 19). Each figure carries its statistical errors explicitly and points to the systematic table that bounds it; each is discussed in its own paragraph below, and each reappears, condensed to a milestone row, in Table 32.




Figure 17: The daily (sidereal) modulation at Fermilab’s latitude (41.8°N , amplitude 1.2 Hz), resolved within the first year. Both panels: two-sidereal-hour centroid bins at 85% duty, both sub-arrays stacked on the line; curve, the SHM expectation (Eq. (48.2)); dashed red, the drawn envelope of the fitted curve (± 1.0 Hz at one month, ± 0.3 Hz at one year; about 1.6σ on the plotted errors); bars, the statistical error of Eq. (48.3) — ± 1.5 Hz per bin after one month (left, a $\sim 2\sigma$ amplitude measurement), ± 0.44 Hz after one year (right, $\sim 6.5\sigma$). The contrast with the single-cavity camp of Figure 13 (bars ± 3.5 Hz even after a full-duty year, larger than the amplitude) is the KAXIO aperture: $P_{\text{sig}}/k_B T_{\text{sys}}$ five times larger and the two streams stacked at 85% duty, i.e. nearly an order of magnitude in centroid precision; the amplitude itself is also slightly smaller here (1.2 vs 1.3 Hz) because Fermilab lies at higher latitude than Daejeon ($\cos \lambda$).

Week one: the line. A camped sub-array saturates immediately: at 80 Hz of signal-to-noise density the matched-line 5σ time is 12.8 s, and a week at 85% duty returns a peak-bin SNR near 2.7×10^2 per sub-array (stat. only; the $\pm 24\%$ amplitude systematic scales it) with the line shape fully resolved — the same picture as the one-day panel of Figure 12, thirteen times deeper — and with both sub-arrays stacked on one rung the exposure doubles. The fence and the address are read in the first week; the count test opens the moment the sub-arrays split.

One week, then one month: the daily fingerprint. Figure 17, two panels. With 80 Hz of signal-to-noise density and both sub-arrays stacked on the line, two-sidereal-hour centroid bins reach ± 1.5 Hz within the *first month* and ± 0.44 Hz over the year: the 1.2 Hz sidereal wobble (Eq. (48.2) at Fermilab’s latitude, with the apex projection $\cos \delta_{\text{apex}} \simeq 0.67$) — out of statistical reach for the single-cavity camp of Part IX — emerges at $\sim 2\sigma$ within month one and stands as a $\sim 6.5\sigma$ measurement by year end, with the sidereal-versus-solar period test as the artifact killer once the two periods have separated — they drift apart by 3.9 min per day, one two-hour bin per month, so that discrimination is a matter of months, and the month-one milestone is a statistics milestone only. (The contrast with Figure 13 is aperture, not physics: five times the signal-to-noise density and the two streams stacked.)

One year: the annual curve as precision astronomy. Figure 18. Fortnightly centroids carry statistical errors of ± 0.64 Hz — the annual amplitude is measured at the few-per-mille level statistically, so the deliverable is still set by the ± 1 Hz engineered floor of R6. The same fit returns




Figure 18: One year of KAXIO: the annual modulation becomes precision halo astronomy (same format as the annual panel of Figure 13, calendar year from January). Fortnightly centroid fits with both sub-arrays stacked on the line: the statistical error per point is ± 0.64 Hz (Eq. (48.3) at 85% duty; drawn bars include a conservative 0.3 Hz fit floor), so the 74 Hz amplitude is measured at the few-per-mille level statistically and the deliverable is set by the ± 1 Hz engineered floor of R6 (red dashed; LO and thermal seasonal drifts, held by continuous calibration-tone injection — and cross-checked by the sidereal/solar discrimination, which no instrument effect passes). The same fit returns $v_\odot$ to $\pm 1.5\%$ and v_0 to $\pm 2\%$: a dark-matter *spectrometer*, not merely a detector.

$v_\odot$ to $\pm 1.5\%$ and v_0 to $\pm 2\%$: the detector doubles as a dark-matter spectrometer, measuring the halo it discovered.

The verdict clock. Figure 19 draws the three running observables against the four milestones of the run plan; the annual centroid’s statistical error crosses below the ± 1 Hz engineered floor near day 6, after which the annual measurement is set by R6, while the daily significance keeps growing all year — the reason the run is a year and not a season.




Figure 19: The verdict clock. The running observables of the four-step plan (Table 32) against run time at 85% duty. Blue: the significance of the daily (sidereal) modulation, 2σ at one month and $\sim 6.5\sigma$ at one year (Figure 17), growing as the square root of the exposure. Red: the statistical precision of the annual centroid, ± 0.64 Hz per fortnight with both sub-arrays stacked (Table 29), against the ± 1 Hz engineered seasonal floor (dashed, R6) — the statistical error crosses below the floor within the first week, after which the annual measurement is systematics-limited and its 74 Hz amplitude (Eq. (48.1)) is read at the floor’s precision across the year. Milestones (dotted): week one, the line with its shape resolved; month one, the daily modulation at $\sim 2\sigma$; three months, the count test complete across the window; one year, the annual modulation measured against the ± 1 Hz floor and the identification complete.

62 Staged coverage of the discrete window

The window is not a continuum to KAXIO; it is the discrete Ξ_{cap} ladder of Part V plus the seams between rungs, and that discreteness is worth exploiting explicitly, because it converts an open-ended scan into an enumerable campaign with a completion date. Table 31 states the staged plan, and its columns deserve a walk-through. The first columns give each rung’s center frequency and the cavity set that reaches it; with independent tuning, a rung is camped by one sub-array or both. The set column is the engineering content: the upper set (bare 1.79 GHz) reaches the four rungs from $\Xi_{\text{cap}} = 42$ to 47 and the upper band of the window, 1.79–2.13 GHz; the lower set (bare 1.55 GHz) reaches the three from 40 down to 35 and the lower band, 1.59–1.79 GHz; and one exchange between them is what the phrase “window coverage” costs in hardware. The schedule column implements the order argued below, band by band, because a band can only be swept with the set that reaches it: the benchmark rung first; the sweep of the upper band with the upper set installed, then that band’s remaining rungs; the one exchange; the sweep of the lower band, then its rungs. Two engineering realities are priced rather than assumed away. At the lowest rung ($\Xi_{\text{cap}} = 35$, 1.59 GHz, the lower set near its bare frequency) the two nine-cavity families still pack the 0.8 m bore, but with the thinnest clearances of the campaign — bore packing, not machining, is the binding constraint at the window’s low edge, and the engineering year verifies that packing first. And quench recovery lives inside the duty budget: an MRI-class persistent magnet quenches rarely,

Table 31: The staged campaign across the Ξ_{cap} ladder (rung values from Part V). One calendar week per rung includes retune, calibration, and the full candidate protocol; the split-array sweep covers the upper band (1.79–2.13 GHz) at DFSZ depth in weeks 2–6 with the upper set installed and the lower band (1.59–1.79 GHz) in weeks 11–13 after the one exchange (both sub-arrays exchange their sets on the null tape; on the discovery tape only the sweeping sub-array does, the camping sub-array keeping the line’s set); each band’s rungs are camped after its sweep, so that the set exchange happens once.

Ξ_{cap}	ν_0 [GHz]	cavity set	schedule
43	1.95	upper	week 1 (benchmark)
sweep	split-array, upper band 1.79–2.13 GHz at DFSZ depth		weeks 2–6 ($g_{\text{DFSZ}}/4.3$)
42	1.90	upper	week 7
45	2.04	upper	week 8
47	2.13	upper	week 9
cavity-set exchange (both sub-arrays; see caption)			week 10
sweep	split-array, lower band 1.59–1.79 GHz at DFSZ depth		weeks 11–13 ($g_{\text{DFSZ}}/4.3$)
40	1.81	lower	week 14
38	1.72	lower	week 15
35	1.59	lower	week 16

recovery is a standard multi-day procedure, and the 85% duty requirement of R4 carries it as an explicit line item (Table 33).

The order matters. A rung is ± 25 MHz wide at fixed f_K and the ladder itself is an ansatz (Part V), while a cavity bandwidth is 20 kHz: a week camped on a rung’s center interrogates 4×10^{-4} of the rung, and the prediction proper is the window. So the sweep comes first. At 80 Hz of signal-to-noise density the matched-line 5σ time is 12.8 s (Table 29); the 544 MHz window is 2.7×10^4 cavity bandwidths — 1.7×10^4 in the upper band the upper set reaches and 1.0×10^4 in the lower — and with the two sub-arrays split across a band, each sweeps its half at 5σ -local depth in 30 h of live time for the upper band and 18 h for the lower: days of calendar time, weeks with the candidate protocol of §37.3 run on everything that flickers. The benchmark rung is camped first, for one week, because it is the highest-prior 50 MHz in the field and because a week there also commissions the protocol. The sweep of the upper band at DFSZ depth follows in weeks 2–6, at a booked dwell of 302 s of live time per cavity bandwidth, which deepens it to $g_{\text{DFSZ}}/4.3$ (Eq. (51.1) without its systematic inflation, the convention of Figure 20; with the 24% inflation the same dwell reads $g_{\text{DFSZ}}/3.9$); the upper band’s remaining rungs are camped in the deep setting in weeks 7–9, one week each; the cavity sets are exchanged in week 10; the lower band is swept at the same dwell in weeks 11–13 and its rungs camped in weeks 14–16 — so that each rung’s central bandwidth carries a completed exclusion, a factor of nearly thirty below DFSZ, rather than a logbook note, while the count test runs whenever the sub-arrays differ. Every row is count-aware by construction, so each rung’s null writes a completed F-Kaxion-1c entry for that rung with the rest of the window interrogated; the remainder of the year rides on the discovery tape (modulations, count) or deepens the exclusion on the null tape. Figure 20 shows what the completed year leaves behind on the null tape: not a limit that grazes the prediction, but one that buries it a factor of nearly thirty deep at every rung’s camped bandwidth and a factor of four across the whole window.




Figure 20: The null-tape deliverable: the window closed a factor of four below the target, the camped rungs nearly thirty. The 90% CL coupling reach of the staged campaign of Table 31 (Eq. (51.1) without its systematic inflation, 85% duty, matched filter of effective width 2.9 kHz, 302 s of live time per cavity bandwidth in the sweep): the eight-week split-array sweep — five weeks on the upper band, three on the lower — holds $g_{\text{DFSZ}}/4.3$ across the entire window, and one camped week per Ξ_{cap} rung drives g_{90} to $g_{\text{DFSZ}}/28$ in the cavity bandwidth it sits on (notches; a camp covers 20 kHz of a ± 25 MHz rung, and the notches are drawn 10^3 wider than scale so that they can be seen — which is why the sweep, not the camps, is the window’s kill). The error bar on the benchmark is $\pm 1.3\%$ in mass at fixed f_K and $\pm 5.5\%$ in coupling. No haloscope result to date has excluded any axion band even a factor of a few below DFSZ across a window; on the null tape KAXIO reaches a factor 4 across the one band where it matters, and a factor ~ 30 in the seven cavity bandwidths where the ladder says to look hardest.

63 The four-step run plan

A month-zero commissioning shakedown — cooldown, tuning-system closure on both sub-arrays, and calibration-tone and synthetic-line injection runs — precedes week one inside the schedule, so the plan’s clock starts with a verified instrument. Table 32 is the contract the run signs in advance, and its two columns should be read as two complete, alternative papers. Down the left column, the discovery narrative accumulates: week one delivers the line at the benchmark with its shape resolved, and opens the count test across the split arrays; month one adds the sidereal fingerprint at $\sim 2\sigma$ — a $\sim 6.5\sigma$ measurement by year end; month three completes the count test across the window — the line alone, the whole window interrogated, the point at which the word “axion-like” hardens toward “Kaxion”; and the year completes the identification with the annual curve and the halo parameters. Down the right column, the null narrative is equally structured — rung by rung, seam by seam, to the executed F-Kaxion-1c of the final row — and it is worth noticing that the right column is *also* a strong physics program: no experiment has ever excluded a pre-registered axion band a factor of four below DFSZ across its full width (Figure 20).

The plan’s logic is the paper’s logic: the discovery case terminates in an *identification* — the fence, the address, and the count all read — and the null case terminates in the cleanest falsification in the dark-matter literature: one instrument, one year, one pre-registered window, closed at depth. Each milestone is also a natural reporting point to the agencies and to the community — four

Table 32: The four-step run plan. Cumulative deliverables at each milestone, on both the discovery tape and the null tape; every entry is a pre-registered number from this paper, none is adjustable after the run starts. The null tape is listed beside the discovery tape because both are deliverables, and only one of them is likely.

Milestone	If the line is there	If it is not
week 1	the line at the benchmark; shape = halo thermometer; count test opened across the split arrays	the benchmark’s cavity bandwidth excluded at $g_{\text{DFSZ}}/28$; the split-array sweep begins
month 1	daily modulation $\sim 2\sigma$ (a statistics milestone; the sidereal/solar test needs months); count test running as the split arrays sweep	the upper band (1.79–2.13 GHz, four of the seven rungs) swept once at 5σ -local DFSZ depth (F-Kaxion-1c’s first pass); the sweep deepens
month 3	count test complete across the window: the line alone, the whole window interrogated	the upper band at $g_{\text{DFSZ}}/4.3$ and its rungs camped (week 9); the sets exchanged (week 10); the lower band swept at the same depth (week 13); F-Kaxion-1c executed at depth across the window
year 1	annual modulation vs ± 1 Hz floor; daily modulation $\sim 6.5\sigma$; $v_\odot$, v_0 measured; identification complete	F-Kaxion-1c closed: the window a factor 4 below DFSZ, all seven rungs a factor 28 (week 16); the calibrated placement, its ladder and C-Kaxion-d retired — the Kaxion owed at $\gtrsim 6$ GHz, or the relic not this sector’s; not the framework

scheduled deliverables in twelve months, each publishable on either tape. On the existence question the year cannot end inconclusive: the two sub-arrays either find the line or close the window at depth. There is no month in which the collaboration is not producing physics.

64 Blinding, data management, and openness

A pre-registered experiment deserves a pre-registered analysis. *Blinding:* the DAQ firmware applies a sealed, common frequency offset (drawn once from $\pm[0, 200]$ Hz) to both streams jointly; the offset preserves every internal relation — the modulation periods, the line shape, the inter-array frequency comparisons of the count test — while blinding the *sub-linewidth* absolute-frequency comparisons against the theory sheet (the line’s placement within its bin and the modulation phases; a ± 200 Hz seal cannot, and is not claimed to, hide which megahertz-scale rung the camp occupies), and is opened only at the pre-declared box-opening after each run stage. *Injection campaigns:* hardware-injected synthetic lines, at blind times, blind frequencies, and blind strengths — including deliberate two-line injections that exercise the count-test veto logic — run through the full chain including the exclusivity trigger; the recovered efficiency updates ϵ_a exactly as in the field’s current practice. *Fixed analysis:* the analysis pipeline is the Part-IX script with the KAXIO parameter block (§52), frozen at run start; any post-start change is logged and reported. *Openness:* raw spectra, calibration records, and the analysis code are released publicly twelve months after each box

opening — on either tape. A theory that publishes its falsifiers in advance is owed an experiment that publishes its data on the same terms.

Two remarks on why this discipline is not ceremonial. First, the scientific weight of the entire program rests on the claim that every number was written down *before* the data — a claim that survives scrutiny only if the analysis provably could not steer toward the theory sheet; the sealed common offset delivers exactly that at the sub-linewidth level, while its relation-preserving construction means the blinded analysis still exercises every surviving test, so unblinding changes one number (the absolute frequency) and no machinery. Second, the injection campaigns do double duty: they measure the efficiency ϵ_a as in current practice, and — because injections are synthetic *lines* at hardware level, placed in one sub-array at a time — they continuously verify the one subsystem no predecessor experiment has run, the cross-array count test, against ground truth throughout the year. Governance is correspondingly simple: box-openings at the four milestones of Table 32, each preceded by a frozen-pipeline tag and followed by the twelve-month data release.

65 Risk register

Table 33 registers the program’s risks with likelihood, impact, and mitigation, graded low/medium/high on the usual project scale. One design principle controls the whole table and should be stated before it: *no single failure removes the instrument’s core deliverable*. Either sub-array alone still camps the benchmark and executes F-Kaxion-1c across the window at depth (more slowly); every risk therefore degrades the program gracefully rather than voiding the *discovery* program.

Three rows deserve individual comment. The seasonal-floor row is the only risk that touches a *claim* rather than a capability: if the calibration-tone residual exceeds ± 1 Hz, the annual-modulation measurement degrades — but the sidereal analysis, whose period no instrument effect shares, is untouched, so the identification suite loses precision, never validity. The schedule row is strategic rather than technical: CAPP’s scan front (§36.9) may reach the window first, but a single-tone detection elsewhere would make KAXIO’s count-test identification *more* urgent, not less — the risk is to priority of discovery, not to purpose. And the quench row is the reminder that the instrument rides an MRI-class magnet in persistent mode: the failure is rare, the recovery is standard, and the cost is duty cycle, which the 85% requirement already budgets.

Table 34: Cost classes. The two dominant capital items of any haloscope — magnet and cryogenic plant — are in hand at the Dark Wave Laboratory. Dollar figures are deliberately not quoted here; they belong to the engineering estimate that follows collaboration review.

Element	Class	Note
9.4 T magnet, cryo plant, building	reuse	EFR baseline, in hand
dilution refrigerator, JPA chains	reuse	EFR baseline
cavity arrays	modify	re-machine as two nine-cavity families
tuning system	modify	independent-tuning control layer (software + piezo)
two-stream DAQ + cal-tone system	new	GPS clock, injection network

Table 33: Risk register. L/M/H = low/medium/high. The controlling design principle: no single failure removes the instrument’s core deliverable, because either sub-array alone still executes the camp, the count test, and F-Kaxion-1c at depth. A register is included because an instrument proposal without one is not a proposal.

Risk	Likelihood	Impact	Mitigation / fallback
photon-counter dark rate high	M	L	staged: the first-generation JPA chain already meets the baseline plan
inter-array mechanical cross-talk	L	M	independent piezo drives; verified in the engineering year
JPA yield across 18 cavities	M	L	shared with the EFR baseline program; spares pool
seasonal floor exceeds ± 1 Hz	L	M	sidereal analysis unaffected; annual deliverable degrades gracefully, never vanishes
cavity-set exchange over-runs (week 10)	M	L	one exchange, scheduled inside the 15% duty budget; the lower band and its rungs slip, the upper band’s sweep and the count do not
long-cavity mode localization (1 m TM_{010} ; TE_{11n} crossings; bore field map)	M	M	mode map and crossings measured on the bench in the engineering year; the bore’s field map over the cavity length verified against the EFR design homogeneity before commissioning; a localized mode costs form factor, not the verdicts
stray-field transients (ramps, quenches)	L	L	SC can + servo; module pauses; duty-cycle entry only
schedule slip vs. CAPP window entry	M	L	the camp option (§37) preserves scientific priority at every stage

66 Cost class, schedule, and organization

Table 34 sorts every element of the instrument into three cost classes, and the sorting is the argument. The *reuse* rows are the budget of a flagship experiment already spent: the 9.4 T MRI magnet, the cryogenic plant, the dilution platform, and the Dark Wave Laboratory building — in any greenfield proposal these dominate the total by an order of magnitude, and here they cost

nothing. The *modify* rows are machine-shop and software line items: re-machining eighteen cavities into two nine-cavity families, each with two cavity sets specified by bare frequency (Table 31) and the independent-tuning control layer. The *new* row is the only genuinely new procurement — a handful of JPA channels from an existing production line and a commodity-FPGA DAQ layer. Deliberately, no dollar figures are quoted: they belong to the engineering estimate that follows collaboration review, and any number printed here would be mistaken for one.

Schedule, assuming the EFR baseline advances on its published course: two-family cavity engineering design in year 1; construction and commissioning in years 2–3, riding EFR’s own installation; the one-year physics run thereafter — squarely inside the discovery calendar of §36.9, where KAXIO takes the upper-edge role projected for EFR in the kill-test timeline (Figure 10) and extends it to the full window through the split-array sweep. The critical path runs through the cavities: the two arrays are conventional work that can proceed in parallel with EFR’s own construction, and the engineering year closes on two quantitative exit criteria, declared here: (i) inter-array microwave isolation verified > 80 dB end to end (Table 30); (ii) the two tuning loops closed on the benchmark rung with their frequencies reported to the common clock at 10^{-8} . Meeting both commits the build.

Organizationally the proposal is an upgrade path *within* the existing DOE Generation-2 framework, not a competing project: every EFR subsystem retains its baseline function, and the KAXIO additions are separable work packages suited to a university-group partnership alongside the host laboratory — cavity metrology and re-machining, DAQ and blinding infrastructure, and the simulation/analysis pipeline, available from the author on request. Each package has a natural single-institution owner class (machine shop and cryogenics for the cavity work, a computing-strong group for DAQ and blinding, the analysis authors for the pipeline), a deliverable that can be reviewed independently, and an interface to the baseline experiment no wider than a microwave line and a clock signal.

Relationship to the ADMX-EFR program. KAXIO is proposed *to* the ADMX-EFR collaboration, as a phase of its program rather than a rival to it: the baseline three-year 2–4 GHz campaign (§36.5) remains EFR’s own science, and the argument here is one of ordering — the Kaxion window overlaps the bottom of the EFR band, so a KAXIO phase executed first (or, at minimum, a count-test analysis carried through the baseline scan) converts EFR’s opening year into the highest-stakes test available to the instrument, with the full 2–4 GHz campaign resuming either enriched by an identified line or unburdened of the window. Whether the camp precedes, interleaves with, or follows the baseline scan is a decision that belongs to the collaboration and the agency jointly; the natural instrument is a memorandum of understanding covering the re-machining window and the data-release terms of §64. This paper supplies the physics case and the pre-registered analysis; the schedule authority is EFR’s.

67 The case for KAXIO

We close by making the case as it would be made to the funding agencies, in three parts: readiness, cost, and the singularity of the scientific opportunity.

Every critical technology is in hand or demonstrated. KAXIO invents nothing. The magnet — the cost driver of every haloscope ever built — is a commercial 9.4 T MRI solenoid already procured for ADMX-EFR and in hand at Fermilab’s Dark Wave Laboratory; the cryogenic plant, dilution platform, and building are EFR-baseline scope, with the experiment in its design-and-prototyping phase targeting operations in the late 2020s (§36.5). The cavity technology is the EFR baseline’s own: TM_{010} copper-and-superconducting-film cavities at $Q_L = 10^5$, re-machined to two radii instead of one. The quantum electronics are the field’s daily bread: flux-pumped JPAs at near-quantum-limited noise are running today in ADMX, CAPP, and HAYSTAC; single-microwave-photon counting has been demonstrated in the transmon architecture. What is genuinely new — the independent tuning loops, the cross-array exclusivity check, the calibration-tone discipline — is control software and microwave plumbing, the least risky tier of experimental technology. No subsystem of the instrument waits on a demonstration.

The cost is modest because the expensive decisions are already made. A Generation-2 dark-matter experiment normally begins with a nine-figure question: the magnet. Here that question is answered and paid. What remains is machining (two cavity families), JPA channels from an existing production line, and a DAQ layer built from commodity FPGA hardware — a university-scale work-package portfolio riding on flagship infrastructure (Table 34). The marginal cost of converting EFR’s scan into KAXIO’s count-test observatory is, by any Gen-2 standard, small; the marginal science — the difference between *detecting something* and *identifying the dark matter* — is the entire point of the field. We know of no cheaper way to buy a decisive answer.

Nothing like this target has existed in fifty years of dark-matter searching. Since the 1970s, every dark-matter campaign has searched a *space*: WIMP experiments swept decades of mass and ten orders of magnitude of cross-section, guided by naturalness bands that moved as they were excluded; axion haloscopes scan decades of frequency because the QCD axion’s mass is unknown by construction. Fifty years of exquisite instruments have therefore shared one structural weakness: no experiment could be pointed at a *single frequency with its coupling and its identification locks fixed in advance*, because no theory staked one. This paper stakes one. The window is named, and its benchmark marked, before the run; the strength sits on the DFSZ line by theorem, not by choice; the line shape, the annual and daily breathing, and the count — one line, no partner anywhere the fence allows — are published in advance with statistical and systematic errors printed on every figure, and the analysis script is available from the author on request. The field’s own candidate-signal history shows what fifty years without such locks has meant: the DAMA annual modulation, the 3.5 keV line, the 1.036 GHz excess of §37.4 — each carried *one* signature, none carried a pre-registered number, and each lingered or died in ambiguity. Here the ambiguity is removed in advance: the curves of Parts IX–X are drawn before the data, and when real data draws

them there will be little left to argue and something left to measure more precisely. KAXIO is, to our knowledge, the first dark-matter experiment in history that can *schedule its own verdict*: one year of steady running returns either the identified discovery of the dark matter — the geometry read twice, the halo measured — or the executed falsification of the sharpest prediction the problem has ever been offered. Both outcomes are publishable, historic, and final; on the existence question — discovery or window closure — “inconclusive” is not on the outcome list (§65). Agencies are usually asked to fund sensitivity. This asks them to fund a verdict. That is what one year of KAXIO buys.

The case in three sentences. The hardware exists and the methods are demonstrated. The cost is a small increment on infrastructure already bought, for the largest qualitative upgrade available in the field: from detection to identification. And the target is, to our knowledge, the first dark-matter prediction pre-registered to a single frequency with its identification locks fixed in advance — an experiment that cannot end without an answer.

That completes the constructive arc: a forced field (Parts II–III), its derived numbers (Parts IV–VI), and the laboratory that reads them (Parts VII–X). What is left is to step back from the apparatus and weigh what it all amounts to — which of the field’s old puzzles a confirmed Kaxion would settle, which it would honestly leave open, and how a single dark-matter line both answers to and is answered by the rest of the twenty-two-paper cohort. That reckoning is Part XI.

68 Part X in plain English: a specific machine, at a specific place, on a specific clock

A prediction that cannot be acted on is a claim about the world that nobody is obliged to test. This Part is the attempt to make the Kaxion actionable: a named instrument, at a named laboratory, with a run plan and a budget shape and a year.

The design is deliberately conservative. It reuses an existing magnet and an existing collaboration’s baseline, and the machining and control work it adds are the smallest changes that deliver the count-test measurement. The proposal’s novelty is in the logic rather than in the hardware, which is the right place for novelty in an experiment whose purpose is to settle a theoretical question.

The run plan is written as deliverables on two tapes: the discovery tape and the null tape. Both are listed because both are results, and the null tape is the likelier one. An experiment that only knows what to say if it wins is not designed.

The risk register at the end is there for the same reason. The controlling principle is that no single failure removes the core deliverable — if one channel is lost, the remaining configuration still executes the primary falsifiers at full depth. That is a design property rather than an assurance, and it can be checked against the table.

Part XI

Discussion

Summary, significance, the classic puzzles resolved and the innovations registered, open items, and the road to a verdict

69 Summary: the Kaxion in theory and in the laboratory

This Part steps back from the derivations and states, at the level a colleague would want in a seminar’s closing minutes, what the paper has shown and why it matters. We summarize the theory and the experiment in turn, mark the honesty ledger that separates theorem from benchmark from conjecture, and then — in the section that follows — argue the significance and uniqueness of the result against the whole dark-matter field.

69.1 The theory: a field the framework could not omit

The single most important sentence about the Kaxion is that it was not added. Every other dark-matter candidate of the last half-century is a sector *installed* to do a job; the Kaxion is a field the consistency of the six-dimensional theory *forces* into the spectrum before the dark-matter question is even posed.

The chain is short and entirely upstream of this paper. GG-A2-GG6 forces the bulk $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S_\theta^1$, with its two non-spatial coordinates — a scale interval I_r and a phase circle S_θ^1 , both dimensionless (Theorems 0 and 2). GG-A3-BI6 proves that anomaly freedom of the forced gauge group $SU(5) \times U(1)_B \times U(1)_{Q_i}$ admits exactly one Green–Schwarz two-form B_2 , of rank one, obeying the modified Bianchi identity $dH_3 = \text{tr } R^2 - \text{Tr } F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Q_i}^2$. On the four-dimensional boundary the field strength H_3 is dual to a pseudoscalar, and the holonomy of the rigid phase circle promotes to a Wilson-line angle; the circle’s holonomies are those of the rank-six group and are lifted; the tensor’s dual, the one angle that stays light, selected and stabilized as the Geometric Attractor of the GRIP runtime, is the Kaxion (Theorem T-Kaxion-1).

The logical force can be put bluntly: *delete the Kaxion and one deletes anomaly freedom* — the consistency of the quantum theory itself. The constitutional rule of Maximum Symmetry, Minimum Freedom (GG-A1-Constitution AX-4) then does the work a model-builder usually does by hand: it forbids introducing a new sector while a forced one stands available, so the search for dark matter is *ordered* to begin at the compact-angle pseudoscalar, and a new sector may be contemplated only if that one fails. Parts III and V show it does not fail.

69.2 The numbers: every observable a gear ratio of one ledger

What converts “the Kaxion exists” into “the Kaxion is at 1.95 GHz” is one calibrated scale and then a sheet of gear ratios of the *same* five integers, $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, fixed in GG-A1-Constitution–GG-A3-B16 before any astronomical input. The mass obeys the exact invariant $m_K f_K = \sqrt{\chi_{\text{QCD}}}$, legitimate because the gluonic portal (Theorem T-Kaxion-2) makes the Kaxion couple to the QCD topological charge with $N_{\text{DW}} = 1$. The decay constant is calibrated, not derived: $f_K = 7.1 \times 10^{11}$ GeV is the benchmark row at which the derived-conditional abundance $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2) = 0.2676$ of GG-A10-Cosmos is saturated by a generic $\mathcal{O}(1)$ misalignment angle, and the closure returns that angle. The photon coupling, $g_{K\gamma\gamma} = (\alpha/2\pi f_K) [8/3 - 1.92(4)]$, lands *exactly* on the canonical DFSZ line because the anomaly ratio $E/N = 8/3$ is a theorem of the SU(5) ledger. And the count is derived-conditional: exactly one light QCD-coupled angle survives the descent (Part VI), conditional on the five registered exposures of the parents (§9.4), so the line arrives alone — the window a discrete ladder around the calibrated benchmark, the address on the DFSZ line, the count one. One ledger, many readouts. A candidate whose every observable is a gear ratio of five integers is either right in a sharply testable way or wrong in a sharply testable way — exactly the epistemic position a predictions paper should occupy.

69.3 One physical line, and the discipline that fixes it

A reader could worry that the bulk and canonical normalizations amount to two predictions. They do not. The raw Green–Schwarz variable is a *bookkeeping* frame — the natural coordinate for bulk normalization — but the abundance integrals of Parts III and V are performed with the single physical decay constant f_K of the laboratory eigenstate, and no apparatus measures a coordinate system; it measures the canonically normalized eigenstate. That eigenstate is the single physical line, $m_K = 8.1 \mu\text{eV}$, $\nu_K = 1.95$ GHz at the benchmark row. The two discrete questions that could have multiplied the prediction — the cosmological history of the compact angle and the frame a detector reads — are closed, the frame by theorem and the history by the named conditional C-Kaxion-d (Part VII, the uniqueness argument), leaving one window, $[1.59, 2.13]$ GHz, and a mandated silence everywhere else.

69.4 The five credentials, derived rather than assumed

The textbook checklist for a dark-matter particle — neutral, cold, collisionless, stable, and produced in the right amount — is, for the Kaxion, four theorems and one calibration rather than five assumptions. Neutrality and colorlessness follow from its being a gauge-singlet boundary pseudoscalar. Coldness is a production statement: a misalignment condensate, pressureless from $t \simeq 0.3 \mu\text{s}$, with velocities $\sim 4 \times 10^{-14}$ at equality, a free-streaming length $\sim 10^{-6}$ pc and a Jeans length $\sim 3 \times 10^{-4}$ pc — nine and six orders below the scales where warm dark matter is constrained. Collisionlessness follows from the $1/f_K$ suppression of every non-gravitational coupling, and is now *computed*, not merely asserted (§73.4). Stability is an output with thirty-two orders of headroom for the physical line. And the abundance target Ω_{DM} is derived-conditional (GG-A10-Cosmos), and the

one scale is calibrated to it and declared as such. The Kaxion meets the dark-matter job description not by construction but by consequence.

69.5 The experiment: one line, two fingerprints, three instruments

The same rigidity that makes the theory falsifiable makes the experiment decisive. The Kaxion sits on the DFSZ line — the community’s canonical hard target — in a single window that two announced programs would divide between them: ADMX’s four-cavity array, designed for $5\text{--}9\,\mu\text{eV}$, over all seven addresses, and ADMX-EFR, targeting DFSZ across $2\text{--}4\,\text{GHz}$, over the topmost two. Neither has published a schedule, and no DFSZ-depth search has yet been published anywhere in the window. On top of a discovery sits an identification layer no generic search is organized to read: the count test — two independently tuned sub-arrays that ask, of every candidate, whether it stands alone anywhere in the window. A lone line in the fenced window, at DFSZ depth, with no partner anywhere the fence allows, would not merely be “an axion”; it would be a laboratory reading of the cap loading of a six-dimensional geometry under the ladder ansatz, and of the count that the boundary condition terminating its gauged baryon symmetry leaves behind.

69.6 The honesty ledger

Finally, the discipline the whole paper has kept: a sharp line between theorem, benchmark, and conjecture. *Theorem-level*: the existence of the forced pseudoscalar, its coldness, the $E/N = 8/3$ photon coupling, the ledger-locked entries of its coupling sheet, and, within the compact-angle sector, the count. *Derived-conditional*: the count across the whole boundary theory — exactly one light line — conditional on the parents’ five registered exposures (§9.4). *Model-level, resting on one named assumption each*: the decay constant, calibrated through the relic closure under an orientation convention the anomaly companion registers as open (O14); and the π -locked abundance target, which GG-A10-Cosmos carries as a conjecture. *Benchmark-level*: the cap loading $\Xi_{\text{cap}} = 43$, hence the single decay constant and the absolute masses, the marked row of the discrete ladder (Ansatz 26.2) and bracketed by the window. *Conjecture-level*: GG-A10-Cosmos’s misalignment-completeness closure (the condensate as the sole saturator of the inventory), carried as the named conjecture it is, and the inflation-era pinning of the angle (C-Kaxion-d) on which the single-patch history rests — while the misalignment angle $\theta_i \simeq 1.1\,\text{rad}$ is the generic $\mathcal{O}(1)$ value read back from the abundance at the calibrated benchmark, not a tuned input and not a derived one. The experiments of Parts VII–VIII test the theorem-level locks even if the benchmark masses shift — which is why the prediction degrades gracefully rather than catastrophically, and why the fingerprints, not the mass, are where identification lives. Earlier revisions of this paper carried readings the present one does not; the main text states only what survives, and the record of what was withdrawn, with the reasons, is Appendix F.

70 Significance and uniqueness

The Kaxion’s significance is not that it adds one more name to a crowded list. It is that, within a framework that fixes everything else — the gauge group, the matter content, the cosmological

history, and the resolution of strong CP — dark matter ceases to be an open slot and becomes a *consequence*: forced and unique. And it is that, unlike any rival, a detection would *identify* the Kaxion rather than merely bound it. We make the case in nine points.

(1) Dark matter as a consequence, not a slot. For forty years dark matter has been treated as an empty slot to be filled by whichever particle was theoretically fashionable, and the search has been organized as a sweep through that particle’s parameter space. GG-Theory reframes the problem at its root. Once a framework has fixed its gauge group, its fermion content, its inflationary history, and the field that relaxes the strong-CP angle, it has no remaining freedom to “choose” a dark-matter particle: whatever plays that role must already be in the spectrum, and the one scale it leaves open is calibrated once and declared. The Kaxion is that field, and the reframing is the deepest significance of the result — dark matter is not something the theory accommodates but something it predicts.

(2) Forced, not installed. Every leading rival is a sector *added* to account for the missing mass: a stable weak-scale multiplet for the WIMP, a global $U(1)_{PQ}$ and a complex scalar for the axion, a gauge-singlet fermion for the sterile neutrino, an ultralight modulus for fuzzy dark matter, a tuned primordial spectrum for primordial black holes. The Kaxion is forced by anomaly freedom before the question is posed — the least-new sector the theory contains. The honest comparison is therefore never “our axion versus their axion” but *a forced sector versus an installed one*, and the distinction is not rhetorical: a forced field carries no adjustable existence, one declared calibration, and no freedom to be relocated when a search comes up empty.

(3) The conditional exclusion of the field, within the framework. The strongest — and most falsifiable — claim of the paper is that, *within GG-Theory*, the rivals are not merely less elegant but excluded, each by a constitutional clause the theory adopted for independent reasons: no supersymmetry and no spare weak-scale matter for the WIMP; no Peccei–Quinn scalar and no residual CP problem for the axion; a fixed three-generation lepton sector against the sterile neutrino; a forced gauge group with no spare $U(1)$ carrying a dial against the dark photon (the one massless vector the parents register has open mixing and is carried as such); a fixed GUT-adjacent inflationary spectrum against primordial-black-hole dominance (Part IV). The candidates the data now disfavor — fuzzy and warm dark matter bounded by JWST and Lyman- α , the WIMP squeezed between the neutrino fog and the LHC — are exactly the ones the constitution forbids. The observational arrow and the structural arrow point the same way, which is the most an exclusion argument can ask for. The claim is conditional on the framework, and we hold it as exactly that — a structural exclusion *within GG-Theory*, not a model-independent one; because the framework is sharply falsifiable, the conditional can be tested rather than merely asserted.

(4) Predicted, not fitted. There is no continuous dial left free anywhere in the Kaxion sheet: one calibrated scale, stated as such, then one mass (with a $\pm 1.3\%$ Standard-Model input), one coupling, one count. This is the anti-landscape posture of the whole program made concrete — a correct theory should look like a compression, not a catalog of adjustable stories. The Kaxion is

precisely what a DFSZ axion becomes when every one of its free parameters — the decay constant, the misalignment angle, the model index E/N , the cosmological history — is replaced by a theorem, a calibration named as one, or a fenced conjecture, with two fingerprints added that no Peccei–Quinn construction can carry.

(5) Identifiable, not merely reachable. The haloscope community has organized itself around *reach* — pushing the excluded coupling ever lower. The Kaxion adds an orthogonal and rarer resource: *shape*. A generic axion detection *starts* an argument — which model, which E/N , which history produced it? — whereas a Kaxion detection *ends* one: a lone line in the fenced window, at DFSZ depth, with no partner anywhere the fence allows, is consistent with nothing else. With respect to specificity this is a sharper template than the Higgs: the Higgs was discovered without its mass predicted in advance, whereas the Kaxion arrives with its window, its coupling depth, and its count all written down. Identification, not reach, is where the Kaxion is falsifiable — and that is by design, the organizing philosophy of the detection program.

(6) A new kind of evidence: a dark-matter measurement that reads a geometry. The count is not an incremental discriminant; it carries information across an extraordinary span of physics. A lone line — one tone in the fenced window, with no partner anywhere the fence allows — would be a direct readout of a compact dimension’s cap loading under the ladder ansatz, and of the boundary condition that terminates the bulk’s gauged baryon symmetry; a *second* line inside the fence would refute the count and name which of the parents’ registered modes is light — the Kaxion, its address and the window would stand, the count would not — and the paper stakes the count on its absence. Few single measurements in physics would carry information from a microwave receiver to a six-dimensional geometry and a unification-scale anomaly structure at once; the Kaxion’s count does.

(7) Cross-domain exposure. The same five integers that fix the Kaxion fix the sister predictions: P1-KK’s 3.85 TeV radial lock cluster at the LHC and P2-N3’s generation count. The cohort is therefore deliberately *cross-exposed* — a falsification of a ledger-level claim anywhere impeaches the ledger everywhere (a window null, which retires a calibrated placement, does not), and a confirmation anywhere is collateral support for the rest. The Kaxion’s ledger-level claims — the address, the count — can be struck not only by a haloscope but by a collider or a neutrino experiment; a framework that can be shot at from three independent fields is, by the falsification standard of this series, the stronger for it.

(8) The redirection. The practical significance, stated as a message to the dark-matter community, is a redirection of the search. If the framework holds, the multi-decade WIMP, sterile-neutrino, primordial-black-hole, and fuzzy-dark-matter programs are looking where the answer is not, and the dark matter is a μeV pseudoscalar at 1.95 GHz, standing alone in a window fenced in advance — reachable by instruments already built or designed. We state this as a falsifiable, two-sided proposition for the community to test, not as a rhetorical flourish.

(9) The two fibers, read on a laboratory dial. The cohort’s recent campaign (Arisaka, 2026, W1-Symmetry; Arisaka, 2026, P5-CPT; Arisaka, 2026, P1-KK) adds a last register of significance. Its instrument is one exact statement: *holonomy requires closed loops, and an interval has none* — the bulk’s compact circle carries everything quantized and protected (a group, the fiber that remembers), while the scale interval carries only one-way transport (a semigroup, the fiber that forgets). The Kaxion is a creature of the circle: its parent is the periodic pseudoscalar direction of the compact geometry (§9), the tensor’s dual, whose circle-holonomy siblings the count lifts; its lightness is the circle’s own protection — a periodic direction admits no polynomial potential — and the quantized features a haloscope would read — the DFSZ address, the alignment that retires strong CP — are ledger integers on the fiber that remembers, which is why neither is adjustable; the fenced window is the ladder those integers set by ansatz around the one calibrated scale, and the count is conditional on the parents’ registered modes. Its *address* on the dial, by contrast, is the descent’s bookkeeping: the decay constant closed from the π -locked inventory, the misalignment abundance accumulated along a one-way cosmic history, the μeV mass and the 1.95 GHz line. A detection would therefore be the same double reading the collider companion (Arisaka, 2026, P1-KK) describes at the opposite end of the energy ladder — *identity* from the circle, *address* from the interval — the TeV lock and the μeV line standing as the two laboratory ends of one anatomy. In the campaign’s shorthand: *the haloscope reads the circle; the descent set the dial*. The full anatomy — every structure of this paper sorted by fiber, the triptych with the electron and the proton, and the inevitability argument — is developed in the body at §10.5 and §10.7, with Figure 6 as its portrait.

71 KAXIO: a flagship for the dark sector

The instrument that Parts IX and X build deserves a word of its own, because it changes what a dark-matter search *is*. For fifty years the field has run as a survey: sweep a parameter space, exclude a band, move on. KAXIO is not a survey. It is a purpose-built observatory aimed at a single named target, designed from the theory outward to read the fence, the address, and the count at once, and engineered so that one year of steady running cannot end without a verdict on the existence question. Part IX draws the data it would return, computed in advance, seed by seed, every error bar fixed before the first photon; Part X shows that the machine to take that data asks for nothing the field does not already have, riding on infrastructure already bought.

The right comparison is not to the previous haloscope but to the instruments that founded new observational sciences. LIGO did not improve on an earlier gravitational-wave detector; it opened a window that had never been open, and a century-old prediction became a measured waveform. JWST did not sharpen an existing infrared survey; it reached a depth and an epoch no instrument had reached, and the early universe became a place one could look at rather than infer. KAXIO is of that kind. If it hears the note, it will not merely add a particle to the table — it will open the *dark sector* to direct laboratory observation for the first time, turning five-sixths of the matter of the universe from a gravitational rumor into a substance with a measured mass, a measured coupling, and a count — one line, standing alone — that reads the terminated geometry it descends from. That is a new window, not a sharper view through an old one.

And the analogy runs deeper than ambition. LIGO and JWST were each staked on a firm prediction and built to a firm specification, so that their first light was a test rather than a fishing trip. KAXIO inherits that discipline from the theory it serves: the frequency is named, the fingerprints are locked, the calendar is printed, and the failure mode is as sharp as the success. A flagship is an experiment whose result — *either* result — is worth the building. By that standard an instrument that will either identify the dark matter or retire the sharpest prediction the problem has ever been offered, within a decade and on infrastructure already in hand, is a flagship for its era.

72 Uniqueness, nobility, and innovation

Three words a physics paper should use carefully, so we define them operationally before spending them. *Uniqueness* is what this paper does that no neighboring construction does, stated so that a reader can check the negative — that case was made in the nine points of Section 70 and is not repeated here. *Problems solved* are the standing puzzles of the dark-matter and axion literature that the construction resolves or dissolves, registered as PS-Kaxion-1 through PS-Kaxion-8, each closing with an honest verdict and its conditionality. *Innovation* is the specific technical moves that did not exist before this paper, registered as I-Kaxion-1 through I-Kaxion-10 in before-and-after form. The two registries carry cohort-global identifiers (Appendix E.2) so that companion papers can cite an entry rather than paraphrase it. A closing subsection states the conduct the paper binds itself to; the three together are related but not the same, and conflating them is how strong claims come to sound like salesmanship.

72.1 The classic dark-matter puzzles, resolved or dissolved

PS-Kaxion-1: Why has half a century of searching found nothing? The puzzle, stated honestly, shadows every dark-matter review: if five sixths of the matter is particles, why have four decades of magnificent instruments — tonne-scale xenon, cryogenic crystals, indirect and collider hunts — returned only bounds? The conventional answers are despairing (it is only gravitational) or evasive (it is weaker, or heavier, or lighter). This paper’s answer is structural: the searches were organized around *installed* candidates, WIMP-first, because installation freedom made those candidates theoretically fashionable — while the *forced* candidate sits at μeV mass and photon-frequency couplings, in a regime instruments have only just learned to reach. On this reading the fifty-year silence is not evidence of absence; it is a statement about which mass decades were listened to, and PS-Kaxion-1 dissolves into an instruction: listen at 1.95 GHz, with the Search Card of Part VII. *Dissolved — conditional on the framework, and testable within the decade.*

PS-Kaxion-2: The axion window as a landscape. The QCD axion is the field’s best-motivated light candidate and simultaneously its least predictive: the decay constant is a free parameter, so the mass window spans 10^{-12} – 10^{-2} eV — ten orders of magnitude, each demanding its own detection technology, a scan program measured in careers. Here the window collapses by one calibration and a ladder: the decay constant’s benchmark is calibrated through the abundance ($f_K \simeq 7.1 \times 10^{11}$ GeV at the cap loading 43), the ladder of loadings around it is read from the integers by ansatz, and

with $m_K f_K = \sqrt{\chi_{\text{QCD}}}$ exact, ten orders of magnitude become the 0.42-octave band $[1.59, 2.13]$ GHz with a marked benchmark rung (Part V). What was a landscape is now a target. *Reduced — to one calibrated scale and a discrete ansatz-graded window; the line at fixed f_K carries $\pm 1.3\%$.*

PS-Kaxion-3: The quality problem. The acknowledged open wound of every Peccei–Quinn construction: a global $U(1)_{\text{PQ}}$ is not respected by quantum gravity, and Planck-suppressed operators generically displace the relaxed angle by far more than the 10^{-10} the neutron bound allows — so the PQ solution to strong CP must be protected by hand against the very UV physics it ignores. The Kaxion is structurally protected: its shift symmetry is not a global symmetry of any scalar potential but the gauge invariance of the two-form, and gauge redundancies cannot be broken by gravitational operators; its only breakings are the ledger’s own topological densities (Section 24.3). The protection is the model-independent string axion’s, with the entire structure forced rather than chosen. *Dissolved at the root — no PQ scalar exists to have a quality problem.*

PS-Kaxion-4: The abundance coincidence. Why should the dark and baryonic densities sit within an order of magnitude of each other — $\rho_{\text{DM}}/\rho_b \simeq 5.3$ — when in every standard picture they are set by unrelated physics, one by a freeze-out or misalignment accident, the other by baryogenesis? Here neither is an accident: the π -locked inventory fixes the ratio as arithmetic, $\rho_{\text{DM}}/\rho_b = (1/\pi - 1/2\pi^2)/(1/2\pi^2) = 2\pi - 1 \simeq 5.28$ (Part III), tested against the CMB inventory at the percent level. Honesty note, carried from Part III: the π -locked fractions are among the *weaker* members of the cohort’s track record — π/e -fraction assignments rather than rigid integer locks — and the verdict is graded accordingly. *Resolved conditional on the inventory locks; killed by any converged departure of the CMB inventory from the locked fractions.*

PS-Kaxion-5: The initial-angle problem. Every pre-inflationary axion cosmology ends at an embarrassment: the relic abundance scales as θ_i^2 , and the angle is a cosmic accident — so the literature scans it, or anthropically excuses it, and the “prediction” inherits a free parameter through the back door. Here the angle is read back rather than scanned: the benchmark is calibrated so that a generic angle saturates the π -locked target, and the closure returns $\theta_i \simeq 1.1$ rad — a generic value, nowhere near the tuned edges (Section 15.3). What is checkable is that it stays generic: a confirmed line whose absolute rate implied a tuned angle would strain the closure visibly. *Recast — the last free parameter of axion cosmology is not derived here, but it is the generic value at which one declared calibration closes.*

PS-Kaxion-6: The defect and isocurvature hazards. The two standard axion scenarios each carry a cosmological tax. Post-inflationary: strings, domain walls, and miniclusters, with an order-unity abundance uncertainty that simulations have fought over for a decade. Pre-inflationary: isocurvature perturbations that CMB bounds turn into a stringent inflation-scale ceiling. The Kaxion pays neither tax *if* one conditional holds: with no radial order parameter anywhere in the spectrum, thermal patch decorrelation has no mechanism, so no strings and no walls; but de Sitter decorrelation during GG-A10-Cosmos’s primary inflation ($H_1 \simeq 1.5 \times 10^{13}$ GeV, from $r = 0.0038$ and $A_s = 2.1 \times 10^{-9}$, as GG-A10-Cosmos now prints) is foreclosed only by the inflation-era pinning of

the angle, Conjecture C-Kaxion-d (Section 13.5), under which the isocurvature amplitude vanishes identically and there is no minicluster spectrum. *Conditional — dissolved if the pinning holds; a Kaxion isocurvature detection at any level, or a confirmed Kaxion minicluster population, kills the pinning and with it the single-patch clause.*

PS-Kaxion-7: The look-elsewhere pathology. Haloscope searching inherits resonance hunting’s statistical tax — scan enough bins and fluctuations are guaranteed, so a blind GHz-decade search must demand $\sim 7.4\sigma$ local for 5σ global — and its sociological one: candidates that evaporate (the 1.036 GHz episode being the field’s freshest lesson) teach the community to discount all excitement. A pre-registered window pays far less: the trials factor collapses from a decade of frequency to the 2.8×10^5 bins of the window, so that a global 5σ costs a local 7.0σ rather than 7.4σ (§37.1), and the prediction was in print, with its rate, its line shape, and its two fingerprints, before the candidate was reported. *Dissolved — for this search alone, which is rather the point.*

PS-Kaxion-8: The identification problem. The endgame nobody plans for: suppose a haloscope finds a line. Which axion is it — KSVZ, DFSZ, which E/N , which history? A generic detection *starts* that argument, and nothing in a single power excess can end it; the community’s roadmaps are candid that post-discovery model identification is an open problem. The Kaxion arrives with the argument pre-settled: a lone line in a window fenced in advance, at DFSZ depth, with no partner anywhere the fence allows, is consistent with little else on the market — the fence and the address read the bulk’s ledger, the count reads its termination, and all three are immune to the halo uncertainties that blur absolute rates (Part VI). Detection and identification collapse into the same run. *Resolved by design — the organizing philosophy of the entire detection program.*

72.2 The structural innovations

The register of what did not exist before, each entry in before-and-after form:

I-Kaxion-1 (the one-calibration dark-matter sheet). *Before:* every dark-matter candidate carried at least one continuous dial — a cross-section, a decay constant, a mixing angle — scanned per analysis. *After:* coupling, abundance target, and count derived; the decay constant calibrated through the abundance at a marked benchmark, with the residual a discrete ladder and a $\pm 1.3\%$ Standard-Model input on the line. To our knowledge the first dark-matter candidate published with no continuous freedom left *free* anywhere in its chain — one scale calibrated and declared, one integer open, and the misalignment angle read back as generic.

I-Kaxion-2 (existence by consistency). *Before:* dark-matter sectors installed to fill the observational slot, then tuned to survive. *After:* the field’s existence forced by anomaly freedom before the dark-matter question is asked (T-Kaxion-1), reversing the burden of proof — it is the *absence* of the Kaxion that would require new physics.

I-Kaxion-3 (the PQ-less strong-CP relaxation). *Before:* solving strong CP meant installing a global symmetry and defending it against gravity. *After:* relaxation by a gauge-topological

angle whose protection is a redundancy, not a symmetry — the quality problem dissolved structurally rather than defended numerically (PS-Kaxion-3).

I-Kaxion-4 (the count test). *Before:* haloscope signals characterized by frequency, power, and line shape — none of which identifies a model. *After:* a pre-registered count — exactly one line inside a fenced window — testable by reanalysis of any scan that crosses the window, at zero hardware cost: the first proposed laboratory observable that reads a compact dimension’s termination (Part VI).

I-Kaxion-5 (the closed coupling sheet). *Before:* axion couplings to photons and matter treated as independent model parameters. *After:* a coupling sheet descending entry by entry from one anomaly ledger — the photon coupling on the DFSZ line by theorem, the baryonic entry fixed at $-2/5$ of X_4 with no channel below M_{lock} — to our knowledge the first proposed dark-matter candidate whose couplings are fixed in advance rather than scanned (T-Kaxion-3).

I-Kaxion-6 (camping statistics). *Before:* look-elsewhere penalties and scan-rate economics accepted as the price of axion searching. *After:* a pre-registered campaign whose trials factor is the declared window’s — 2.8×10^5 bins, local 7.0σ for a global 5σ , and 2.6×10^4 bins, local 6.7σ , on the benchmark rung, with no one-bin collapse claimed (§37.1) — quantified against the blind-scan alternative, and executed live — the 1.036 GHz candidate faced F-Kaxion-1d as a real-time test and the framework survived it (§37).

I-Kaxion-7 (the reproducibility-grade discovery simulation). *Before:* sensitivity projections quoted as reach curves, with the analysis machinery private. *After:* the data of discovery week computed in advance — spectra, modulations, count-test acceptance — from fixed seeds, with complete input tables, justified systematics, worked equations, a pipeline flowchart, and a standard-library script, available from the author on request, that an expert can re-run and dispute line by line (Part IX).

I-Kaxion-8 (KAXIO, the count-test observatory). *Before:* haloscopes as single-channel scanning instruments, upgraded for speed. *After:* an identification instrument — two sub-arrays independently tuned across a fenced window, one common-clock analysis, one year of steady running — designed so that the run cannot end inconclusive on the existence question: every exit (discovery, count-kill, window-kill, location-kill) is a verdict (Part X).

I-Kaxion-9 (the falsification calendar). *Before:* falsifiability claimed in principle, undated — axion windows in particular treated as open-endedly patient. *After:* four registered kill criteria with checkbox form (the Search Card), a discovery calendar tied to the published scan fronts, and a window that closes, on this paper’s reading of those roadmaps, by the early 2030s, with no relocation clause (Part VII).

I-Kaxion-10 (the identification taxonomy). *Before:* the haloscope program organized around *reach* — lower couplings, wider scans — with post-discovery questions deferred. *After:* a division of labor derived from the theory’s own observables: camp for discovery, split the

arrays for the count, run the year for the modulations — each step costed, ordered by expense, and assigned its falsifier (Parts VII–X).

72.3 The nobility clause

“Nobility” is not a word for the theory; it is a word for the *conduct* the paper binds itself to, and it can be stated as four commitments, each checkable. *Pre-commitment*: every falsifier is registered with its decision rule, its cost, and where possible its date, before the data — the window was staked in print while the scan fronts were still years below it, and the blinding protocol seals the analysis against the target’s own publicity (§64). *Honest grading*: every quantitative claim carries its status in the master ledger, and the least flattering entries are displayed most prominently — the named conjectures of the deck (Table 41) and the calibration appear in the abstract’s own register, and the π -locked fractions are flagged as the track record’s weaker members at every use. *No retreat*: when the 1.036 GHz candidate appeared, the framework’s response was to submit to F-Kaxion-1d in public rather than accommodate the line — and the window carries no relocation clause by construction, because a confirmed line outside it kills the calibrated placement and its ladder rather than moving the prediction. *Symmetric stakes*: the paper commits to what failure means — a swept and silent window retires the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d, and sends the Kaxion to $\gtrsim 6$ GHz or the relic out of this sector — as explicitly as to what success means. These commitments cost rhetorical safety and buy the only currency a fifty-year-old search still respects: a claim whose author has arranged, in advance, to be unable to escape.

73 Future directions, open items, and recommendations

73.1 The experimental road to a verdict

The path from prediction to verdict is staged and, at every stage, decisive. *Now*: a reanalysis of existing 1–2 GHz data tests the count at zero hardware cost. *When the programs arrive*: the announced programs — ADMX’s four-cavity array over all seven addresses, with no stated sensitivity, and ADMX-EFR’s 18-cavity array at DFSZ depth over the topmost two — together with a CAPP-class cavity entering from below would cover [1.59, 2.13] GHz at the depth the Kaxion requires — discovery, or closure of the window. None has published a date for the window; the dates this paper prints for it (Table 32, the timeline) are its own extrapolation from the programs’ stated rates, and are labeled as such. *2030s*: KAXIO (Part X), the count-test observatory, integrates all of it — the fenced window, the live count test, both modulations — into a single continuous year, so the run cannot end inconclusive on whether the Kaxion is there. The falsifiers are pre-declared and hard: DFSZ-depth exclusion of the window (F-Kaxion-1c), a compact-angle-like line at any other frequency (F-Kaxion-1d), or a second line inside the fence, which the count (T-Kaxion-5) forbids. The first two retire the calibrated placement, its ladder and Conjecture C-Kaxion-d — the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinned branch, or the relic is not this sector’s, and neither is a kill of the framework; the third refutes the count and names a light registered mode.

73.2 Cross-paper interfaces and the division of labor

Two upstream interfaces are exercised by this derivation and are recorded here for the reader who moves between the papers. First, the misalignment normalization: the abundance pipeline is stated once, in Part V, with the integrated prefactor 0.135 at the $f_K/10^{12}$ GeV normalization; GG-A10-Cosmos prints the same integrated prefactor at its misalignment equation (Section 25.5 records the one disagreement the two papers still carry); the division of labor is single-decay-constant — the benchmark is calibrated here, the observed abundance returns the misalignment angle, and neither paper is free to renormalize independently of the other. Second, the runtime vocabulary: the Ground Geometric Attractor (GGA, defined in the capstone GG-A11-Origin and used throughout this paper, Section 10.4) lives in GG-A6-GRIP’s runtime vocabulary, whose GFF-ST7 carries its content, together with the two-regime definitions GRIP-In/GRIP-Out (def:GRIPin, def:GRIPout) on which this paper’s stability and coldness arguments lean. Neither interface touches the existence theorem, the coldness, the π -locked target, the invariant $m_K f_K = \sqrt{\chi_{\text{QCD}}}$, or the fingerprints: the load-bearing results are local to their proofs, which is what lets twenty-two papers share one ledger without sharing one point of failure. A third, quantitative interface concerns the mass window itself: under the calibration convention of Part V the line carries $\pm 1.3\%$ at fixed f_K and the discrete window as its real residual; GG-A10-Cosmos has withdrawn the $\pm 6\%$ ($\pm 0.5 \mu\text{eV}$) band of the retired reading, in which the abundance pipeline’s thermodynamic errors were propagated into f_K , and prints the present budget ($\pm 4\%$, on the angle). The two papers are one prediction and print one budget.

73.3 The three named open items

Honesty about what would close the remaining gaps, in increasing order of difficulty. (i) *The misalignment angle — generic, by construction of the calibration.* With the bulk normalization calibrated at the benchmark, the observed abundance returns $\theta_i \simeq 1.1$ rad (Section 15.3), the generic $\mathcal{O}(1)$ value; the single-patch history that makes a single angle meaningful is a named conditional (no Peccei–Quinn scalar to restore thermally, reheating a factor 12 below f_K , and the inflation-era pinning of Conjecture C-Kaxion-d, Section 13.5).

This item is not closed by derivation but by declaration: the angle is read back from a calibration, the history it presumes is a conjecture, and the closure itself presumes misalignment completeness (MC, Table 41) — that no second cold component and no late entropy injection share the inventory; what would close MC is that accounting in GG-A10-Cosmos’s history, and what would refute it is a measured second dark component or a found line with Ω_K short of the inventory (Appendix E.5). (ii) *The mass-uniqueness index.* $\Xi_{\text{cap}} = 43$ is the marked (benchmark) row of the ladder, fixing the single decay constant; a defining computation for the cap loading, together with the orientation computation of O14 (Section 26.3), would replace the discrete window $[1.59, 2.13]$ GHz by a point — or retire the ladder. (iii) *The baryonic magnitude — now computed.* What survives the termination at laboratory energies is the chiral-level nucleon coupling; its absolute magnitude we evaluate in the next subsection.

The astrophysical (coring) half closes *negatively*: the dark-baryon portal is tens of orders of magnitude too weak to core a dwarf (§73.4), so the Kaxion predicts standard cuspy cold dark matter

and the cusp–core problem is, for it, baryonic. The nucleon-spin couplings are now computed on the KSVZ-class sheet of Part VIII ($C_p = -0.47 \pm 0.03$, $C_n = -0.02 \pm 0.03$), with one fence left standing: the sheet is completion-class dependent, and no spin port has a reachable target at this f_K .

73.4 Closing open item (iii): the dark-baryon coring magnitude

Open item (iii) and the cusp–core lever of Section 22 wait on a single number — the magnitude of the Kaxion’s dark–baryon energy transfer. We compute it here, in the cleanest baryon-number channel, and find it negligible by tens of orders of magnitude. The item closes with a definite, if negative, answer.

The threshold to beat. To core a gas-poor classical dwarf in a Hubble time by dark-matter–baryon energy exchange requires, in the optically-thin (single-scatter) regime, that each dark-matter particle exchange order-unity energy with the baryons it traverses: $n_b \sigma_{\chi b} v t_H \gtrsim 1$. With a central baryon (stellar) density $\rho_b \simeq 0.1 M_\odot \text{pc}^{-3}$ ($n_b \simeq 4 \text{cm}^{-3}$), a halo velocity $v \simeq 30 \text{km s}^{-1}$, and $t_H \simeq 1.4 \times 10^{10} \text{yr}$,

$$\sigma_{\chi b}^{\text{thr}} \simeq \frac{1}{n_b v t_H} \simeq 2 \times 10^{-25} \text{cm}^2 \quad \text{per nucleon.} \quad (73.1)$$

The portal cross-section. The dark–baryon vertex that survives at laboratory energies is the chiral-level nucleon coupling of Part VIII’s KSVZ-class sheet, $\propto 1/f_K$; the induced dark-matter–nucleon interaction is a derivative (hence momentum-suppressed) contact operator with scale f_K . Dimensional analysis gives

$$\sigma_{\chi b} \sim \frac{m_K^2 v^2}{\pi f_K^4} \simeq 3 \times 10^{-110} \text{cm}^2 \quad (v \simeq 300 \text{km s}^{-1}), \quad (73.2)$$

the double $1/f_K^2$ (with $f_K \simeq 7.1 \times 10^{11} \text{GeV}$) and the momentum factor $(m_K v)^2$ (with $m_K = 8.1 \mu\text{eV}$, at the halo’s $v \simeq 300 \text{km s}^{-1}$; at the dwarf’s 30km s^{-1} of the threshold it is a hundred times smaller still) together rendering it astronomically small. Even the most generous coherent enhancement — promoting the condensate’s full de Broglie occupation $N = \rho_{\text{DM}} \lambda_{\text{dB}}^3 / m_K \simeq 2.0 \times 10^{26}$ to a coherent multiplier and, over-generously, *squaring* it — lifts this only to $\sigma_{\chi b} \lesssim 1 \times 10^{-57} \text{cm}^2$.

The verdict. The portal therefore falls short of the coring threshold (73.1) by

$$\frac{\sigma_{\chi b}^{\text{portal}}}{\sigma_{\chi b}^{\text{thr}}} \simeq 1.5 \times 10^{-85} \text{ (single-particle)} \quad \text{to} \quad \lesssim 6 \times 10^{-33} \text{ (extreme coherent)}, \quad (73.3)$$

between thirty-two and eighty-five orders of magnitude, as Part IV states. No coherence bookkeeping closes a gap of that size, and the dominant suppression — $1/f_K^4$ with the physical decay constant $7.1 \times 10^{11} \text{GeV}$ — is robust against any plausible $O(1)$ uncertainty in the prefactor. *The dark-baryon portal does not core dwarfs; the magnitude is negligible.*

What this settles, and what it leaves. Three consequences follow, and they tighten the paper rather than weaken it. (1) *A definite astrophysical prediction.* The Kaxion forms standard cuspy cold-dark-matter halos; the cusp-core problem and the rotation-curve diversity are, for the Kaxion, baryonic and free of dark self-interaction — it neither solves nor worsens them (Section 22 now reads in this branch). (2) *The laboratory program is untouched.* The fence, the address, and the count are ledger and theorem data, not rate data — none of them depends on the portal’s absolute magnitude — so the program of Parts VII–X stands regardless of how feeble the baryonic portal is. (3) *Consistency, not contradiction.* The same $1/f_K$ weakness explains why no spin port is staged (Part VIII): a coupling far too feeble to move a galaxy is also, on the computed KSVZ-class sheet at this f_K , too feeble for any current laboratory port to read — the collider, not the receiver, carries the baryonic sector’s exposure. The chiral matrix elements themselves are no longer open: the sheet quotes them ($C_p = -0.47 \pm 0.03$, $C_n = -0.02 \pm 0.03$), fenced by completion class. (The early-universe streaming estimate of GG-A10-Cosmos rests on the same coupling but in a far denser, higher-velocity regime and through a long-range channel rather than the contact scattering computed here; GG-A10-Cosmos fences that estimate with its own kill criteria, and we do not adjudicate it.)

73.5 The path to few-percent

Section 29.2 located the line’s continuous error, $\pm 1.3\%$, in the lattice susceptibility and its real residual in the discrete window. The improvement program is therefore on the geometry’s side and internal to the series: the orientation computation behind O14 and a defining computation for the cap loading would replace the window by a rung — at which point ν becomes a few-linewidth target and the Stage-II scan of Part VIII shortens by an order of magnitude. Lattice progress sharpens the angle’s budget, not the line’s; the address was exact from the start, the ladder is an ansatz, and the count is derived-conditional — each stated at its grade.

73.6 Relation to the sister predictions

The phenomenology branches read as one geometry tested at three scales: P1-KK at $M_{\text{lock}} \approx 3.85$ TeV (the radial lock cluster, colliders), the present paper at $6.6\text{--}8.8\ \mu\text{eV}$ (the compact-angle zero-mode, haloscopes), and P2-N3 at the generation count (the chirality ledger). All three draw on the same five integers; a falsification of any one’s ledger-level claim impeaches the ledger everywhere (a Kaxion window null retires a calibrated placement, not the ledger), and a confirmation of any one is collateral support for the other two — the cohort is deliberately exposed.

The discussion, then, tightens the claim rather than qualifying it. The one astrophysical lever the Kaxion owns is far too weak to move a galaxy, so it makes a clean, cuspy prediction instead of a hedge; the line’s absolute placement is a calibration declared as one, with the window as its honest residual; and of the tests that actually decide the matter, the address is exact, the ladder is an ansatz and the count a derived-conditional, each stated at its grade and each with its own killer. Nor does the Kaxion stand alone: it is one reading of the same five integers that fix a collider cluster in GG-P1-KK and a generation count in GG-P2-N3, so a verdict on the line is a verdict on the ledger, and through it on the whole construction.

What remains is the simplest question and the largest, and it deserves plain words rather than machinery: why does any of this matter — to physics, and to us? That is where Part XII closes the paper.

74 Part XI in plain English: what the paper has settled, and what it has not

Three things are kept apart in this Part, and a reader in a hurry could reasonably start here.

What is established: a light compact-angle pseudoscalar is forced by the framework, its portal is the only one the charge assignment allows, and the four headline quantities follow, after one declared calibration, with no continuous input from the framework side. What is fenced: several steps rest on named computations that nobody has carried out, and the ledger says which, with the margin each currently has. What is open, and stated as open rather than folded into the other two: the items in the roadmap table, ordered by how sharply they can be attacked.

The comparison with the QCD axion deserves its own sentence, because it is the closest relative and the easiest confusion. The Kaxion sits on the same line for a reason the axion literature would recognize, and the difference is not the mechanism but the freedom: the axion's decay constant is a free parameter and this one is not. That is a claim about provenance, and it is worth exactly as much as the provenance argument of the early Parts turns out to be worth.

The falsification section is the operative content. Each channel names a number and the measurement that would end it. A framework that could only be supported would not be doing physics.

Part XII

Conclusions

Why the Kaxion matters — and why finding it matters to us

75 Conclusions

75.1 A mystery a century old

For nearly a hundred years, the largest entry in any honest inventory of the universe has been a blank. In 1933 Fritz Zwicky weighed a cluster of galaxies and found it far heavier than its stars could account for; in the 1970s Vera Rubin watched spiral galaxies turn too fast to hold themselves together on visible matter alone. The conclusion has only hardened since: about five-sixths of all the matter in the cosmos is something we have never seen — something that sculpts galaxies, bends starlight, and shaped the infant universe, yet shines with no light of its own. We have weighed it to a few percent and mapped where it lies, across the sky and back to the first moments after the Big Bang. We still do not know what it is. That one unanswered question — *what is the dark matter?* — has stood at the center of physics and astronomy for a century, and a generation of magnificent experiments, built to catch the particle almost everyone expected, has so far caught nothing. The field has been handed back its oldest and widest question: not where it is, but what it is.

75.2 The first exact answer

The claim of this paper is that, for the first time, a candidate complete theory of nature answers that question with a *number* rather than a range. Most proposals name a kind of particle and then leave its mass, its couplings, and its abundance free — a vast plane to be swept, not a single point to be checked. GG-Theory does something different in kind. Because it has already fixed the forces, the particles, the history of the early universe, and the field that tames the strong nuclear force, it has no freedom left to *choose* a dark-matter particle: whatever plays that part must already be in the theory.

That field is the *Kaxion*, and the theory does not merely assert that it exists — it pins it down. The Kaxion should be a particle about a hundred-billionth the mass of an electron, whose gentle oscillation we should hear as a faint tone at a single radio frequency, near 1.95 gigahertz, inside a narrow band.

And it should arrive carrying signatures written into the geometry, none of them adjustable: the tone should stand alone — no companion anywhere the fenced window allows, because the terminated geometry, on the parents' registered accounting, leaves exactly one light angle — and it should sit at the depth the anomaly ratio $E/N = 8/3$ fixes, on the line the community already calls

its hardest target. One scale was calibrated to the observed abundance, and that is said in plain words wherever it matters; not one of the numbers after it was adjusted to fit anything. Each is a consequence, read off that calibration and a ledger of five whole numbers the theory fixed long before dark matter was on the page. That is what it means for a prediction to be *exact*: there is a right answer and a wrong one, and the universe will tell us which.

75.3 What this paper has established

Stripped of narrative, the paper’s content is nine numbers and the machinery that stands behind them — each with its status, its proof location, and the measurement that can kill it:

Table 35: The paper in nine numbers. Status grades follow the master ledger (Table 13): *calibrated* (one scale set by the observed abundance and declared), *derived* (theorem + Standard-Model input), *derived-conditional* (resting on a named open computation), *ladder* (the discrete window, an ansatz), *fitted* (the angle read back from the abundance).

Quantity	Value	Status	Proved	Killed by
1 mass m_K	$8.1 \mu\text{eV} (\pm 1.3\%)$	calibrated (via f_K)	Part V	F-Kaxion-1c/1d
2 frequency ν_K	$1.95 \text{ GHz} (\pm 25 \text{ MHz})$	calibrated (via f_K)	Part V	F-Kaxion-1c/1d
3 window	$[1.59, 2.13] \text{ GHz}$	discrete ladder	Part V	F-Kaxion-1c
4 line count	single (no partner tone)	derived-cond. (T-Kaxion-5)	Part VI	(count test, Part VI)
5 coupling $g_{K\gamma\gamma}$	$(1.2 \pm 0.07) \times 10^{-15} \text{ GeV}^{-1}$	derived (DFSZ; via f_K)	Part V	F-Kaxion-1c
6 decay constant f_K	$7.1 \times 10^{11} \text{ GeV}$	calibrated; bench. $\Xi_{\text{cap}}=43$	Part V	(via 1–3)
7 misalignment angle θ_i	$1.18 \pm 0.04 \text{ rad}$	the convention’s output	Part V	(via 8)
8 abundance Ω_{DM}	$1/\pi - 1/(2\pi^2) = 0.2676$	derived (GG-A10-Cosmos, cond.)	Part III	CMB
9 lifetime τ_K	$\simeq 4 \times 10^{32} t_U$	derived	Part V	(stability)

Around the nine numbers stand the structures this paper built to make them decidable: the existence theorem that puts the field in the theory before dark matter is mentioned (Part II); the cold, quiet condensate cosmology with its isocurvature accounting (Part III); the constitutional exclusion of every rival within the framework (Part IV); the camping strategy that confines the look-elsewhere penalty to the declared window (Part VII); the simulated discovery week with statistical and systematic errors printed in advance, and its script, available on request (Part IX); and KAXIO, the count-test observatory whose one-year run returns discovery or window closure — never an inconclusive existence verdict (Part X). None of it waits on future theory; all of it waits on data.

75.4 An answer we can check — soon

What lifts this above a conjecture is that it can be tested now, and decisively, on a timescale of a few years and at most a decade. The tone the Kaxion would make falls squarely inside the band that two programs — IBS-CAPP in Korea and ADMX-EFR at Fermilab — would listen to, at exactly the faintness the theory demands, between now and the early 2030s. A modest addition to those instruments — a receiver with a second mode listening elsewhere in the band — would not only hear the tone but confirm that it stands alone, and a year of steady listening would read its yearly and daily breathing, turning a detection into an unambiguous identification. And the test cuts both ways, cleanly. If the Kaxion is there, it will ring at the predicted pitch, at the predicted faintness, alone — and we will have found the dark matter and, with it, the first laboratory trace of an extra, non-spatial dimension of the world. If a thorough search of that band hears only silence, the placement is wrong: the calibrated pitch, the ladder that fences it, and the conjecture that held the angle through inflation are retired, and the Kaxion is then owed at a higher pitch, above 6 gigahertz, or the dark matter is not this field. That is not a kill of the framework, and we do not claim one. We say this plainly and without hedging: the prediction may be refuted, and a clean null result would be a real and valuable answer, not a defeat. It is rare in fundamental physics to meet a question this old with a test this sharp and this near. That is the opportunity.

75.5 KAXIO: a new window on the dark universe

There is a larger way to see what is at stake, and it is worth saying plainly. Every so often physics does not merely answer a question — it opens a *window*, builds an instrument that lets us observe a part of nature we had only been able to infer. When LIGO first felt two black holes collide, it did not improve an old measurement; it opened the gravitational-wave sky, and a prediction a century old became something we could hear. When *Webb* looked back toward the first galaxies, it did not sharpen a photograph; it reached an epoch we had never seen, and the early universe became a place to look at rather than a thing to guess at.

The dark matter is the last great component of the universe still on the far side of such a window. We know it is there — it outweighs everything we can see, five to one — but we have only ever felt its gravity, never touched the substance. The observatory this paper proposes, KAXIO, is built to open that window. If the Kaxion is real, one year of KAXIO would turn the dark sector from a gravitational rumor into a measured thing: a particle with a known mass and a known voice, carrying in its two fingerprints — the window and the single-line count — the depth of an unseen dimension. That is not a better dark-matter detector; it is the first telescope pointed at the dark half of the universe — and, like LIGO and *Webb* before it, its first light would either vindicate a bold prediction or force us to think again. Either way, a door closed for a century would finally stand open.

75.6 The objective limitations

Honesty, one last time, in one place. The relic pipeline rests on the named conjectures of the deck (Table 41) — misalignment completeness (the condensate is the sole saturator of the π -locked share),

compact-angle mass-uniqueness (the cap index selects the marked row), the π -locked target's own condition in GG-A10-Cosmos (C-Kaxion-c), and the inflation-era pinning of the angle (C-Kaxion-d), on which the single-patch history rests — flagged wherever they enter and carried in the master ledger; the benchmark's absolute placement is a calibration (O14) and the window is its residual under a ladder ansatz; the count is derived-conditional on the parents' five registered exposures, none of whose masses has been computed; the line's $\pm 1.3\%$ and the coupling's $\pm 5.5\%$ at fixed f_K are Standard-Model-side and shrink with lattice QCD, not with argument. No spin port is staged: on the computed KSVZ-class sheet at this f_K none has a reachable target. The modulation magnitudes are standard-halo characteristics, quoted at model level. And the whole construction stands inside a twenty-two-paper deductive structure whose foundations are argued elsewhere: a reader who rejects the constitution owes this paper nothing — except, we would argue, a week of camping time, since the falsifiers pay out regardless of what one thinks of the axioms. These are the limits as we see them; Part XIV answers the sharper versions readers will raise.

75.7 Why it matters that we find out

It is worth ending with why the answer matters beyond the satisfaction of curiosity — because the stakes are, quite literally, ourselves. In the first hundred million years after the Big Bang, ordinary matter was too smooth, and too buffeted by light, to gather into anything. What pulled the first structures together — early enough for the first stars to ignite and forge the carbon and oxygen we are built from — was the patient, invisible gravity of the dark matter, the scaffolding on which every galaxy, every star, and every planet was later raised. Remove it, and the universe stays a featureless haze: no galaxies, no suns, no worlds, and no one to wonder about any of it. The dark matter is not a curiosity at the edge of the cosmic ledger. It is the reason there is any structure at all, and therefore the reason there is anyone here to ask the question. To learn what it is made of is to learn something about the origin of our own existence. If the Kaxion is real, then the very stuff that made us possible has been humming faintly all along — a single, pure note we have only now learned how to listen for. To catch that note would be to hear, at last, an answer to the oldest question of all — what the universe is made of — and, with it, part of the story of how there came to be anyone to ask.

75.8 A personal note

Papers rarely say plainly what their author expects, so let it be said. I do not know whether the Kaxion exists. What I know is that after ninety years the dark-matter question deserves an answer that can be forced, and that writing down numbers a machine can kill is the honesty physics respects most. Every number in this paper was fixed before the experiments that will judge it; the script that simulates its discovery is available to anyone who asks; the programs that can find it or kill it are named, and their calendars are theirs to set. If the note is there, it has been sounding since the first microsecond of the universe, and the credit will belong to the experimentalists who camp on it. If it is not, this paper will have failed in the most useful way a theory can fail: quickly, cleanly, and in public. I can ask nothing better of either outcome.

For ninety years the question was where to look; from here on it is only when to listen. The frequency is the bet, the cavity is the table — and the note, if it is real, has been playing since the universe's first microsecond, at a pitch we now know by heart.

Part XIII

Technical Appendices

A graduate machine room in seven blocks — the reference dictionaries and cross-cohort naming, the conceptual on-ramp (primers and the six foundational questions), the from-scratch derivation courses behind every headline number, the detection-and-KAXIO course, the registries, audits, and promotion roadmaps, and the record of the withdrawn readings, closed by the notation

How to read these appendices — a note for the incoming graduate student. This Part is written to be read on its own, at the level of a first- or second-year graduate course in physics or astronomy, so that a newcomer can reconstruct every headline number of the paper from scratch. It is organized in seven blocks:

- **Block A — Reference dictionaries.** The abbreviations, the symbol glossary, the cohort-wide naming convention, and a one-stop table of every derived quantity with a pointer to where it is proved.
- **Block B — The conceptual on-ramp.** Ideas before formalism: four one-paragraph primers, then a six-question foundational essay answering the questions a reader actually asks.
- **Block C — The derivation courses.** The machine room — seven courses, each taking one result (the pseudoscalar from a two-form, the GFF Hessian at the attractor, the $8.1\,\mu\text{eV}$ mass and photon coupling, the baryonic ledger entry, the cap index $\Xi_{\text{cap}} = 43$, the misalignment relic abundance, and the π -locked inventory) and deriving it step by step, with the Kaxion's numbers inserted at the end. The material that derived the withdrawn doublet's split and channel ratio is kept as an archived record in Appendix F, marked as such where it stands.
- **Block D — Detection and the KAXIO instrument.** The haloscope, transfer-function, matched-filter, instrument-sensitivity, and error-budget courses behind Parts VII–X.
- **Block E — Registries, audits, and roadmaps.** The bookkeeping: the end-to-end benchmark audit, the falsifier registry, the cross-paper inheritance table, the theorem-and-conjecture deck, and the promotion roadmap for the open items.
- **Appendix F — The record.** The readings earlier revisions carried and this one does not — the doublet, the channel ratio with its spin-port instrument, the two-lock simulation — collected in one place with the reasons for their withdrawal.
- **Block G — Notation.** The symbol and object conventions used throughout.

Suggested paths. A reader who wants the ideas should read Block B and stop; a reader who wants the dark-matter physics should begin at the mass course (Appendix C.3) and read Block C forward;

a reader who wants the geometry should begin at the dualization course (Appendix C.1); an experimentalist should read Block D.

Standard graduate texts that develop this material at book length are cited in each course as it opens: for the field-theory and geometric background Nakahara (2003); Peskin & Schroeder (1995); Weinberg (1996); for axion cosmology and model building Grilli di Cortona et al. (2016); Kim & Carosi (2010); Marsh (2016); and for the cosmological background Dodelson & Schmidt (2020); Kolb & Turner (1990).

A Reference dictionaries

This block is the paper’s reference desk: the abbreviations and symbols used throughout, the program-wide naming convention that lets any item be cited from outside its home paper, and a single table listing every derived quantity with the course that proves it. A reader can return here at any point to decode a symbol or locate a derivation.

A.1 Summary of abbreviations

This appendix collects the abbreviations used in the present paper. *All program-defined names are rendered in sans-serif type* (GG-Theory, GG-Constitution, GG-6, BI-6, OGN-5, GU-5, GDOS, GRIP, GGA, ...) so the reader can see at a glance which terms belong to the GG-Theory program; standard physics abbreviations imported from the wider literature (QCD, DFSZ, WIMP, ...) are rendered in plain uppercase Roman, following the convention of the physics literature. The third column gives the *primary source* — the paper in the Arisaka-GG series, or “this paper,” where the term is introduced. Figure 21 shows how the program-defined names nest, from the framework down to the Kaxion.

Table 36: Summary of abbreviations. Program-defined names in sans serif; standard physics abbreviations in plain uppercase Roman. The typographic convention is the quickest way to tell a term this program defines from one it borrows, which matters most in the sections where both appear in the same sentence. The third column points at the paper that owns each definition, so an unfamiliar term can be chased to its source in one step rather than by searching the cohort.

Symbol	Meaning	Primary source
<i>Framework name and core program objects</i>		
GG-Theory	Grand Geometric Theory: the unification program of which this paper is the dark-matter entry	GG-A1-Constitution, Part I
GG-Constitution	The constitutional charter: four axioms (AX-1–AX-4) and seven postulates (POS-1–POS-7)	GG-A1-Constitution
GG-6	Geometry–Gauge locking in six dimensions: the 6D bulk $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$	GG-A2-GG6
BI-6	Green–Schwarz-modified Bianchi identity in 6D (forces the two-form B_2 ; fixes the trace ratio 2/5)	GG-A3-BI6

continued on next page

Table 36 continued from previous page

Symbol	Meaning	Primary source
OGN-5	The five-integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ behind every derived number	GG-A2-GG6, GG-A3-BI6
GU-5	Geometric Universal Five Constants: $\{c, G_N, \hbar, k_B, Q_B\}$ reconstructed from the bulk geometry	GG-A4-Euclid
GDOS	Geometric Dynamics of Open Systems: the 6D-bulk-to-4D-boundary projection law	GG-A5-GDOS
Kaxion	The boundary compact-angle pseudoscalar this paper identifies as cold dark matter	this paper
<i>Axioms and postulates</i>		
LC	AX-1 Locality and Causality	GG-A1-Constitution
GGEP	AX-2 Geometry–Gauge Equivalence Principle	GG-A1-Constitution
OGGEP	AX-3 Operational/Open Geometry–Gauge Equivalence Principle	GG-A1-Constitution
MSMF	AX-4 Maximum Symmetry, Minimum Freedom (no hidden continuous knobs)	GG-A1-Constitution
POS-1–POS-7	The seven admissibility postulates; POS-6 is Occam Gauge Normalization (the OGN-5 ledger) and POS-7 assigns the strong-CP resolution to the Kaxion	GG-A1-Constitution
<i>The runtime: GRIP and its parts</i>		
Gi-4	Geometric-Invariance four-pillar structure (TGG, CCI, GFF, GRIP)	GG-A6-GRIP
TGG	Total Geometric–Gauge object (Pillar I of Gi-4)	GG-A6-GRIP
CCI	Covariant Conservation Identity (Pillar II of Gi-4)	GG-A6-GRIP
GFF	Geometric Free-Energy Functional, $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$ (Pillar III)	GG-A6-GRIP
GEP	Geometry–Gauge Energy Principle (energetic block of GFF)	GG-A6-GRIP
CEP	Curvature–Entropy Principle (entropic block of GFF)	GG-A6-GRIP
GRIP	Gauge-Runtime Invariance Pipeline (Pillar IV): the descent that selects the vacuum	GG-A6-GRIP
DGI	Dynamical Gauge Invariance (initial stage of GRIP)	GG-A6-GRIP
GSSB	Geometric Spontaneous Symmetry Breaking (middle stage of GRIP)	GG-A6-GRIP
GA	Geometric Attractor (a stable endpoint of GRIP)	GG-A6-GRIP
GGA	Ground Geometric Attractor: the unique deepest GA, the vacuum the Kaxion occupies	GG-A6-GRIP, GG-A11-Origin
GRIP-In, GRIP-Out	Fixed-landscape (descent-to-GGA) and landscape-reorganizing regimes of the runtime	GG-A6-GRIP
Qi-4	Quantum-Information four-pillar structure (the boundary informational flank)	GG-A7-Qi4
QGI	Quantum-geometric single-mode harmonic-oscillator readout (detector model)	this paper, Parts VII–VIII

continued on next page

Table 36 continued from previous page

Symbol	Meaning	Primary source
<i>Named results of this paper</i>		
T-Kaxion-1	Existence theorem: the forced compact-angle pseudoscalar	this paper, Part II
T-Kaxion-2	The gluonic portal, $m_K f_K = \sqrt{\chi_{\text{QCD}}}$	this paper, Part V
T-Kaxion-3	The anomaly ratio $E/N = 8/3$ (the Kaxion on the DFSZ line)	this paper, Part V
F-Kaxion-1	The master falsifier (window excluded at DFSZ depth, a confirmed line at the wrong location, or an in-window line off the DFSZ address; the second-line clause is the count test)	this paper, Part VI
<i>Standard physics abbreviations (plain Roman, literature convention)</i>		
QCD, QED	Quantum Chromodynamics, Quantum Electrodynamics	—
EW, GUT, SM	Electroweak, Grand Unified Theory, Standard Model	—
SUSY	Supersymmetry	—
Λ CDM	Standard cosmology: cosmological constant + cold dark matter	—
CMB, BBN, BAO	Cosmic Microwave Background, Big-Bang Nucleosynthesis, Baryon Acoustic Oscillations	—
WIMP	Weakly Interacting Massive Particle	—
ALP	Axion-Like Particle	—
FDM, WDM	Fuzzy / Warm Dark Matter	—
SIDM	Self-Interacting Dark Matter	—
PBH	Primordial Black Hole	—
PQ	Peccei–Quinn (mechanism/symmetry)	—
DFSZ, KSVZ	The two benchmark QCD-axion coupling models (Dine–Fischler–Srednicki–Zhitnitsky; Kim–Shifman–Vainshtein–Zakharov)	—
NMR	Nuclear Magnetic Resonance	—
JPA	Josephson Parametric Amplifier	—
SQL	Standard Quantum Limit	—
NFW	Navarro–Frenk–White (halo density profile)	—
JWST	James Webb Space Telescope	—
ADMX, ADMX-EFR	Axion Dark Matter eXperiment; its Extended Frequency Range program	—
HAYSTAC	Haloscope at Yale Sensitive to Axion Cold dark matter	—
CAPP	Center for Axion and Precision Physics (IBS-CAPP)	—
CAST	CERN Axion Solar Telescope	—
DMRadio, MADMAX	Lumped-element and dielectric-haloscope programs	—




Figure 21: Hierarchy of the program-defined names. The framework GG-Theory and its charter GG-Constitution (axioms AX-1–AX-4 and postulates POS-1–POS-7) force the bulk objects GG-6, BI-6, OGN-5, and the constant-set GU-5; the bulk projects to the boundary law GDOS; the runtime GRIP (with free energy $GFF = GEP + CEP$ and stages $DGI \rightarrow GSSB \rightarrow GA \rightarrow GGA$) selects the deepest vacuum — the Kaxion. All names shown are program-defined and appear in sans-serif throughout the paper; this figure is the visual key to Table 36.

A.2 Glossary of key symbols

Table 37 collects the principal symbols of this paper, with their meaning and the section in which each is used. Program-defined names appearing inside symbols (e.g. the subscript on F_{GFF}) follow the sans-serif convention of Appendix A.1.

Table 37: Glossary of frequently used symbols. Each symbol is listed once, with its meaning and the context it usually appears in. Symbols that occur only inside a single proof are omitted deliberately, on the grounds that a glossary long enough to include them would be too long to consult. Where a symbol is defined in a companion paper rather than here, the context column says so, because the convention is that paper’s and not this one’s.

Symbol	Meaning	Typical context
<i>Geometry of the bulk</i>		
$\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S_\theta^1$	The 6D bulk: spacetime, a non-spatial scale interval, and a compact phase circle	GG-6; Part II
$r \in [0, K]$	Scale (renormalization) coordinate, $r = \ln(E_{GUT}/\mu)$	Part II

continued on next page

Table 37 continued from previous page

Symbol	Meaning	Typical context
θ, S_θ^1	The compact phase angle and its circle (the non-spatial sixth dimension)	Part II
B_2, H_3	Green–Schwarz two-form and its field strength	Bl-6; Part II
$\theta_{\text{GS}}, \theta_W$	Tensor shadow (dual of H_3) and Wilson-line holonomy echo	Part II
<i>The Kaxion and its couplings</i>		
a_K	The Kaxion field (the light compact-angle eigenstate)	Parts II, V
m_K	Kaxion mass ($8.1 \mu\text{eV}$)	Part V
f_K	Kaxion decay constant ($\simeq 7.1 \times 10^{11} \text{ GeV}$)	Part V
ν_K	Laboratory frequency, 1.95 GHz	Parts V–VIII
$g_{K\gamma\gamma}$	Kaxion–photon coupling (on the DFSZ line)	Part V
χ_{QCD}	QCD topological susceptibility, $\chi_{\text{QCD}}^{1/4} = 75.5 \text{ MeV}$	Part V
θ_i	Initial misalignment angle, $\theta_i \simeq 1.1 \text{ rad}$ (read from the abundance at the calibrated benchmark)	Part III
N	OGN-5 second entry, $N = 3$; the ledger identities and the cap sums are written with it	Part II
N_c	Color triplicity, the three colors of $\text{SU}(3)_C$ — what that entry counts. Distinct from the axion color-anomaly coefficient $N = 1$ of Part V	Part II
N_g	The generation (family) count, 3. <i>Not</i> a ledger integer	GG-P2-N3
N_f	Quark-flavor count of a beta function, $N_f = 2N_g = 6$ at high scale	PDG
Ξ_{cap}	Cap renormalization index ($\Xi_{\text{cap}} = K + N + L = 43$, benchmark)	Part V
β	Rotation angle of the light-against-heavy diagonalization (a rotation that preserves normalization, Part V)	Part V
<i>Integers, scales, and cosmology</i>		
$(M, N, L; J, K)$	The OGN-5 ledger, $= (1, 3, 5; 7, 35)$; $J = 2N + 1$, $K = LJ$	Part II
$E/N = 8/3$	Electromagnetic-to-color anomaly ratio (a theorem of the $\text{SU}(5)$ ledger)	Part V
M_{lock}	Lock-cluster scale, $M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K} \approx 3.85 \text{ TeV}$	GG-A9-SM; GG-P1-KK
$E_{\text{GUT}}, E_{\text{Pl}}, M_{\text{Pl}}$	GUT, Planck energy, and Planck mass	Part V
$\Omega_m, \Omega_b, \Omega_{\text{DM}}, \Omega_\Lambda$	Matter, baryon, dark-matter, and dark-energy density fractions (π -locked)	GG-A10-Cosmos; Part III
$F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$	The runtime free-energy functional and its two blocks	GG-A6-GRIP; Part II
t_U, τ	Age of the universe; Kaxion lifetime ($\tau \gg t_U$)	Part V
<i>Retired quantities (kept as the record of the withdrawn doublet configuration; Appendix F)</i>		
g_{KBB}	Retired readout: the Kaxion–baryon coupling of the withdrawn doublet configuration; the ledger entry survives as $-2/5$ of X_4 (Part VI)	App. F
$\Delta\nu/\nu = 1/K$	Retired: the withdrawn doublet’s split (Part VI)	App. F

continued on next page

Table 37 continued from previous page

Symbol	Meaning	Typical context
$\varepsilon = 1/K$	Retired: the single-transit overlap of the two compact-angle profiles, not present in the present construction	App. F

A.3 Global naming convention for cross-cohort citation

Following Arisaka (2026, A9-SM) (its Appendix C) and beginning with Arisaka (2026, A2-GG6), each load-bearing item of the GG-Theory program carries a globally unique identifier of the form *Prefix-Suffix-N*, where the Prefix names the item type and the Suffix names the paper. For the present paper the Suffix is *Kaxion*, and the families are:

- **T-Kaxion-N** for the *N*th Theorem (T-Kaxion-1 is the existence theorem of the forced compact-angle pseudoscalar, Part II; T-Kaxion-2 the gluonic portal $m_K f_K = \sqrt{\chi_{\text{QCD}}}$, Part V; T-Kaxion-3 the anomaly ratio $E/N = 8/3$ that puts the Kaxion on the DFSZ line, Part V).
- **L-Kaxion-N** for the *N*th Lemma, **C-Kaxion-N** for the *N*th Corollary (the lettered **C-Kaxion-b, c, d** of the deck, Table 41, are the fenced conjectures, lettered to keep them apart from the numbered corollaries), **D-Kaxion-N** for the *N*th Definition, and **R-Kaxion-N** for the *N*th Remark.
- **P-Kaxion-N** for the *N*th Prediction and **PS-Kaxion-N** for the *N*th Problem-Solved entry (§72.1).
- **O-Kaxion-N** for the *N*th Objection (Part XIV) and **I-Kaxion-N** for the *N*th structural Innovation (§72.2).
- **F-Kaxion-N** for the *N*th Falsifier: F-Kaxion-1c (DFSZ-depth exclusion of the unique window), 1d (a confirmed line outside the window); 1a (the partner tone) and 1b (the channel ratio) are retired with the doublet, their identifiers kept unassigned.

The canonical Suffix list for the twenty-two cohort papers is: *Constitution* (A1), *GG6* (A2), *BI6* (A3), *Euclid* (A4), *GDOS* (A5), *GRIP* (A6), *Qi4* (A7), *QM* (A8), *SM* (A9), *Cosmos* (A10), *Origin* (A11), together with the phenomenology preprints *KK* (P1), *N3* (P2), *Electron* (P3), *Proton* (P6), the foundations preprints *Real* (P4) and *CPT* (P5), and the consequence preprints *Kaxion* (C1, the present paper), *Hubble* (C2), *BH* (C3), and the white papers *Symmetry* (W1) and *Holonomy* (W2). The four axioms (AX-1...AX-4) and seven postulates (POS-1...POS-7) of Arisaka (2026, A1-Constitution) keep their existing cohort-wide names; the global identifier is the same item viewed from outside its home paper.

A.4 Comprehensive summary tables

Table 38 is the one-stop index of every quantity the paper predicts or uses, with its symbol, its value, its status (theorem-level, benchmark, or conjecture-fenced), and the course in this Part where

it is derived from scratch. A reader who wants to check any single number can enter the machine room directly at the row that names it.

Table 38: Every headline quantity of the paper: symbol, value, status, and the derivation course that proves it. “T” = theorem-level (forced by geometry and standard physics); “B” = data-calibrated benchmark; “A” = ansatz (the cap-loading ladder); “C” = conjecture-fenced (one named computation from theorem).

Quantity	Symbol	Value	Status	Derived in
Kaxion mass	m_K	$8.1 \mu\text{eV} (\pm 1.3\%)$	T (given f_K)	App. C.3
Laboratory frequency	ν_K	1.95 GHz	T (given f_K)	App. C.3
Photon coupling	$g_{K\gamma\gamma}$	$1.2 \times 10^{-15} \text{ GeV}^{-1}$	T (DFSZ line; given f_K)	App. C.3
Anomaly ratio	E/N	8/3	T	App. C.3
Baryonic ledger entry	—	$-2/5$ of X_4	T	App. C.4
Decay constant	f_K	$7.1 \times 10^{11} \text{ GeV}$	calibrated; B	App. C.6
Cap index	Ξ_{cap}	$K + N + L = 43$	C	App. C.5
Discrete mass window	ν_K	[1.59, 2.13] GHz (7 rungs)	A	App. C.5
Misalignment angle	θ_i	$\simeq 1.1 \text{ rad}$	fitted (generic)	App. C.6
Dark-matter fraction	Ω_{DM}	$1/\pi - 1/(2\pi^2) = 0.2676$	C	App. C.7
Matter fraction	Ω_m	$1/\pi = 0.3183$	C	App. C.7
Baryon fraction	Ω_b	$1/(2\pi^2) = 0.0507$	C	App. C.7
Lifetime	τ_K	$1.7 \times 10^{50} \text{ s}$	T	App. E.1
Lock scale (cousin)	M_{lock}	3.85 TeV	T	Arisaka (2026, P1-KK)
<i>Retired quantities*</i>				
Twin-peak split (retired)	$\Delta\nu/\nu$	withdrawn (Part VI)	—	App. F
Channel ratio (retired)	P_B/P_γ	withdrawn (Part VI)	—	App. F.10

* Rows marked *retired* belong to the withdrawn doublet configuration of earlier revisions and are kept as the record of it; Part VI states the withdrawal and the count (T-Kaxion-5, derived-conditional) that replaces them.

B The conceptual on-ramp: primers and foundational questions

Before any formalism, this block answers the questions a newcomer actually asks. The four primers give the one-paragraph intuition for each headline fact; the six-question foundational essay that follows then answers, at full rigor and with the physical picture beside the algebra, the questions a graduate student brings to a seminar and a skeptical colleague brings to a seminar. Each primer and each question points forward to the derivation course where the same result is proved from the metric.

B.1 Primer: Why is the Kaxion not just another axion?

Every axion model since 1977 begins the same way: a new complex scalar field is written down, given a symmetry, and arranged so that its relaxation solves the strong-CP problem; its mass

and couplings then depend on a decay constant the model is free to choose, which is why the axion band on every exclusion plot is a *band* — decades wide in both directions. The Kaxion begins at the opposite end. No field is added: the six-dimensional geometry of GG-A2-GG6 and the anomaly identity of GG-A3-BI6 already contain a two-form whose four-dimensional dual *is* a pseudoscalar, present in the theory before anyone asks about dark matter or strong CP. The constitution then forbids inventing anything further (MSMF), assigns this inherited field the strong-CP job (Problem P7), and calibrates its decay constant once, to the observed abundance, then holds it fixed — so where an axion model has a dial, the Kaxion has a declared number. The practical difference is everything in this paper: a band becomes a point, a scan becomes a camp, and “consistent with an axion” becomes “identified as the Kaxion, or dead.” The rigorous statement is Theorem T-Kaxion-1 (Theorem 8.2); the candidate-by-candidate contrast is Part IV.

B.2 Primer: Why exactly $8.1 \mu\text{eV}$?

Three numbers meet, and only one of them is set by hand — once, and declared. First, the decay constant: the scale throat of the six-dimensional bulk is $K = 35$ deep, its cap is taken to carry the discrete index $\Xi_{\text{cap}} = K + N + L = 43$, and the bulk normalization is *calibrated* through the cosmic abundance so that the benchmark loading gives $f_K \simeq 7.1 \times 10^{11} \text{ GeV}$ — the one continuous choice in the chain, stated as such (Section 26.3). Second, quantum chromodynamics fixes the conversion rate between a decay constant and a mass: because the Kaxion couples to the gluon’s topological charge, the exact invariant $m_K f_K = \sqrt{\chi_{\text{QCD}}}$ holds, and the topological susceptibility χ_{QCD} is a measured, lattice-computed property of ordinary matter — $(75.5 \pm 0.5 \text{ MeV})^4$, nature’s own dial reading (Appendix C.3). Third, the cosmic abundance closes the loop: demanding that the relic condensate saturate the π -locked share $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ at a generic misalignment angle is what calibrated the benchmark, and the angle read back is 1.1 radians (Section 15.3). Divide $\sqrt{\chi_{\text{QCD}}}$ by f_K and the answer is $8.1 \mu\text{eV}$ — a microwave tone at 1.95 GHz, on the marked rung of a window the integers span. At fixed f_K every per-cent of quoted uncertainty is Standard-Model-side; the placement of f_K itself is the calibration.

B.3 Primer: What happens the day a line appears?

Suppose a haloscope camped on 1.95 GHz records a persistent narrow excess tomorrow. Day zero: the standard vetting runs — does it persist across re-tunes, vanish with the magnet off, stay absent from the environmental antenna? Hours. Day one: the receiver asks the one question no generic signal can answer — does the line stand alone? A single-mode cavity is blind to most of the window (the tuned mode suppresses a detuned tone by tens of decibels), so this is a deliberate act: split the sub-arrays and interrogate the other rungs. If no partner appears anywhere the fence allows, that is the count’s prediction confirmed, and a result worth reporting at its depth; if a second line is there, the count is refuted and one of the parents’ registered modes is light (Part VI) — the line, its address and the window stand, and the count does not. The following weeks: the line’s shape sharpens into the halo’s thermal profile, and its center begins to drift with the Earth’s

orbit — a few tens of hertz over the seasons, a couple of hertz with the sidereal day — motions no laboratory artifact shares. The following months: the split sub-arrays finish interrogating the window, and the line is still alone. At that point the logbook does not say “an axion-like particle”; it says the Kaxion, with the cap loading read in the fence under the ladder ansatz and the count read in the silence around the line. Part IX simulates this entire week with errors printed on every panel; Part X designs the observatory that runs it as one program.

B.4 Primer: If the window is excluded, why can the frequency not simply be adjusted?

Because, after one calibration, there is nothing left to adjust. In an ordinary axion model the mass traces back to a decay constant chosen from a range spanning many decades; excluding one frequency band merely moves the expectation elsewhere, which is why axion exclusions never kill axion models. The Kaxion’s frequency traces back, step by step, to one calibrated scale, five integers and a measured QCD number: $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ fix the throat and, under the ladder ansatz, the cap index; the cap index and the calibrated bulk scale fix f_K within a seven-rung window, and $\sqrt{\chi_{\text{QCD}}}/f_K$ fixes the frequency inside $[1.59, 2.13]$ GHz — with the benchmark row marked and argued (Section 26.3). The integers are already spent elsewhere: $K = 35$ sets the TeV lock cluster of GG-P1-KK, the ledger sets the fine-structure constant and the weak angle in GG-A9-SM, the π -locks set the cosmic inventory in GG-A10-Cosmos. What is not spent is the one calibrated scale, and that is where the falsifiers bite: a swept-and-silent window (F-Kaxion-1c), or a confirmed line anywhere else (F-Kaxion-1d), retires the calibrated placement, the ladder that fences it, and Conjecture C-Kaxion-d — after a null the Kaxion is owed at $\gtrsim 6$ GHz on the unpinned branch, or the relic is not this sector’s; after a line elsewhere, the same objects are retired and the line’s own address is the new placement. Neither is a kill of the framework, and this paper does not claim one (Appendix E.2).

B.5 Six foundational questions about the Kaxion

An extended tutorial in six questions. The four primers above answer, in a paragraph each, the questions a newcomer asks first. This appendix answers the six questions a newcomer asks *next* — the ones a graduate student brings to a seminar and a skeptical colleague brings to a seminar — at full rigor, with the derivation and the physical picture side by side. Each is self-contained; each points to the machine-room appendix where the same result is proved from the metric. Nothing here is new physics: it is the paper’s own sheet, reassembled around the questions the reader actually has.

B.5.1 What is the DFSZ line, and why does the Kaxion sit exactly on it?

What is the DFSZ line for a QCD axion, and how is it computed accurately? Why does the Kaxion land exactly on it — and why is the Kaxion tied to QCD at all?

Every QCD axion couples to two photons through

$$\mathcal{L}_{a\gamma\gamma} = \frac{1}{4} g_{K\gamma\gamma} a F \tilde{F}, \quad g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K} \left(\frac{E}{N} - 1.92 \right). \quad (\text{B.1})$$

Two pieces sit inside the bracket. The ratio E/N of the electromagnetic to the color anomaly of the Peccei–Quinn current is the *model-dependent* (ultraviolet) part: it depends on the charges the axion’s parents carry. The number 1.92 is the *model-independent* (infrared) part, shared by *every* QCD axion — it comes from the axion mixing with the π^0 and η through the light quarks, and chiral perturbation theory fixes it to 1.92(4) at next-to-leading order (Grilli di Cortona et al., 2016) (the leading value $\frac{2}{3} \frac{4+z}{1+z} \approx 2.0$, with $z = m_u/m_d$, shifted down by NLO corrections).

Why a QCD axion is a line, not a point. Because the mass is also set by $1/f_K$,

$$m_K f_K = \sqrt{\chi_{\text{QCD}}}, \quad \chi_{\text{QCD}}^{1/4} \simeq 75.5 \text{ MeV (lattice (Borsanyi et al., 2016))}, \quad (\text{B.2})$$

one may eliminate the decay constant and write the coupling as a function of the mass alone,

$$g_{K\gamma\gamma} = \frac{\alpha m_K}{2\pi \sqrt{\chi_{\text{QCD}}}} \left(\frac{E}{N} - 1.92 \right). \quad (\text{B.3})$$

In the $(m_K, g_{K\gamma\gamma})$ plane this is a straight line of unit logarithmic slope whose height is fixed *entirely* by E/N . The two canonical benchmarks are KSVZ ($E/N = 0$, so $|E/N - 1.92| = 1.92$) and DFSZ ($E/N = 8/3$, so $|8/3 - 1.92| = 0.747$). The *DFSZ line is the $E/N = 8/3$ locus*; it sits a factor $1.92/0.747 \approx 2.6$ below KSVZ — fainter, and therefore the hardest experimental target. That faintness is a partial cancellation: $8/3 = 2.667$ is close to the infrared value 1.92, so the ultraviolet and pion-mixing contributions nearly subtract.

Computing the line accurately. Three inputs, each with a controlled error. (i) The mass anchor: the zero-temperature topological susceptibility from lattice QCD, $\chi_{\text{QCD}}^{1/4} = 75.5(5) \text{ MeV}$, computed from the fluctuation $\langle Q^2 \rangle / V$ of the topological charge, giving $m_K = 5.70(7) \mu\text{eV}$ ($10^{12} \text{ GeV} / f_K$). (ii) The infrared coupling 1.92(4) from NLO chiral perturbation theory, its error dominated by the light-quark mass ratios. (iii) E/N , a rational fixed by the model (8/3 here). Evaluating at $m_K = 8.1 \mu\text{eV}$, so $f_K = 10^{12} \text{ GeV} \times (5.70/8.1) = 7.0 \times 10^{11} \text{ GeV}$,

$$g_{K\gamma\gamma}^{\text{DFSZ}} = \frac{\alpha}{2\pi f_K} (0.747) = \frac{(1.16 \times 10^{-3})(0.747)}{7.0 \times 10^{11} \text{ GeV}} \simeq 1.2 \times 10^{-15} \text{ GeV}^{-1}, \quad (\text{B.4})$$

exactly the value of the main text (§27; Appendix C.3). The dominant uncertainty on the line’s *height* at fixed E/N is the ± 0.04 on 1.92: because $8/3 - 1.92 = 0.747$, that error is $\pm 5\%$ of the bracket — which is why the paper quotes $\pm 5.5\%$ on $g_{K\gamma\gamma}$ at fixed f_K , the chiral bracket dominating. The geometry contributes no continuous error at fixed f_K ; what it contributes is the calibrated placement of f_K itself and the discrete window around it.

Why the Kaxion lands exactly on it. “DFSZ line” is shorthand for “ $E/N = 8/3$,” and GG-6 *forces* $E/N = 8/3$. That value is the anomaly ratio of a complete generation — equivalently, of an $SU(5)$ multiplet ($\bar{5} + 10$) — famously both the DFSZ two-Higgs-doublet value (family-universal, independent of $\tan\beta$) and the grand-unified value. In DFSZ this is a model-building *choice*; in GG-6 it is *forced*, because the boundary fermion content and the BI-6 anomaly ledger embed the Standard-Model charges in the $SU(5)$ -structured bookkeeping (GG-A9-SM), so the electromagnetic-to-color anomaly ratio comes out $8/3$ with no freedom (evaluated as pure group theory in §27). Same E/N , same line. The Kaxion is not *tuned* onto the DFSZ benchmark; it reproduces $8/3$ for the same reason a grand-unified axion does.

Why it is tied to QCD at all. Two logically separate facts, in sequence. *First, it must couple to gluons.* The constitution demands that strong CP be solved (GG-A1-Constitution P7). The only way a periodic pseudoscalar drives $\bar{\theta}_{\text{QCD}} \rightarrow 0$ is by coupling to the gluon topological density $G\tilde{G}$ — the Peccei–Quinn mechanism. When the Green–Schwarz two-form is dualized and descended to the boundary, the BI-6 ledger carries a color entry, so the compact-angle pseudoscalar inherits a $G\tilde{G}$ coupling with unit color anomaly ($N = 1$, $N_{\text{DW}} = 1$; Appendix C.3). *Second, QCD then sets its mass.* Once it couples to $G\tilde{G}$, QCD instantons tilt its potential, and the curvature of that tilt is the susceptibility: $m_K = \sqrt{\chi_{\text{QCD}}}/f_K$. The geometry supplies f_K ; QCD supplies $\sqrt{\chi_{\text{QCD}}}$.

The picture in one line. The compact non-spatial angle is a dial; to cancel the QCD vacuum angle it must be the knob θ_{QCD} slides along to zero, so it is bolted to the gluons; anything bolted to the gluons gets a mass from the instanton tilt, $m \sim \sqrt{\chi}/f$; and its glow into light carries an unavoidable infrared piece (-1.92 , the borrowed π^0 coupling) plus a UV piece E/N that GG-6 fixes at the grand-unified $8/3$. “The Kaxion sits on the DFSZ line, the hardest target in the field” is thus the shadow of grand unification cast onto the photon coupling, with QCD holding the mass fixed underneath.

B.5.2 Why does the Kaxion carry a baryonic ledger entry, and what became of the 5:2 readout?

[Partly archived. The ledger derivation below — baryon number gauged, the locked vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$, the $2/5$ trace ratio — stands. Its last step, reading the entry as a light-channel coupling ratio $g_{KBB}:g_{K\gamma\gamma} = 5:2$, is withdrawn: the boundary condition that terminates $U(1)_B$ leaves no baryonic channel below M_{lock} (Part VI), and the entry’s laboratory-scale consequence is its absence.]

Why does the Kaxion carry a baryonic ledger entry at all, unlike a QCD/PQ axion? And what became of the 5:2 readout?

The difference is upstream, in the gauge group. GG-6’s bulk group is forced to be

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

and the decisive factor is $U(1)_B$: *baryon number is gauged, not global*. This is forced by the postulate that forbids accidental global symmetries — the observed proton stability ($\tau_p > 10^{34}$ yr) must be protected by a *local* symmetry, and gauging B forbids every proton-decay operator outright (curing the disease that kills ordinary $SU(5)$, whose X, Y bosons drive $p \rightarrow e^+ \pi^0$).

Why the Kaxion inherits it. The Green–Schwarz two-form exists to cancel anomalies, and it does so by coupling to the *entire* anomaly polynomial,

$$S_{\text{GS}} = \int_{\mathcal{M}_6} B_2 \wedge X_4, \quad X_4 \text{ with locked coefficient vector } (1, -1, -\frac{2}{5}, -\frac{2}{5}) \quad (\text{B.6})$$

over (gravity, $SU(5)$, $U(1)_B$, $U(1)_{Qi}$). The Kaxion is the boundary dual of B_2 , so it inherits a coupling to every factor B_2 talks to — including the gauged $U(1)_B$. That is the baryonic portal $\frac{c_B}{4} a_K F_B \tilde{F}_B$, with c_B the ledger-level coefficient (Appendix C.4). A QCD/PQ axion lives in a theory where baryon number is an accidental *global* symmetry, so there is no baryon gauge field for the axion to have an anomaly with; standard axions have model-dependent nucleon couplings through the light-quark sector, but not a locked gauge portal.

Why exactly 5:2: the trace-norm ratio. In the fundamental of $SU(5)$ the hypercharge and baryon generators are

$$Y_0 = \text{diag}(-\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}, \frac{1}{2}, \frac{1}{2}), \quad B_0 = \text{diag}(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0); \quad (\text{B.7})$$

baryon number lives only on the three colored (quark) entries, hypercharge on all five. Their trace-norm ratio is a two-line computation,

$$\frac{\text{Tr}(B_0^2)}{\text{Tr}(Y_0^2)} = \frac{3 \times (1/3)^2}{3 \times (1/3)^2 + 2 \times (1/2)^2} = \frac{1/3}{5/6} = \frac{2}{5} = \frac{L - N}{L}, \quad (\text{B.8})$$

with the ledger integers $L = 5$ (the $SU(5)$) and $N = 3$ (color), so $L - N = 2$ (this is Eq. (6.17) of the main text). This fixes the two Abelian entries of the coefficient vector at magnitude $2/5$.

Honestly fenced. The $2/5$ trace ratio is theorem-level exact at the anomaly/gauge level (§27); the 5:2 readout built on it is withdrawn (Part VI), because the terminated $U(1)_B$ leaves no light vector for the condensate to drive. What survives at the chiral level is the nucleon- and electron-spin sheet of Part VIII, computed and out of reach at this f_K .

The picture in one line. Think of B_2 as the theory’s anomaly bookkeeper: every gauge factor writes its curvature into the ledger with a fixed weight, and because **GG-6** gauges baryon number (to keep the proton stable by law rather than by luck), $U(1)_B$ is one of the entries. The Kaxion is the bookkeeper’s phase, so at the ledger level it talks to whatever the ledger records — but the baryonic entry opens no channel below M_{lock} , so at laboratory energies that conversation is silent. (Earlier revisions read the entry “faintly written, loudly read” by a source/response inversion; that was a mistake — dualization preserves the coefficient vector, so the ledger ratio is 2:5 for source and

response alike — and the record is Appendix F. What follows in this paragraph is retained only to mark where the inversion entered: it was said that a faint entry is read loudly by the dual field, so that baryon number — written as 2, read as 5 — would be the channel the Kaxion shouts to while whispering to photons. It is not: the entry is written as 2/5 and read as 2/5, and below M_{lock} it is not read at all.)

B.5.3 What is the cap index Ξ_{cap} , and why is it $K + N + L = 43$?

What is Ξ_{cap} , physically? Why is it expected to be $K + N + L = 43$ — and is that not just three integers added by hand?

It is fair to press here, because Ξ_{cap} is the one number in the whole mass derivation that the paper fences as a *conjecture*, not a theorem (§26.3). There is a real, non-arbitrary skeleton, and the arbitrariness is confined to exactly one unevaluated integral. Let us separate the layers.

What Ξ_{cap} is. Not a new parameter — by ansatz, a wavefunction-renormalization integer relating two normalizations of the *same* field,

$$f_K = \frac{f_{\text{bulk}}}{\Xi_{\text{cap}}}, \quad m_K = \frac{\sqrt{\chi_{\text{QCD}}}}{f_K} \propto \Xi_{\text{cap}}. \quad (\text{B.9})$$

Here f_{bulk} is the Green–Schwarz direction’s normalization deep in the bulk; f_K is the decay constant of the physical eigenstate a haloscope couples to. Because $m_K \propto \Xi_{\text{cap}}$, it sets the *absolute* mass — and nothing else. Crucially it moves only *where* the line sits, never the coupling sheet: the DFSZ depth and the ledger entries are scale-independent theorems; only the address is conjectural.

Where the factor comes from — the throat sets the normalization. On the canonical metric $ds_6^2 = e^{-2r}\eta dx^2 + dr^2 + d\theta^2$, $r \in [0, K]$, the tensor’s dual and a circle holonomy acquire wildly asymmetric kinetic norms,

$$Z_{\text{GS}} = 2\pi \int_0^K e^{2r} dr = \pi(e^{2K} - 1) \simeq \pi e^{2K}, \quad Z_\theta = 2\pi \int_0^K e^{-2r} dr \simeq \pi, \quad (\text{B.10})$$

differing by $e^{2K} = e^{70} \sim 10^{30}$ (Appendix C.1). Canonical normalization divides each field by $\sqrt{Z}$, so the map from the bulk variable to the physical eigenstate is enormous, and Ξ_{cap} is that map expressed in the throat’s own quantized units. The Green–Schwarz variable is normalized after the warp has run through its full K e-folds; the boundary eigenstate a detector reads sits at the throat’s mouth; the ansatz reads the conversion between them as the throat depth in integer units, plus whatever the cap adds.

The skeleton — one derived fact and an ansatz. One fact is derived from the canonical metric, and two are the ansatz of §26.3 (Ansatz 26.2), stated there with its grade:

1. $\Xi_{\text{cap}} \neq 1$: the two face-value chiral norms (B.10) differ by some thirty orders and neither reproduces the calibrated f_K , so the boundary eigenstate is not either bulk norm (O14 open).

(This says the ratio is not unity; it does *not* set the value, and the exponential itself proves nothing about Ξ_{cap} .)

2. $\Xi_{\text{cap}} = K + \Sigma$, an *additive integer* form — the ansatz. The throat depth $K = 35$ is an e-fold count of the warp, and the ultraviolet terminus of the interval (the “cap”) is *taken* to contribute an integer Σ ; the admissible class carries no orbifold cap, and no quantized flux equal to K is exhibited. This is the load-bearing step, and it is an ansatz: it is what makes the residual an integer rather than a continuous knob, *if* it holds.
3. Σ is a subset-sum of the remaining ledger integers $\{N, L, J\} = \{3, 5, 7\}$, and the cap can host at most a pair of the three flux types, so

$$\begin{aligned} \Sigma \in \{0, 3, 5, 7, 8, 10, 12\} &\Rightarrow \Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\} \\ &\Rightarrow \nu_K = 1.95 (\Xi_{\text{cap}}/43) \in [1.59, 2.13] \text{ GHz.} \end{aligned} \tag{B.11}$$

So, under the ansatz, the absolute mass is not a free parameter with an error bar — it is one of *seven discrete values* (Table 11), the benchmark row calibrated through the relic closure and the others ratios to it. The discreteness is inherited from the ansatz, not derived; it is the experimentally decisive content if the ansatz holds.

Why 43 specifically — the flux-bookkeeping argument. The whole conjecture reduces to one choice: which subset Σ ? The benchmark $\Sigma = N + L$ is not “add three nice integers”; it is a UV/IR sorting of the ledger, suggested by how these same integers function everywhere else in the series. The family-projector flux (N) and the gauge-rank flux (L) encode the $\text{SU}(5) \rightarrow \text{SM}$ breaking — *ultraviolet* data, fixed where the gauge symmetry is whole, i.e. *at the cap* (the UV terminus of the throat). The dyadic seam depth (J) is *infrared* data — it sets $\alpha^{-1} = 2^J$ and the lepton dyadics of GG-A9-SM *at the boundary*, so it terminates at the boundary, not at the cap. The cap therefore collects the two UV fluxes and not the IR one (Eq. (26.5)):

$$\Sigma = N + L \implies \Xi_{\text{cap}} = K + N + L = 35 + 3 + 5 = 43. \tag{B.12}$$

Read this way, 43 is the throat depth K , plus the two symmetry-breaking fluxes (N, L) that end at the cap, with the seam flux (J) *excluded* for a reason — which is exactly what an arbitrary “sum of three integers” would not have.

Where the softness genuinely remains. Three joints are heuristic, not proven, and they are the honest locus of the “arbitrary” feeling. (1) The selection $\Sigma = N + L$ rests on the UV/IR narrative; what would settle it is a defining computation for the cap loading, which has yet to be formulated on the admissible class (it carries no orbifold cap, and the manifold, class and integrand are not written down) — and until such a computation exists and returns $N + L$ (rather than, say, $N + J$, or something not of the form $K + \Sigma$) the assignment is a flux-bookkeeping ansatz, the single sharpest open item of the paper (§26.4). (2) “At most a pair of the three flux types” (which trims the eighth subset $N + L + J = 15$) is a stated codimension constraint on the cap, the least-defended

step. (3) The exponential-asymmetry step shows only that the eigenstate is not either face-value norm and contributes nothing to selecting the value — so the impressive e^{70} is a red herring for “why 43”; the value is pure integer counting.

Why this is a sharp prediction, not a hedge. Under the ansatz the uncertainty is a *list of seven frequencies*, disposable two independent ways. *By theory*: a defining computation for the cap loading, once formulated, returns Σ and collapses the window to a point, or retires the ladder. *By experiment*: since $\nu_K = 1.95 (\Xi_{\text{cap}}/43)$ GHz, the discovery frequency *measures* Ξ_{cap} *outright* — the line’s position within the window *is* the value of the integer (the rungs sit 45–136 MHz apart against a ± 25 MHz line, so every pair is resolved but 42/43, separated by 1.8σ ; §26.3). A tuned parameter hides forever inside its error bar; this integer is read off the dial the moment the bell is heard.

The picture in one line. Ξ_{cap} is the ratio f_{bulk}/f_K between the bulk normalization and the boundary eigenstate; the ansatz reads it in the form $K + \Sigma$ with Σ a subset-sum of $\{N, L, J\}$, so the mass is one of seven discrete frequencies about a pitch calibrated once; and $43 = K + N + L$ is the row picked out by the cap collecting the two ultraviolet symmetry-breaking fluxes while the infrared seam flux stays at the boundary — a motivated, falsifiable benchmark whose only genuinely open step is a single named cohomology integral.

B.5.4 Why is the Kaxion mass so small — and why can GG-6 predict it when the axion and the WIMP cannot?

Why is the Kaxion mass so tiny compared with a normal particle such as the electron? And why can GG-6 predict it so accurately, unlike the QCD/PQ axion or the WIMP?

Both the smallness and the predictability come from the same fact: the Kaxion is a shift-protected zero-mode whose one free scale GG-6 calibrates once, declares, and then holds fixed.

Why so small. Everything rests on the pseudo-Nambu–Goldstone relation

$$m_K f_K = \sqrt{\chi_{\text{QCD}}} \implies m_K = \frac{\sqrt{\chi_{\text{QCD}}}}{f_K}. \quad (\text{B.13})$$

The numerator is a *hadronic* scale, $\sqrt{\chi_{\text{QCD}}} = (0.0755 \text{ GeV})^2 = 5.7 \times 10^{-3} \text{ GeV}^2$; the denominator is a *geometric*, *GUT-adjacent* scale, $f_K \simeq 7 \times 10^{11} \text{ GeV}$. Their ratio is

$$m_K = \frac{5.7 \times 10^{-3} \text{ GeV}^2}{7 \times 10^{11} \text{ GeV}} \simeq 8 \times 10^{-15} \text{ GeV} = 8 \mu\text{eV}. \quad (\text{B.14})$$

Compare the electron: $m_e = y_e v / \sqrt{2}$ with $v = 246 \text{ GeV}$ and Yukawa $y_e \simeq 3 \times 10^{-6}$, giving 0.511 MeV, so $m_e/m_K \sim 6 \times 10^{10}$ — the Kaxion is eleven orders of magnitude lighter. The reason is structural, not a small coupling. The electron is a *wiggle*: a charged fermion permits a mass term $m_e \bar{e}e$ the moment electroweak symmetry breaks, and it costs weak-scale energy to make. The Kaxion is a *dial position* — a compact-angle zero-mode. Rotating the whole dial is a shift symmetry, so its mass is

forbidden to all orders in perturbation theory; it is exactly massless until topology breaks the shift symmetry, and the only thing that does is QCD instantons, “paying in exponentially suppressed coin” (§9.6). That feeble tilt, $\sqrt{\chi_{\text{QCD}}}$, is shared out over the huge field range f_K . The lightness is a *see-saw*: a small explicit-breaking scale over a large spontaneous-breaking scale, $m \sim \Lambda^2/f$; the bigger f_K , the lighter the boson.

Why so accurate — and why the axion and WIMP cannot. Write the mass as $m_K = \sqrt{\chi_{\text{QCD}}}/f_K$ with $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$, and look at where each factor comes from. The susceptibility $\sqrt{\chi_{\text{QCD}}}$ is lattice QCD ($\chi^{1/4} = 75.5(5)$ MeV) — the entire $\pm 1.3\%$ on the line at fixed f_K ; the $\chi(T)$, $g_*(T)$ transport enters the closure angle, not the mass. Of the geometric factors, $\Xi_{\text{cap}} = 43$ is a benchmark integer and carries *no continuous error* — its residual is the discrete ladder, not a percentage. The bulk scale f_{bulk} is different in kind and the difference should not be blurred: it is a continuous normalization, not an integer, and this paper does not compute it; it is calibrated through the abundance closure at the benchmark row (§36.2), so the benchmark’s absolute placement is a calibration, not a prediction. What is true, and is the point, is that no continuous quantity in the chain is left *free*: the one calibrated scale is declared as such, and everything after it is arithmetic. That is why the line is printed as $8.1 \mu\text{eV}$ to $\pm 1.3\%$ at the benchmark row, with the window as the geometry’s residual.

Now the contrast. For the *QCD/PQ axion* the relation $m_a = \sqrt{\chi}/f_a$ is equally precise — QCD is QCD for everyone — but f_a is a *free ultraviolet parameter*, unconstrained across the axion window 10^9 – 10^{17} GeV, so the mass is unknown by eight orders of magnitude: the theory predicts the *line*, not the *point*, which is exactly why the field has spent decades *scanning*. GG-6’s advance is to tie the one number the axion program leaves free, $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$, to a bulk scale calibrated once and an integer ladder; supply the calibration and the same precise QCD relation collapses the band to a window of seven rungs (§B.5.3) with one rung marked. The *WIMP* is worse off: it has no sharp mass at all. The “WIMP miracle” is a loose coincidence of scales — a weak-scale annihilation cross-section gives roughly the right relic density — fixing only a cross-section-times-mass combination and leaving the mass anywhere from GeV to ~ 100 TeV, while requiring new weak-scale physics (SUSY) the LHC has not found. So the three cases line up as: WIMP, a scale coincidence with no formula; QCD axion, an exact formula with one free scale (hence a band); Kaxion, the exact formula with that scale calibrated once (hence a window with a marked rung).

The picture in one line. Imagine a heavy flywheel (the six-dimensional geometry, scales 10^{12} – 10^{16} GeV) with a frictionless dial on its rim. Spinning the dial costs nothing — the shift symmetry — so its natural frequency would be zero; only a faint notch cut by QCD instantons gives it a restoring force, and because the notch is spread over the enormous circumference f_K , the pitch is a μeV whisper. A normal particle is a ripple *in* the flywheel and rings at the flywheel’s own huge frequency; the Kaxion is the dial’s rotation and rings eleven orders lower. And it is predictable because the pitch is (notch depth)/(circumference): QCD measures the notch to a percent, and GG-6 — unlike every prior axion theory — reads the circumference off one calibrated scale and an integer count of the throat, the ladder ansatz.

B.5.5 Why is the cosmic inventory locked to π , and what becomes of the coincidence problem?

Why are the cosmic fractions given so precisely — $\Omega_m = 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$? And “why now”: what of the cosmic coincidence problem?

By the paper’s own grading these are the *weakest* entries in the program — the “ π/e -fraction” rows, explicitly flagged as suggestive rather than decisive (Table 2), in contrast to the rigid integer locks like $\alpha^{-1} = 2^J$. Here is the real derivation, marked where it rests on a conjecture, and the coincidence stated honestly.

What is claimed, and how well it does. The present-epoch fractions are

$$\Omega_m = \frac{1}{\pi} = 0.3183, \quad \Omega_b = \frac{1}{2\pi^2} = 0.0507, \quad \Omega_{\text{DM}} = \frac{1}{\pi} - \frac{1}{2\pi^2} = 0.2676, \quad \Omega_\Lambda = 1 - \frac{1}{\pi} = 0.6817, \quad (\text{B.15})$$

against Planck 0.315, 0.0493, 0.265, 0.685 — every slice within $\sim 3\%$ (Eq. (15.1)). In GG-A10-Cosmos this is a conditional theorem, and the condition is a single named open one: the *closed-packet conjecture* that paper carries as its Conjecture VIII — the Hopf-fiber identification of its Lemma VIII, with T-Cosmos-17 the conditional theorem (C-Kaxion-c in this paper’s deck). The form is derived-conditional, not asserted, and not fitted — the integers and geometry were fixed upstream for reasons that never mention cosmology.

The derivation: one Hopf-fibered 3-sphere. The engine is OGGE/GDOS: the four-dimensional world is a lawful projection of the six-dimensional bulk, and the cosmic budget is how the projection’s measure divides. The relevant compact fiber is the Hopf fibration $S^1 \hookrightarrow S^3 \rightarrow S^2$, and the inventory is then volumes of the unit 3-sphere and its Hopf pieces. With the unit-sphere volumes $\text{Vol}(S^1) = 2\pi$, $\text{Vol}(S^2) = 4\pi$, $\text{Vol}(S^3) = 2\pi^2$, spread the total (critical) density uniformly over S^3 and assign matter to the fiber, dark energy to the base:

$$\begin{aligned} \Omega_m &= \frac{\text{Vol}(S^1)}{\text{Vol}(S^3)} = \frac{2\pi}{2\pi^2} = \frac{1}{\pi}, & \Omega_\Lambda &= 1 - \frac{1}{\pi}, \\ \Omega_b &= \frac{1}{\text{Vol}(S^3)} = \frac{1}{2\pi^2} \text{ (baryons = one arc unit)}, & \Omega_{\text{DM}} &= \frac{2\pi - 1}{2\pi^2} = \frac{1}{\pi} - \frac{1}{2\pi^2}. \end{aligned} \quad (\text{B.16})$$

The whole inventory falls out of one figure: matter is the Hopf fiber S^1 , a fraction $2\pi/2\pi^2 = 1/\pi$ of the 3-sphere; along that fiber circle baryons are one unit of arc and dark matter the remaining $2\pi - 1$. The two “flexible symbols” π and $2\pi^2$ are not chosen — they are a circle’s circumference and a 3-sphere’s volume, the only measures a Hopf projection makes available. The one genuinely conjectural step is the *identification* that the projection fiber is this Hopf S^3 , with matter→fiber, dark energy→base, baryon→one unit; settle its cohomology and the π -shape becomes a theorem.

Timeless versus snapshot. A fraction of the *critical* density is time-dependent: $\Omega_m(z)$ runs from near 1 in the past toward 0 in the future, so “ $\Omega_m = 1/\pi$ ” is a statement about *now*, not a constant of nature. The parts that hold at *every* epoch — because dark matter and baryons both

dilute as a^{-3} — are the ratios within matter,

$$\frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} \text{ (baryon share of matter),} \quad \frac{\rho_{\text{DM}}}{\rho_b} = 2\pi - 1 \simeq 5.28, \quad (\text{B.17})$$

both epoch-independent and both matching the sky (0.156 vs $1/2\pi = 0.159$; 5.37 vs $2\pi - 1 = 5.28$). These are the robust π -locks: on a circle of circumference 2π , baryons take one unit of arc and dark matter the remaining $2\pi - 1$. The absolute $\Omega_m = 1/\pi$ is the present-day split — which is exactly where the coincidence lives.

Why now — demoted, not solved. GG-A10-Cosmos does not make the coincidence vanish; it demotes it, by factoring “why now” into three questions. *Form*: why this shape? Forced by the geometry. This is the big win — because $\Omega_\Lambda = 1 - 1/\pi$ is a *residual share* of a fixed inventory, the notorious 10^{120} fine-tuning of the cosmological constant *dissolves*: one no longer tunes a tiny number, one reads a leftover. *Scale*: why is the absolute vacuum energy meV-sized? A see-saw residue, not a free input. *Epoch*: why do we observe *when* $\Omega_m \approx 1/\pi$? This is the only piece read from the sky; since $\Omega_m(t)$ crosses $1/\pi$ at a definite time, and the universe is 13.8 Gyr old today versus 13.8 Gyr at the crossing, “now” sits only ~ 75 Myr — a few percent — past it. So the honest scorecard: the 120-order fine-tuning of Λ is genuinely removed (the form is forced), and what remains is a mild, few-percent timing statement. The paper does not claim to close that last few percent; it shows it demoted from a fine-tuning to a coincidence of *when*.

The picture in one line. Picture the energy budget painted uniformly on a unit 3-sphere, the natural closed arena of the projection; that sphere is Hopf-fibered, a circle through every point organized over a 2-sphere. Dark energy is the base, matter is the fiber — a fraction $2\pi/2\pi^2 = 1/\pi$ of the whole — and along the fiber circle baryons are one tick of the clock and dark matter the remaining $2\pi - 1 \approx 5.3$, exactly the measured dark-to-ordinary ratio, holding at every epoch because it is a ratio of two things that dilute alike. Those shares of the sphere are fixed geometry; a *fraction of the critical density* is a moving snapshot; so the only accident left is the mild timing one — we read the dial a few percent past the moment matter’s slice equals $1/\pi$.

B.5.6 Why is the Kaxion cold, and why are there no galaxies or stars without it?

Why is the Kaxion cold (non-relativistic)? How can it contribute to the early formation of galaxies and stars? And is it true that without it no galaxies or stars form?

Coldness, the JWST head start, and the scaffold argument are three faces of one story.

Why the Kaxion is cold. Three facts stack up, and together they make coldness a theorem. (a) *It is never thermal.* Every coupling is divided by $f_K \sim 7 \times 10^{11}$ GeV, so the Kaxion never reaches thermal equilibrium; it is made by vacuum misalignment. (b) *Misalignment makes a zero-momentum condensate.* In the early universe the field is frozen by Hubble friction ($3H \gg m_K$) at $\theta_i \simeq 1.1$, spatially uniform on super-horizon scales; when the expansion slows to $3H \simeq m_K$

(onset at $T \simeq 1$ GeV, $t \simeq 0.3 \mu\text{s}$) it oscillates. A spatially uniform oscillating field is a condensate of zero-momentum quanta — intrinsically cold. *(c) Dust equation of state, exactly.* The Bessel envelope gives $a_K \propto t^{-3/4} \propto R^{-3/2}$, so

$$\rho_K = \frac{1}{2} m_K^2 a_K^2 \propto R^{-3}, \quad (\text{B.18})$$

the matter dilution law, exactly, as the asymptotics of a special function (§16). The instantaneous equation of state swings between $+1$ and -1 at $2\nu_K$, but the oscillation period today is $\sim 10^{-8}$ s, so cosmology sees only the average, $w = 0$ — pressureless dust. *(d) Velocities microscopic.* The velocity dispersion at equality is $\sim 4 \times 10^{-14}$; the free-streaming length is $\lambda_{\text{fs}} \sim 10^{-6}$ pc and the Jeans length $\lambda_J \sim 3 \times 10^{-4}$ pc — nine and six orders below the dwarf-galaxy scales where warm dark matter is constrained; and the de Broglie wavelength in a halo is $\lambda_{\text{dB}} \simeq 154$ m. It erases no structure. Its occupation number is $\mathcal{N} \sim 2 \times 10^{26}$: a magnificently classical field. It is not a boiling gas but a still pond — one coherent, universe-filling oscillation in which every quantum carries essentially zero momentum.

How it helps make the first galaxies and stars. Two levels. *Necessary:* coldness preserves all the small-scale power. Warm dark matter free-streams away small-scale structure; fuzzy dark matter erases it with quantum pressure. The Kaxion, with $\lambda_{\text{fs}} \sim 10^{-6}$ pc and $\lambda_{\text{dB}} \sim 154$ m, does neither, keeping the full ΛCDM spectrum. The JWST surprise is a demand for *more* early structure, a one-sided discriminant against the suppressing candidates and for cold; the Kaxion clears their exclusion bounds by sixteen orders of magnitude (§21). *Sufficient:* what survives of the GG-6 head start after the termination is one lever, not three — the collapse-threshold working value its cosmology companion carries for the baryonic environment ($\delta_c^{\text{eff}} \simeq 1.43\text{--}1.33$ in place of 1.686, a chronology shift $\Delta z \sim 2\text{--}3$ on the high- ν tail, graded qualitative there and quoted at model level here). The streaming-velocity damping and the early power boost that earlier readings drew from the gauged $U(1)_B$ sector are withdrawn with the termination: a Dirichlet-terminated vector that endpoint-localized quarks do not produce at tree level carries no drag on baryons (§21). (Honesty fence: the coldness is theorem-level; the surviving lever is GG-A10-Cosmos’s working value, not first-principles-exact; and if the JWST tension dissolves into baryonic systematics, the Kaxion simply predicts standard cold ΛCDM early structure — still on the favorable side.)

Without the Kaxion, no galaxies or stars? It is the cold-dark-matter *scaffold* that is indispensable, and within GG-6 the Kaxion is the only thing allowed to be it. Before recombination ($z \simeq 1100$) baryons are Thomson-locked to photons; radiation pressure forbids their collapse, so baryonic perturbations oscillate (the acoustic peaks) rather than grow, and small scales are Silk-damped. Dark matter, neutral and decoupled from photons, begins growing at matter–radiation equality ($z \simeq 3400$), unimpeded. Perturbations grow linearly, $\delta \propto a$, and the CMB fixes $\delta \sim 10^{-5}$ at recombination. *Baryons alone*, starting only at recombination, would reach today

$$\delta_{\text{today}} \sim 10^{-5} \times (1 + z_{\text{rec}}) \sim 10^{-5} \times 1100 \sim 10^{-2} \ll 1 \quad (\text{B.19})$$

— still linear: no collapse, no galaxies, ever. With cold dark matter the potential wells are already deep at recombination (they grew, undamped, since equality); baryons fall into the pre-carved wells and the combined perturbation grows to $\delta \sim 1$ — nonlinear collapse — in time to make galaxies at the observed epochs. So the observed early, abundant structure requires a cold, neutral, pressureless component collapsing *before* recombination: dark matter. The honest statement has two tiers: without cold dark matter, no galaxies or stars (textbook); and *within* GG-6, which forbids the WIMP, the sterile neutrino, the dark photon, and primordial black holes and forces the Kaxion as the unique admissible cold relic, no Kaxion means no scaffold, no early galaxies, no stars, no carbon and oxygen, no us. The first tier is secure; the second is the falsifiable GG-6 claim — which is why a clean null in the Kaxion window retires the calibrated placement and its ladder, not the framework: the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinned branch, or the relic is not this sector’s, and the second tier is tested again there.

The picture in one line. Dark matter is the riverbed and baryons are the water: water cannot pool into lakes until the beds are carved, and radiation pressure keeps the baryonic water sheeting flat until recombination; the dark matter, immune to that pressure, carves the wells early, and after recombination the gas rushes in and lights the first stars. Take the riverbeds away and the water spreads thin and never gathers into a galaxy in the age of the universe. The Kaxion is GG-6’s name for the riverbed — a still, cold, coherent condensate that did its gravitational work in silence from the first microsecond, and is only now, thirteen billion years later, quiet enough and near enough to be heard at 1.95 GHz.

C From-scratch derivation courses

This is the machine room. Each course below is self-contained and worked at graduate-course pace: it states the setup, carries the algebra without skipped steps, boxes the result, and inserts the Kaxion’s numbers only at the end, so that a reader can reconstruct each headline quantity independently. The courses proceed from the geometry (how a tensor field becomes the Kaxion, where the throat normalization Ξ_{cap} comes from, why the attractor is stable) through the three pieces of physics every axion dark-matter paper rests on (the QCD mass and photon coupling, the misalignment abundance) and the two entries unique to the Kaxion (the baryonic ledger entry and the cap index Ξ_{cap}), closing with the inherited π -locked inventory.

C.1 Dualization and periodicity

Why a two-form is a scalar in four dimensions. The structural fact underneath the whole construction is elementary field theory and worth stating cleanly for the newcomer. In d spacetime dimensions a p -form gauge potential with field strength H_{p+1} is physically equivalent (“Hodge-dual”) to a $(d - p - 2)$ -form: the equation of motion $d \star H_{p+1} = 0$ and the Bianchi identity $dH_{p+1} = 0$ are the same pair of equations one would write for the dual field strength $\tilde{H}_{d-p-1} = \star H_{p+1}$, with the roles of dynamics and identity exchanged. In four dimensions a two-form potential B_2 has field

strength H_3 , and its dual is a $(4 - 2 - 2) = 0$ -form — a single scalar. Counting degrees of freedom confirms it: a massless two-form in 4D propagates $\binom{d-2}{2} = \binom{2}{2} = 1$ polarization, exactly one real scalar. That scalar is a *pseudoscalar* (parity-odd) because $\star$ carries one factor of the Levi-Civita tensor, and it is periodic because the flux of H_3 is quantized (below). This is the model-independent reason every theory with a Green–Schwarz two-form contains a light pseudoscalar “string axion”; the Kaxion is that pseudoscalar for GG-6.

Reduction. The bulk tensor kinetic term reduces on the canonical metric (6.2) as $\int_{\mathcal{M}_6} \sqrt{-G} |H_3|^2 \rightarrow \int_{\mathcal{M}_4} \sqrt{-g} Z |H_3^{(4)}|^2$ with $Z = \int_0^K dr \int_0^{2\pi} d\theta e^{-2r} |\psi_0(r)|^2$ the zero-mode normalization; Z defines f_{GS}^2 and is finite and nonzero on the rigid geometry (no modulus can send it anywhere). The boundary dualization is Eq. (9.1) of the main text.

Periodicity. The differential-cohomology refinement of GG-A3-B16 quantizes $\oint_{\Sigma_3} H_3 \in 2\pi\mathbb{Z}$ over closed three-cycles. Under (9.2) a large dual transformation winds a_{GS} by $2\pi f_{\text{GS}}$, identifying the target as a circle: the tensor shadow is an angle, as used throughout. The Wilson echo’s periodicity is the large-gauge invariance of (6.4). Both periodicities are gauge facts, which is the structural root of the quality-problem immunity (Section 24).

The reduction in components, and the hierarchy of the two decay constants. The zero-mode normalizations can be exhibited explicitly on the canonical metric, and the computation — three lines per field — explains structurally why the tensor’s dual and a circle holonomy carry hierarchically different normalizations. For the tensor: with $G_{\mu\nu} = e^{-2r} \eta_{\mu\nu}$, $G_{rr} = G_{\theta\theta} = 1$, the volume factor is $\sqrt{-G} = e^{-4r}$, while raising the three boundary indices of $H_{\mu\nu\rho}$ costs e^{+6r} ; the boundary-directed kinetic term therefore reduces with the weight

$$\int_0^K dr \int_0^{2\pi} d\theta e^{-4r} e^{+6r} = 2\pi \int_0^K e^{+2r} dr = \pi(e^{2K} - 1) \simeq \pi e^{2K}, \quad (\text{C.1})$$

an *infrared-dominated, exponentially enhanced* norm: $f_{\text{GS}}^2 \propto e^{2K}$ in bulk units. For the Wilson line: the relevant component is $F_{\mu\theta}$, whose kinetic term carries $\sqrt{-G} G^{\mu\nu} G^{\theta\theta} = e^{-4r} e^{+2r} = e^{-2r}$, reducing with

$$2\pi \int_0^K e^{-2r} dr = \pi(1 - e^{-2K}) \simeq \pi, \quad (\text{C.2})$$

an *ultraviolet-dominated, order-one* norm: $f_\theta \sim M_6$ with no enhancement. The structural conclusion is twofold and is used in the main text without its coefficients: (i) which decay constant the canonical remnant carries depends on the dualization convention: for the 2π -periodic scalar $B_{r\theta}$ the decay constant scales as $\sqrt{Z}$, for the dual of a two-form with unit-normalized large gauge transformations as $1/\sqrt{Z}$, and the anomaly companion’s O14 records that the infrared-weighted chirality read in the second convention lands near the TeV scale, which is excluded, while the first convention with a Planckian cell lands super-Planckian. Neither convention reproduces the calibrated $f_K = 7.1 \times 10^{11}$ GeV; the ordering $f_K \gg M_{\text{lock}}$ is therefore not something this appendix derives. (ii) The *ratio* of the two norms is exponentially sensitive to the region at $r = K$ where (C.1) accumulates, which is why the absolute normalization is the hard, throat-sensitive quantity

that O14 names, and why this paper *calibrates* f_K through the abundance closure at a benchmark loading rather than attempting the normalization directly. The residual of that reading is discrete (the Ξ_{cap} ladder) on top of an undeclared-by-computation absolute scale; the $\pm 1.3\%$ on the mass at fixed f_K is the Standard-Model input (χ_{QCD}). The geometry shouts the ordering of the two norms and is silent, so far, on the coefficient.

C.2 The GFF Hessian at the compact-angle attractor

On the stratum (a_K, a_H) the relevant potential is $V = \chi_{\text{QCD}}[1 - \cos(a_K/f_K)] + \frac{1}{2}M_{\text{cap}}^2 a_H^2 + \lambda a_K a_H$, with M_{cap} the throat-scale cap mass and λ a residual mass mixing, kept general here although on the intended class no scalar–scalar bilinear appears at all (§9.3), so that $\lambda = 0$ and what follows is an upper bound on the effect of any residual mixing. At the attractor $a_K = a_H = 0$:

$$\mathcal{H} = \begin{pmatrix} \chi_{\text{QCD}}/f_K^2 & \lambda \\ \lambda & M_{\text{cap}}^2 \end{pmatrix}, \quad \det \mathcal{H} = m_K^2 M_{\text{cap}}^2 \left(1 - \frac{\lambda^2}{m_K^2 M_{\text{cap}}^2}\right) > 0 \quad (\text{C.3})$$

since $\lambda = \epsilon m_K^2$ for an admixture $\epsilon < 1$ of the heavy direction, so that $\lambda \ll m_K M_{\text{cap}}$ by the full hierarchy m_K/M_{cap} : both eigenvalues are positive, the light one equal to m_K^2 up to a fractional correction $\epsilon^2 m_K^2/M_{\text{cap}}^2$, negligible for any ϵ , and exactly on the intended class ($\lambda = 0$). The attractor is a true GA (Morse index 0), and the eigenvalue statement is Eq. (10.1) of the main text with its accuracy now quantified.

The eigenvalues, step by step. For a reader meeting two-state mixing for the first time, the diagonalization is worth doing explicitly, because it is the same 2×2 algebra that governs neutrino mixing, coupled pendula, and avoided level crossings, and the lesson transfers. Write the diagonal entries as $A = \chi_{\text{QCD}}/f_K^2$ (the light, QCD-set curvature) and $B = M_{\text{cap}}^2$ (the heavy, throat-set curvature), with off-diagonal λ . The eigenvalues of $\begin{pmatrix} A & \lambda \\ \lambda & B \end{pmatrix}$ are

$$\mu_{\pm} = \frac{A+B}{2} \pm \sqrt{\left(\frac{B-A}{2}\right)^2 + \lambda^2}. \quad (\text{C.4})$$

The physics lives in the hierarchy $A \ll B$ (the light mass-squared is $\sim (8.1 \mu\text{eV})^2$, the lifted direction’s mass-squared is set at the scale the holonomies are lifted, $\gtrsim (2 \times 10^{15} \text{ GeV})^2$ on the coset zero-mode reading of GG-A9-SM’s PS-SM-22, vastly larger) and in the smallness of the mass mixing: a term $\lambda a_K a_H$ can arise only from expanding the QCD cosine in an admixture ϵ of the heavy direction, $V = \chi[1 - \cos((a_K + \epsilon a_H)/f)] + \frac{1}{2}M_{\text{cap}}^2 a_H^2$, giving $\lambda = \epsilon \chi/f^2 = \epsilon m_K^2$. Expanding (C.4) to second order in the small ratio $\lambda/(B-A) \simeq \lambda/B$,

$$\mu_- = A - \frac{\lambda^2}{B-A} + \dots = \frac{\chi_{\text{QCD}}}{f_K^2} \left(1 - \frac{\lambda^2}{m_K^2 M_{\text{cap}}^2} + \dots\right), \quad \mu_+ = B + \frac{\lambda^2}{B-A} + \dots \quad (\text{C.5})$$

The light eigenvalue μ_- is the physical Kaxion mass-squared; the fractional shift away from the bare QCD value is $\lambda^2/(m_K^2 M_{\text{cap}}^2) = \epsilon^2 m_K^2/M_{\text{cap}}^2$, which for any admixture $\epsilon < 1$ — even the recorded single-transit overlap $\epsilon = 1/K$ of Appendix F, the most generous value ever entertained — is far

below the lattice uncertainty on χ_{QCD} itself, so to the accuracy of this paper $m_K^2 = \chi_{\text{QCD}}/f_K^2$ exactly — and exactly without qualification on the intended class, where $\lambda = 0$. This is the content of Eq. (10.1): the heavy cap mode decouples, and the Kaxion inherits its curvature from QCD alone.

What “Morse index 0” means, and why it matters. A critical point of a smooth potential is classified by the number of *negative* eigenvalues of its Hessian — its Morse index: index 0 is a local minimum (all curvatures positive, a stable valley floor), index equal to the dimension is a maximum, and anything in between is a saddle. Because both $\mu_{\pm} > 0$ in (C.5), the compact-angle critical point is a genuine minimum, index 0 — a stable resting place for the field, not a hilltop it would roll off. In the GRIP language of Part II this is precisely the statement that the point is a Geometric Attractor (GA): the runtime’s descent of F_{GFF} ends here and stays. The same Morse-theoretic bookkeeping, applied across the whole Wilson-line stratum, is what singles out the *deepest* such minimum as the Ground Geometric Attractor (GGA, A11-Origin), the unique global vacuum the Kaxion condensate actually occupies.

C.3 The QCD-coupled axion: the mass relation and the photon coupling, derived

This appendix derives the two relations that turn a decay constant into an observable mass and an observable photon coupling. They are common to every QCD axion; the Kaxion differs only in that its decay constant is fixed by cosmology (Appendix C.6) and its anomaly ratio $E/N = 8/3$ is a theorem of the SU(5) ledger (GG-A9-SM) rather than a model choice. A reader who follows the four steps below can reconstruct the Kaxion mass and coupling from a single QCD number.

Step 1 — the gluon portal generates a periodic potential. The defining property of any axion is its coupling to the QCD topological density,

$$\mathcal{L} \supset \frac{a}{f_a} \frac{\alpha_s}{8\pi} G_{\mu\nu}^b \tilde{G}^{b\mu\nu}, \quad \tilde{G}^{b\mu\nu} = \frac{1}{2} \epsilon^{\mu\nu\rho\sigma} G_{\rho\sigma}^b, \quad (\text{C.6})$$

which is precisely the operator multiplying the strong-CP angle $\bar{\theta}$, so the axion field enters physics only through the combination $\bar{\theta} + a/f_a$. Nonperturbative QCD — instantons at weak coupling, the lattice in general — makes the vacuum energy a periodic function of this combination,

$$V(a) = \chi_{\text{QCD}} \left[1 - \cos\left(\frac{a}{f_a}\right) \right] + \mathcal{O}(\cos 2(a/f_a)), \quad (\text{C.7})$$

where the curvature scale χ_{QCD} is the QCD *topological susceptibility*, a pure-glue-and-light-quark quantity with no reference to the axion. The Vafa–Witten theorem (Vafa & Witten, 1984) — invoked in its licensed, vector-like domain — certifies that the minimum of (C.7) sits at $a = 0$, i.e. at $\bar{\theta}_{\text{eff}} = 0$; the potential itself is exhibited directly by lattice computation, so the relaxation needs no theorem beyond its existence: relaxation to the minimum *is* the solution of the strong-CP problem.

Step 2 — the mass. The mass is the curvature of the potential at its minimum,

$$m_K^2 = \left. \frac{d^2 V}{da^2} \right|_{a=0} = \frac{\chi_{\text{QCD}}}{f_K^2} \implies \boxed{m_K f_K = \sqrt{\chi_{\text{QCD}}}}, \quad (\text{C.8})$$

the single most important relation in axion physics: the product of mass and decay constant is a fixed QCD number, so a small mass demands a large decay constant and vice versa. At zero temperature the chiral Lagrangian fixes the susceptibility in terms of measured hadronic quantities,

$$\chi_{\text{QCD}} = \frac{m_u m_d}{(m_u + m_d)^2} m_\pi^2 f_\pi^2 = (75.5 \pm 0.5 \text{ MeV})^4, \quad (\text{C.9})$$

with $f_\pi \simeq 92 \text{ MeV}$ the pion decay constant and $m_\pi \simeq 135 \text{ MeV}$; the lattice (Borsanyi et al., 2016) confirms this to percent accuracy. Inserting (C.9) into (C.8) gives the convenient form (Grilli di Cortona et al., 2016)

$$m_K = 5.70(7) \mu\text{eV} \times \left(\frac{10^{12} \text{ GeV}}{f_K} \right), \quad (\text{C.10})$$

which, once f_K is known, fixes the mass and hence the laboratory frequency $\nu = m_K/2\pi\hbar$.

Step 3 — the photon coupling and the model-independent -1.92 . The axion-photon coupling has two physically distinct contributions,

$$g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K} \left(\frac{E}{N} - \frac{2}{3} \frac{4+z}{1+z} \right), \quad z \equiv \frac{m_u}{m_d} \simeq 0.462, \quad (\text{C.11})$$

and the parenthesis is worth dissecting because students often see only its numerical value. The first term, E/N , is the *ultraviolet* anomaly: the ratio of electromagnetic to color anomaly coefficients of the heavy fermions that carry the axion's charge — a short-distance, model-dependent integer ratio. The second term is *infrared* and model-independent: because the operator (C.6) mixes the axion with the neutral pion and the η , removing the $aG\tilde{G}$ term by a chiral rotation necessarily shifts the axion's coupling to two photons by a hadronic amount, evaluated from χPT as $\frac{2}{3}(4+z)/(1+z) = 1.92(4)$. Every axion, regardless of its ultraviolet completion, carries this -1.92 .

Step 4 — DFSZ, KSVZ, and where the Kaxion lands. The two canonical benchmarks differ only in E/N : the KSVZ (heavy-quark) model has $E/N = 0$, giving a parenthesis of -1.92 ; the DFSZ (two-Higgs) model has $E/N = 8/3$, giving $8/3 - 1.92 = 0.75$. In GG-Theory the value $E/N = 8/3$ is not chosen but *derived* — it is fixed by the same $\text{SU}(5)$ hypercharge normalization that gives $\sin^2 \theta_W = 25/108$ (GG-A9-SM) — so the Kaxion sits *exactly on the DFSZ line*. Numerically, with the closed f_K of Appendix C.6,

$$g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K} (0.75) \implies 1.2 \times 10^{-15} \text{ GeV}^{-1}, \quad (\text{C.12})$$

evaluated at the physical $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$ of the laboratory eigenstate (the bulk normalization scale is a bookkeeping intermediate of the diagonalization, not the decay constant of a propagating field

— Part V). The practical upshot for an experimentalist: the Kaxion’s depth target is the DFSZ line, the community’s hardest sensitivity goal, entered at the named frequencies.

C.4 The baryonic ledger entry, derived — and the withdrawn 5:2 readout

[Steps 1–3 below — baryon number gauged, the anomaly ledger, the 2/5 trace ratio — derive the Kaxion’s baryonic entry: $-2/5$ of X_4 under the anomaly companion’s normalization. The inversion of earlier revisions into a light-channel coupling ratio, and the 25/4 power ratio built on it, are withdrawn (Part VI) and recorded in Appendix F.]

The six-question essay (Appendix B.5.2) gives the physical story; this course carries the group theory and the dualization in full, so that a student can reconstruct the ledger entry from the gauge content alone. The background on anomalies is standard (Peskin & Schroeder, 1995; Weinberg, 1996); the Green–Schwarz mechanism is developed at book length in Nakahara (2003).

Step 1 — baryon number is gauged. The bulk gauge group forced by the constitution is

$$G_{\text{Bulk}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}, \quad (\text{C.13})$$

and the entry that has no analog in a Peccei–Quinn axion theory is $\text{U}(1)_B$: baryon number is a *local* symmetry, not the accidental global one of the Standard Model. This is not an aesthetic choice — it is forced by the postulate that forbids accidental global symmetries together with the observed proton stability $\tau_p > 10^{34}$ yr. Gauging B annihilates every dimension-six proton-decay operator (they are not $\text{U}(1)_B$ -invariant), curing the disease that kills ordinary $\text{SU}(5)$, whose X, Y leptoquarks drive $p \rightarrow e^+ \pi^0$. A gauged $\text{U}(1)_B$ requires its own gauge boson Z_B and, crucially, its own anomaly bookkeeping.

Step 2 — the Kaxion couples to whatever the two-form cancels. In six dimensions the Green–Schwarz two-form cancels the reducible gauge and gravitational anomalies by coupling to the full anomaly four-form,

$$S_{\text{GS}} = \int_{\mathcal{M}_6} B_2 \wedge X_4, \quad X_4 = \sum_a c_a \frac{1}{8\pi^2} \text{Tr}(F_a \wedge F_a) + c_{\text{grav}} p_1, \quad (\text{C.14})$$

where the sum runs over the gauge factors of (C.13), p_1 is the first Pontryagin form, and the coefficient vector c_a is *fixed* by anomaly factorization — it is not adjustable. For GG-6 that vector is

$$(c_{\text{grav}}, c_{\text{SU}(5)}, c_{\text{U}(1)_B}, c_{Qi}) = \left(1, -1, -\frac{2}{5}, -\frac{2}{5}\right). \quad (\text{C.15})$$

Because the Kaxion is the four-dimensional dual of B_2 (Appendix C.1), it inherits a coupling to *every* factor B_2 talks to. The color entry gives the QCD portal $a_K G \tilde{G}$ that fixes the mass; the $\text{U}(1)_B$ entry gives the baryonic ledger entry

$$\mathcal{L} \supset \frac{c_B}{4} a_K F_B \tilde{F}_B, \quad (\text{C.16})$$

with c_B the ledger-level coefficient of the $U(1)_B$ entry, $-\frac{2}{5}$ of the hypercharge entry's (Step 3; the retired readout g_{KBB} of the withdrawn doublet configuration is not this coefficient, Appendix F) — a term which a standard axion simply cannot have, because in a theory where B is global there is no F_B for the axion to be anomalous with.

Step 3 — the trace-norm ratio, in full. The magnitudes $2/5$ in (C.15) are the ratio of the baryon and hypercharge anomaly coefficients, and they are pure $SU(5)$ representation theory. In the fundamental **5** the two Abelian generators, normalized so that hypercharge reproduces the Standard-Model assignments, are

$$Y_0 = \text{diag}\left(-\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}, \frac{1}{2}, \frac{1}{2}\right), \quad B_0 = \text{diag}\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0\right), \quad (\text{C.17})$$

where the first three entries are the color (quark) directions and the last two the weak-doublet (lepton) directions; baryon number lives only on the quarks, hypercharge on all five. The anomaly coefficient of an Abelian factor is the trace of the square of its generator over the chiral fermions; for one $\bar{\mathbf{5}} \oplus \mathbf{10}$ generation the **5** trace carries the ratio, and

$$\frac{\text{Tr}(B_0^2)}{\text{Tr}(Y_0^2)} = \frac{3 \times (1/3)^2}{3 \times (1/3)^2 + 2 \times (1/2)^2} = \frac{1/3}{1/3 + 1/2} = \frac{1/3}{5/6} = \frac{2}{5} = \frac{L - N}{L}, \quad (\text{C.18})$$

with the OGN-5 ledger integers $L = 5$ (the rank data of $SU(5)$) and $N = 3$ (color), so that $L - N = 2$. This is the origin of the $2/5$ in the coefficient vector: it is not fitted, it is $\text{Tr}(B_0^2)/\text{Tr}(Y_0^2)$ evaluated on one generation. The same integer pair $(L, L - N) = (5, 2)$ is what GG-A7-Qi4 locks into the four-current identity, which is why the identical ratio surfaces elsewhere in the cohort in a biological costume — one number, two theaters.

What is theorem and what is fenced. The ledger entry $-2/5$ is exact at the gauge/anomaly level: it follows from (C.15) and (C.18) with no continuous input. Dualizing B_2 to a_K preserves the coefficient vector — there is no “reciprocal” of a coupling ratio under a p -form/dual swap, which inverts a kinetic coefficient and nothing else — so the entry is $2/5$ of hypercharge's for source and response alike; the 5:2 of earlier revisions was an error, recorded in Appendix F. Two joints stay flagged. The coefficient vector (C.15) is read from the BI-6 factorization of Arisaka (2026, A3-BI6); and converting any gauge-level entry into an actual nucleon-*spin* reading would require the chiral matrix elements of Part VIII's KSVZ-class sheet. The number the geometry hands you is the ledger entry; the 25/4 once built on it rested on the inversion error and on a light vector the termination does not supply (Part VI).

C.5 The cap index $\Xi_{\text{cap}} = K + N + L = 43$: the throat normalization and its integer

This course makes explicit what the six-question answer (Appendix B.5.3) sketches: why the single conjecture-fenced number of the mass derivation is nonetheless *discrete*, and why the benchmark

value is 43. The relevant mathematics — flux quantization over cycles, wavefunction normalization on a warped throat — is standard differential geometry (Nakahara, 2003).

Step 1 — what Ξ_{cap} is: a wavefunction-renormalization integer. Ξ_{cap} is not a new parameter; it relates two normalizations of the *same* field,

$$f_K = \frac{f_{\text{bulk}}}{\Xi_{\text{cap}}}, \quad m_K = \frac{\sqrt{\chi_{\text{QCD}}}}{f_K} = \frac{\sqrt{\chi_{\text{QCD}}}}{f_{\text{bulk}}} \Xi_{\text{cap}}, \quad (\text{C.19})$$

where f_{bulk} is the decay constant of the Green–Schwarz direction deep in the bulk and f_K is that of the physical eigenstate a haloscope couples to. Because $m_K \propto \Xi_{\text{cap}}$, this single integer sets the *absolute* mass and nothing else: it moves *where* the line sits but never the coupling sheet: the DFSZ depth and the ledger entries are scale-independent theorems (Appendices C.3 and C.4); only the address carries the conjecture.

Step 2 — why the map is enormous: the throat sets the normalization. On the canonical warped metric $ds_6^2 = e^{-2r} \eta_{\mu\nu} dx^\mu dx^\nu + dr^2 + d\theta^2$ with $r \in [0, K]$, the tensor’s dual and a circle holonomy acquire wildly asymmetric kinetic normalizations (Appendix C.1),

$$Z_{\text{GS}} = 2\pi \int_0^K e^{2r} dr = \pi(e^{2K} - 1) \simeq \pi e^{2K}, \quad Z_\theta = 2\pi \int_0^K e^{-2r} dr \simeq \pi, \quad (\text{C.20})$$

differing by $e^{2K} = e^{70} \sim 10^{30}$. Canonical normalization divides each field by $\sqrt{Z}$, so the map from the bulk variable to the boundary eigenstate is exponentially large. Ξ_{cap} is that map expressed in the throat’s own units: the Green–Schwarz variable is normalized after the warp has run its full K e-folds, the boundary eigenstate sits at the throat’s mouth, and the ansatz reads the conversion between them as the throat depth in integer units, plus whatever the cap adds — an e-fold count, not a flux count. The exponential e^{70} says only that the boundary eigenstate is not either face-value bulk norm — it proves nothing about Ξ_{cap} and contributes *nothing* to the integer value; the value is the ansatz’s counting, which is the point students most often miss.

Step 3 — the additive integer form, and its grade. One fact is derived from the metric, and two are the ansatz of Part V (Ansatz 26.2):

1. $\Xi_{\text{cap}} \neq 1$: the two face-value chiral norms (C.20) differ by some thirty orders and neither reproduces the calibrated f_K (O14 open). (This says the ratio is not unity; it does not set the value, and the exponential proves nothing about Ξ_{cap} .)
2. $\Xi_{\text{cap}} = K + \Sigma$, an *additive* integer form — the ansatz. The throat depth $K = 35$ is an e-fold count of the warp, not a flux count, and the ultraviolet terminus of the interval is *taken* to contribute an integer Σ ; the admissible class carries no orbifold cap, and no quantized flux equal to K is exhibited. This is the load-bearing step and it is an ansatz: it makes the residual an *integer*, not a continuous knob, if it holds.

3. Σ is a subset-sum of the remaining ledger integers $\{N, L, J\} = \{3, 5, 7\}$, and a codimension count limits the cap to hosting at most a *pair* of the three flux types. Hence

$$\begin{aligned} \Sigma \in \{0, 3, 5, 7, 8, 10, 12\} &\Rightarrow \Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\} \\ &\Rightarrow \nu_K = 1.95 (\Xi_{\text{cap}}/43) \in [1.59, 2.13] \text{ GHz}. \end{aligned} \quad (\text{C.21})$$

So, under the ansatz, the absolute mass is not a continuous parameter with an error bar — it is one of *seven discrete frequencies* (Table 11), the benchmark row calibrated and the others ratios to it. The discreteness is inherited from the ansatz, not derived.

Step 4 — why 43: the ultraviolet/infrared sorting of the ledger. The whole conjecture reduces to one choice — which subset Σ ? The benchmark $\Sigma = N + L$ is not “add two nice integers”; it is a UV/IR sorting suggested by how these same integers function everywhere else in the series. The family-projector flux N and the gauge-rank flux L encode the $\text{SU}(5) \rightarrow \text{SM}$ breaking — *ultraviolet* data, fixed where the gauge symmetry is whole, i.e. *at the cap*, the UV terminus of the throat. The dyadic seam depth J is *infrared* data: it sets $\alpha^{-1} = 2^J$ and the lepton dyadics of Arisaka (2026, A9-SM) *at the boundary*, so it terminates at the boundary, not at the cap. The cap therefore collects the two UV fluxes and excludes the IR one:

$$\Sigma = N + L \quad \Longrightarrow \quad \boxed{\Xi_{\text{cap}} = K + N + L = 35 + 3 + 5 = 43}. \quad (\text{C.22})$$

Read this way, 43 is the throat depth plus the two symmetry-breaking fluxes that end at the cap, with the seam flux J excluded for a stated reason — exactly what an arbitrary sum of three integers would not have.

Where the softness genuinely remains, and how it is disposed of. Three joints are heuristic, not proven, and they are the honest locus of the “arbitrary” feeling: the selection $\Sigma = N + L$ rests on the UV/IR narrative rather than on the integral; the “at most a pair” codimension count trims the eighth subset $N + L + J = 15$; and the exponential-asymmetry step contributes nothing to the value. What would settle all three at once is a defining computation for the cap loading, which has yet to be formulated on the admissible class of Arisaka (2026, A3-BI6) (it carries no orbifold cap) — and until such a computation exists and returns $N + L$ the assignment is a flux-bookkeeping ansatz, the single sharpest open item of the paper (its formulation is entry 1 of the roadmap, Appendix E.5). The uncertainty is disposable two independent ways: *by theory*, the defining computation collapses the window to a point or retires the ladder; *by experiment*, since $\nu_K = 1.95 (\Xi_{\text{cap}}/43)$ GHz, the discovery frequency *measures* Ξ_{cap} *outright* — the line’s position within the window is the value of the integer. A tuned parameter hides forever inside its error bar; this one is read off the dial the moment the bell is heard.

C.6 Misalignment production: the relic abundance from the equation of motion

This is the derivation that makes a feebly coupled, never-thermalized field into all of the cold dark matter. It is the classic “vacuum misalignment” mechanism (Abbott & Sikivie, 1983; Dine &

Fischler, 1983; Preskill, Wise, & Wilczek, 1983), presented here at tutorial pace; the modern review is Marsh (2016). The output is the abundance–decay-constant law that, set equal to the derived Ω_{DM} , closes f_K in Part V.

Step 1 — the field equation in an expanding universe. The spatially homogeneous zero mode, written as an angle $\theta(t) = a(t)/f_K$, obeys the Klein–Gordon equation in a Friedmann–Robertson–Walker background,

$$\ddot{\theta} + 3H\dot{\theta} + m^2(T)\theta = 0, \quad (\text{C.23})$$

where $H = \dot{R}/R$ is the Hubble rate (R the scale factor) and the small-angle form $V'(a) \simeq m^2 f_K^2 \theta$ has been used. Equation (C.23) is a damped oscillator with a *time-dependent* restoring frequency $m(T)$ and a Hubble “friction” $3H$.

Step 2 — frozen at early times. While $H \gg m$ the friction term dominates and the field is overdamped: θ is frozen at its initial value θ_i , the $\mathcal{O}(1)$ misalignment angle ($\theta_i \simeq 1.1$ rad, read from the abundance at the calibrated benchmark). Its energy density $\rho = \frac{1}{2}f_K^2\dot{\theta}^2 + V \simeq \frac{1}{2}m^2 f_K^2 \theta_i^2$ is nearly constant — the field behaves momentarily like a small vacuum energy, contributing nothing to the matter budget yet.

Step 3 — oscillation onset. As the universe cools, H falls (in the radiation era $H = \sqrt{\pi^2 g_*/90} T^2/M_{\text{Pl}}$) while $m(T)$ rises. Oscillation begins at the temperature T_{osc} where friction can no longer hold the field,

$$3H(T_{\text{osc}}) \simeq m(T_{\text{osc}}). \quad (\text{C.24})$$

After this moment $H \ll m$ and the field rolls down and oscillates about the minimum.

Step 4 — the adiabatic invariant: misalignment energy is cold. Once $m \gg H$ the mass changes slowly over one oscillation period, so the oscillator is *adiabatic* and its action — equivalently the comoving particle number $N_a = n_a R^3$ with $n_a = \rho_a/m$ — is conserved. Hence the number density redshifts as $n_a \propto R^{-3}$ and the energy density as $\rho_a = m n_a \propto R^{-3}$ (up to the slow growth of m until it saturates). *This is the crux:* a coherently oscillating scalar dilutes exactly like pressureless cold matter, $w = \langle p \rangle / \langle \rho \rangle = 0$, with no thermal velocities. The Kaxion is born cold, not cooled.

Step 5 — the yield. The number density at onset is $n_a(T_{\text{osc}}) = \frac{1}{2}m(T_{\text{osc}})f_K^2\theta_i^2$ (times an $\mathcal{O}(1)$ anharmonic correction for large θ_i). The quantity conserved thereafter is the *yield* $Y \equiv n_a/s$, with the entropy density $s = \frac{2\pi^2}{45} g_{*s} T^3$. The present dark-matter density is then

$$\Omega_K h^2 = \frac{m_K n_{a,0}}{\rho_c/h^2} = \frac{m_K Y s_0}{\rho_c/h^2}, \quad (\text{C.25})$$

with $s_0 = 2891 \text{ cm}^{-3}$ and $\rho_c/h^2 = 1.05 \times 10^{-5} \text{ GeV cm}^{-3}$ today.

Step 6 — the temperature dependence and the $f_K^{1.165}$ law. The only model input is how the mass turns on. The lattice gives a topological susceptibility falling as a power of temperature

above the QCD crossover, $\chi(T) \propto T^{-n}$ with $n \simeq 8.16$ (Borsanyi et al., 2016; Gorghetto & Villadoro, 2019), so $m(T) = m_K (T_{\text{QCD}}/T)^{n/2}$. Solving the onset condition (C.24),

$$T_{\text{osc}}^{2+n/2} \propto m_K T_{\text{QCD}}^{n/2} M_{\text{Pl}} \implies T_{\text{osc}} \propto m_K^{2/(4+n)} \propto f_K^{-2/(4+n)}, \quad (\text{C.26})$$

using $m_K \propto 1/f_K$ from (C.8). Carrying this through (C.25) with $n_a(T_{\text{osc}}) \propto m(T_{\text{osc}}) f_K^2 \propto T_{\text{osc}}^2 f_K^2$ and $Y \propto n_a/T_{\text{osc}}^3 \propto f_K^2/T_{\text{osc}}$ gives

$$\Omega_K h^2 \propto m_K \frac{f_K^2}{T_{\text{osc}}} \propto \frac{1}{f_K} f_K^2 f_K^{2/(4+n)} = f_K^{(6+n)/(4+n)} \theta_i^2. \quad (\text{C.27})$$

With $n = 8.16$ the exponent is $(6+n)/(4+n) = 14.16/12.16 = 1.165$ — the law quoted throughout Part V. Restoring the constants with the instantaneous-onset rule gives the estimate $\Omega_a h^2 \simeq 0.090 \theta_i^2 (f_K/10^{12} \text{ GeV})^{1.165}$, and the direct integration of Part V raises the prefactor to 0.135; the per-step arithmetic of the estimate with the benchmark $f_K = 7.1 \times 10^{11} \text{ GeV}$ is carried in the audit, Appendix E.1, and the abundance returns $\theta_i \simeq 1.1$ rad at the integrated normalization.

Step 7 — the closure. Because the target is inherited at derived-conditional grade (GG-A10-Cosmos), $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2) = 0.2676$ (so $\Omega_K h^2 = 0.1215$), and the decay constant is at its calibrated benchmark, $f_K = 7.1 \times 10^{11} \text{ GeV}$, the law (C.27) contains *one* unknown — the misalignment angle. The mass follows from the susceptibility, the angle from the budget,

$$m_K = \frac{\sqrt{\chi_{\text{QCD}}}}{f_K} = 8.1 \mu\text{eV} (\pm 1.3\%), \quad \theta_i = 1.18 \pm 0.04 \text{ rad}, \quad (\text{C.28})$$

the angle being read from the calculation of record (Section 25); the closed-form scaffold above returns 1.16. And the frequency and coupling follow from the mass. The $\pm 1.3\%$ on the mass is the lattice susceptibility; the thermodynamics ($\chi(T)$, g_*) enters the angle, and the absolute placement is the calibration — the honest residual of a prediction with one calibrated scale and no free parameter after it.

C.7 The π -locked inventory: an inherited result, with the key steps reproduced

The cosmic fractions the Kaxion must saturate are not derived in this paper — they are an output of Arisaka (2026, A10-Cosmos), whose Technical Appendices carry the full topological derivation and its one named open lemma. To keep the present Part self-contained we do not reproduce that derivation in full; we reproduce only the key steps a reader needs to see *why* the target the abundance closure (Appendix C.6) aims at is $\Omega_{\text{DM}} = 0.2676$, and we mark exactly where the conjecture sits. The standard cosmology background is Dodelson & Schmidt (2020); Kolb & Turner (1990).

What is claimed. The present-epoch fractions are

$$\Omega_m = \frac{1}{\pi} = 0.3183, \quad \Omega_b = \frac{1}{2\pi^2} = 0.0507, \quad \Omega_{\text{DM}} = \frac{1}{\pi} - \frac{1}{2\pi^2} = 0.2676, \quad \Omega_\Lambda = 1 - \frac{1}{\pi} = 0.6817, \quad (\text{C.29})$$

against Planck 0.315, 0.0493, 0.265, 0.685 — every slice within $\sim 3\%$. In Arisaka (2026, A10-Cosmos) this is a *conditional theorem*: derived, not fitted (the integers and geometry were fixed upstream for reasons that never mention cosmology), but resting on one named open condition, the closed-packet conjecture that paper carries as its Conjecture VIII (Lemma VIII, the Hopf-fiber identification; T-Cosmos-17 the conditional theorem; C-Kaxion-c in this paper’s deck).

The key reproduction — one Hopf-fibered three-sphere. The engine is the boundary projection GDOS: the four-dimensional world is a lawful projection of the six-dimensional bulk, and the budget is how the projection measure divides. The relevant compact fiber is the Hopf fibration $S^1 \hookrightarrow S^3 \rightarrow S^2$, and the inventory is then unit-sphere volumes: $\text{Vol}(S^1) = 2\pi$, $\text{Vol}(S^2) = 4\pi$, $\text{Vol}(S^3) = 2\pi^2$. Spreading the critical density uniformly over S^3 , assigning matter to the fiber and dark energy to the base, gives

$$\Omega_m = \frac{\text{Vol}(S^1)}{\text{Vol}(S^3)} = \frac{2\pi}{2\pi^2} = \frac{1}{\pi}, \quad \Omega_b = \frac{1}{\text{Vol}(S^3)} = \frac{1}{2\pi^2}, \quad \Omega_{\text{DM}} = \frac{2\pi - 1}{2\pi^2} = \frac{1}{\pi} - \frac{1}{2\pi^2}. \quad (\text{C.30})$$

The two “flexible symbols” π and $2\pi^2$ are not chosen; they are a circle’s circumference and a three-sphere’s volume, the only measures a Hopf projection makes available. The one conjectural step is the *identification* that the projection fiber is this Hopf S^3 , with matter→fiber, dark energy→base, baryon→one unit of arc; settle its cohomology and the π -shape becomes a theorem (roadmap entry 2, Appendix E.5).

The robust part, and the honest residual. A *fraction of the critical density* is time-dependent, so $\Omega_m = 1/\pi$ is a statement about *now*. The parts that hold at every epoch — because dark matter and baryons both dilute as R^{-3} — are the ratios within matter,

$$\frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} = 0.159 \quad (\text{sky: } 0.156), \quad \frac{\rho_{\text{DM}}}{\rho_b} = 2\pi - 1 = 5.28 \quad (\text{sky: } 5.37), \quad (\text{C.31})$$

both epoch-independent and both matching observation. These are the robust π -locks. The absolute $\Omega_m = 1/\pi$ is the present-day split, which is where the coincidence lives: Arisaka (2026, A10-Cosmos) demotes “why now” from the 10^{120} fine-tuning of Λ (dissolved, because $\Omega_\Lambda = 1 - 1/\pi$ is a residual share of a fixed inventory) to a mild few-percent timing statement (we read the dial ~ 75 Myr past the moment $\Omega_m = 1/\pi$). For the present paper the operative fact is simply the number: the abundance closure of Appendix C.6 must deliver $\Omega_K h^2 = 0.1215$, and it does at the calibrated benchmark, returning the one initial condition $\theta_i \simeq 1.1$ rad.

D Detection, discovery, and the KAXIO instrument

This block turns the theory into an experiment. It derives how loud the Kaxion signal is and how long it takes to hear (the haloscope power and the radiometer wall), how a finite- Q cavity shapes the recorded line, and then two further courses: the KAXIO instrument sensitivity and its one-year coverage of the window, and the discovery-week error budget behind the verdicts of Part IX.

D.1 Haloscope sensitivity: the Sikivie power and the radiometer scan-rate wall

The last piece a reader needs to judge feasibility is how loud the Kaxion signal is and how long it takes to hear. Both follow from two equations every haloscope experimentalist carries in their head.

Step 1 — the conversion power. In a tunable cavity of volume V and mode form factor C (the overlap of the cavity mode with the external field, $\mathcal{O}(0.5)$ for the fundamental mode), immersed in a static field B_0 and tuned to the Kaxion mass, the resonant signal power is the Sikivie result (Sikivie, 1983)

$$P_{\text{sig}} = g_{K\gamma\gamma}^2 \frac{\rho_{\text{DM}}}{m_K} B_0^2 V C Q_L \frac{\beta}{1 + \beta}, \quad (\text{D.1})$$

with Q_L the loaded quality factor, $\rho_{\text{DM}} \simeq 0.45 \text{ GeV cm}^{-3}$ the local density, and β the antenna coupling. Two conventions are fixed once here. First, the survey rows of Table 17 quote the power at critical coupling with the delivered fraction absorbed into the stated Q_L (so $\beta/(1 + \beta) \rightarrow 1$ in (D.1), the convention of the simulation script); the over-coupled KAXIO budget of Part X carries its $\beta \simeq 2$ — a factor $\beta/(1 + \beta) = 2/3$ — explicitly. Second, the natural-to-SI conversion is worked here once, so that every table in the paper can be traced to it. Three constants suffice:

$$1 \text{ T} = 1.9535 \times 10^{-16} \text{ GeV}^2, \quad 1 \text{ cm} = 5.068 \times 10^{13} \text{ GeV}^{-1}, \quad 1 \text{ GeV}^2 (\text{power}) = 2.434 \times 10^{14} \text{ W}. \quad (\text{D.2})$$

Inserting the Kaxion's numbers at the CAPP-MAX-class row of Table 17 ($g_{K\gamma\gamma} = 1.2 \times 10^{-15} \text{ GeV}^{-1}$, $m_K = 8.1 \mu\text{eV}$, $B_0 = 12 \text{ T}$, $V = 6.4 \text{ L}$, $C = 0.60$, $Q_L = 10^5$), the factors of (D.1) become

$$\begin{aligned} \frac{\rho_{\text{DM}}}{m_K} &= \frac{0.45 \text{ GeV cm}^{-3}}{8.1 \times 10^{-15} \text{ GeV}} = 4.27 \times 10^{-28} \text{ GeV}^3, \\ B_0^2 &= (12 \times 1.9535 \times 10^{-16} \text{ GeV}^2)^2 = 5.50 \times 10^{-30} \text{ GeV}^4, \quad V = 6.4 \text{ L} = 8.33 \times 10^{44} \text{ GeV}^{-3}, \end{aligned} \quad (\text{D.3})$$

using $1 \text{ cm}^3 = 1.302 \times 10^{41} \text{ GeV}^{-3}$ from the second constant, so that

$$P_{\text{sig}} = g_{K\gamma\gamma}^2 \frac{\rho_{\text{DM}}}{m_K} B_0^2 V C Q_L = 1.69 \times 10^{-37} \text{ GeV}^2 = 4.1 \times 10^{-23} \text{ W}, \quad (\text{D.4})$$

the P_{sig} entry of that row — the characteristic yoctowatt level of all DFSZ searches, and the reason this is hard. One pitfall in this arithmetic deserves its own sentence: the $1/m_K$ in (D.1) is one over the *mass* — equivalently $\hbar\omega$ with $\omega = 2\pi\nu$ — and evaluating it with the frequency ν in place of ω silently inflates every power in the table by 2π .

Step 2 — the noise. The signal competes with thermal and amplifier noise of system temperature $T_{\text{sys}} = T_{\text{phys}} + T_{\text{amp}}$. The irreducible quantum floor at the Kaxion frequency is $T_q = h\nu/k_B = 0.094\text{ K}$ at 1.95 GHz, approached by Josephson parametric amplifiers and *beaten* by the squeezed-vacuum technique of HAYSTAC (HAYSTAC Collaboration, 2025).

Step 3 — the Dicke radiometer equation. The signal-to-noise ratio after integration time t on the matched line is

$$\text{SNR} = \frac{P_{\text{sig}}}{k_B T_{\text{sys}}} \sqrt{\frac{t}{\Delta\nu_{\text{eff}}}}, \quad \Delta\nu_{\text{eff}} = \left(\int g^2 d\delta \right)^{-1} \simeq 2.9\text{ kHz}, \quad \Delta\nu_a = \frac{\nu}{Q_a} \sim \nu \frac{\langle v^2 \rangle}{c^2} \sim 10^{-6} \nu, \quad (\text{D.5})$$

the celebrated radiometer equation with the matched filter’s effective width for the lab-frame template (Part IX): sensitivity improves only as the square root of time, and the axion line is intrinsically narrow ($Q_a \sim 10^6$, set by the halo velocity dispersion; the condensate linewidth $\Delta\nu_a \simeq 2\text{ kHz}$ is the scale, $\Delta\nu_{\text{eff}}$ the filter).

Step 4 — the scan-rate wall. Solving (D.5) for the time to reach a target SNR at one tuning, and dividing the band into resolution widths, gives the scan rate

$$\frac{d\nu}{dt} \propto \frac{1}{\text{SNR}^2} \left(\frac{P_{\text{sig}}}{k_B T_{\text{sys}}} \right)^2 \frac{1}{\Delta\nu_{\text{eff}}} \propto \frac{g_{K\gamma\gamma}^4 B_0^4 V^2 C^2 Q_L}{T_{\text{sys}}^2}. \quad (\text{D.6})$$

The steep dependences are the design drivers a newcomer should internalize: the scan rate goes as the *fourth* power of the coupling, so the DFSZ line (coupling ~ 2.5 times smaller than KSVZ) costs ~ 40 times the integration time; it goes as $B_0^4 V^2$, which is why haloscopes chase the strongest magnets and largest bores; and it falls as T_{sys}^2 , which is why every gain against the quantum limit — squeezing, single-photon counting — directly buys scan speed. Covering the Kaxion’s unique window [1.59, 2.13] GHz at DFSZ depth is, by (D.6), a multi-year campaign on a CAPP-class instrument (Section 36.4) — demanding, but demonstrated, not speculative. This is the “scan-time wall” referred to in Part VIII.

D.2 Finite- Q transfer functions

Single port. The steady-state response of the QGI oscillator (35.3) to a drive at frequency ω is obtained by Fourier-transforming the driven, damped harmonic oscillator $\ddot{x} + (\omega_0/Q_L)\dot{x} + \omega_0^2 x = F(t)/M$. With $x, F \propto e^{-i\omega t}$ the algebra gives $(-\omega^2 - i\omega\omega_0/Q_L + \omega_0^2)x(\omega) = F(\omega)/M$, so $x(\omega) = \chi(\omega)F(\omega)$ with the standard transfer function

$$|\chi(\omega)|^2 = \frac{1}{(\omega_0^2 - \omega^2)^2 + (\omega_0\omega/Q_L)^2}, \quad (\text{D.7})$$

a Lorentzian of fractional width $1/Q_L$ about the mode. The recorded power spectrum is the condensate template filtered by (D.7) plus noise,

$$S_{\text{out}}(\nu) = \mathcal{G} |\chi(\nu)|^2 [\mathcal{T}(\nu) + S_{\text{noise}}], \quad (\text{D.8})$$

with $\mathcal{T}$ the line template of Part IX and $\mathcal{G}$ the (independently calibrated) gain chain. Two standard consequences, quantified in the main text: on resonance the line is transmitted faithfully because the condensate width ν/Q_a is far inside the cavity width ν/Q_L ; off resonance, a tone detuned by the count test's 2.5% is suppressed by $[2Q_L\Delta\nu/\nu]^{-2} \simeq 4 \times 10^{-8}$ (−74 dB) at $Q_L = 10^5$.

D.3 The KAXIO instrument and the one-year projection

Part X proposes KAXIO, a second-generation instrument built to read the count across the fenced window and to cover the discrete mass window in one observing year. This course carries the sensitivity budget from the cavity parameters to the coverage claim, so that the proposal's headline can be reconstructed from the Sikivie power (Appendix D.1).

Step 1 — the instrument in one line of physics. KAXIO is two over-coupled resonator families reading the same condensate — two nine-cavity sub-arrays, independently tunable across the window. The design point is the ADMX-EFR magnet and cryostat with the cavities reorganized (Table 28),

$$B_0 \simeq 9.4 \text{ T}, \quad V \simeq 95 \text{ L per sub-array}, \quad C \simeq 0.40, \quad Q_L \simeq 10^5, \quad \beta \simeq 2, \quad T_{\text{sys}} \simeq 0.15 \text{ K}, \quad (\text{D.9})$$

i.e. an EFR-class electromagnetic channel, split in two. Inserting (D.9) and the Kaxion's DFSZ coupling $g_{K\gamma\gamma} = 1.2 \times 10^{-15} \text{ GeV}^{-1}$ into the Sikivie power (D.1) (with the over-coupling factor $\beta/(1+\beta) = 2/3$ carried explicitly) gives, per sub-array, $P_{\text{sig}}/k_B T_{\text{sys}} \simeq 80 \text{ Hz}$ (Table 29), worked from the same formula as Appendix D.1.

Step 2 — the time to a peak, and the scan-rate wall. From the radiometer equation (D.5), the integration time to reach a target significance at one tuning is

$$t_{\text{peak}} = \left(\frac{\text{SNR } k_B T_{\text{sys}}}{P_{\text{sig}}} \right)^2 \Delta\nu_{\text{eff}}, \quad \Delta\nu_{\text{eff}} \simeq 2.9 \text{ kHz}, \quad (\text{D.10})$$

so that reaching the design SNR on the matched line takes seconds of live time per sub-array (12.8 s to 5σ at the benchmark, Table 29), not the multi-year single-instrument campaign of a first-generation search, and the two sub-arrays pool their exposure when stacked. The scan cost to cover a band B follows from the scan-rate wall (D.6), $t_{\text{scan}} = B (dt/d\nu)$, and scales as $g_{K\gamma\gamma}^{-4} B_0^{-4} V^{-2}$ — the steep dependences that dictate the design.

Step 3 — one year covers the window. The prediction is the window $\nu_K \in [1.59, 2.13] \text{ GHz}$, with the seven rungs of (C.21) marked on it under the ladder ansatz. KAXIO's run plan (Part X) camps the benchmark rung first, then sweeps each band of the window at DFSZ depth with the cavity set that reaches it and the sub-arrays split across the band, camping the band's rungs before the one set exchange. Because the window is 0.54 GHz wide, a cavity bandwidth is 20 kHz, and the matched-line 5σ dwell is 12.8 s, the sweep at the wall rate (D.6) for the KAXIO parameters is a

matter of days of live time, and the full program with its protocol overheads is a one-year program:

$$t_{\text{window}} = \frac{B}{\Delta\nu_{\text{cav}}} t_{5\sigma} \simeq \frac{544 \text{ MHz}}{20 \text{ kHz}} \times 12.8 \text{ s} \simeq 97 \text{ h (live, one sub-array)}, \quad t_{\text{program}} \lesssim 1 \text{ yr, (D.11)}$$

after which either the line has been recorded — on a rung, reading Ξ_{cap} under the ladder ansatz, or between rungs, retiring the ansatz, with the count test run across the rest of the window — or the entire DFSZ-depth window is excluded, firing F-Kaxion-1c. This is the sense in which KAXIO is built to *settle* the Kaxion, not merely to search for it.

D.4 The discovery-week error budget and the verdicts

Part IX simulates the data of discovery week in advance and reads its verdicts from it. This course collects the error budget behind those verdicts in one place, so that a reader can see which systematics move a verdict and which cancel.

The two error classes. Every observable carries a *statistical* error that falls as $1/\sqrt{t}$ (the radiometer law) and a *systematic* error that does not. The systematic budget is *halo-dominated*: the local dark-matter density is known only to $\pm 22\%$ ($\rho_{\text{loc}} = 0.45 \pm 0.10 \text{ GeV cm}^{-3}$), which dwarfs the $\pm 7.7\%$ amplitude-instrument total. But — and this is the design principle — the halo systematic is *common mode*: it rescales every haloscope’s signal identically, so it moves detection *times*, never the identification verdicts.

The verdicts, and their budgets. The verdicts map onto the falsifiers and the count test:

1. **Shape** (halo confirmation). The camped line matches the boosted-Maxwellian condensate shape at the predicted frequency. Statistical only after calibration; the theory-side window on ν_K is a band-scan cost, not a within-panel error, and is priced separately (Appendix C.5).
2. **Modulation** (halo confirmation). The line’s energy centroid rides $\simeq 1375 \text{ Hz}$ above its edge at ν_K and breathes annually with amplitude $\simeq 74 \text{ Hz}$; the sidereal daily signature is $\simeq 1.3 \text{ Hz}$. These are halo-model characteristics with $\pm 10\%$ model systematics; their diagnostic power is that no laboratory artifact reproduces either period. The daily signature is out of a single cavity’s reach (below 1σ even in a full-duty year) — the observable KAXIO’s aperture (Appendix D.3) exists to buy.
3. **Count** (the count test of Part VI). A single-mode cavity tuned to one rung is blind to a detuned tone by $\simeq -74 \text{ dB}$ at $Q_L = 10^5$; the $\pm 5\%$ on Q_L shifts this by $\mp 0.4 \text{ dB}$ on a 74 dB suppression — systematics-proof by four orders. The test is a deliberate act — split the sub-arrays and interrogate the other rungs — and at camped SNR each rung’s verdict arrives within a day.

Why the discovery cannot be mistaken — and why the refutation is unarguable. The verdicts that carry the geometry (the location, the window, and the count) are engineered to be free of the halo systematic — the fence and the address are frequency data, the count is a yes-or-no interrogation of the window — so they sit in the shortest, cleanest rows of the budget. The same

figures that make a real discovery unmistakable make a refutation a death certificate: a line that fails the shape, sits still through the seasons, sits outside the window, or sits off the DFSZ locus does not become an anomaly — it falsifies the calibrated placement; a partner inside the window refutes the count and names a registered mode. A theory confirmed by locked numbers can be killed by any one of them, and the Kaxion accepts both verdicts from the same data.

E Registries, audits, and roadmaps

This block is the bookkeeping a reader and a future co-author need: the end-to-end numerical audit that re-executes the closure pipeline row by row, the registry of claims and falsifiers, the table of what each Kaxion number inherits from which parent paper, the deck of named theorems and conjectures with their status, and the promotion roadmap that says what one computation would turn each open conjecture into a theorem.

E.1 Consolidated benchmark audit: the closure pipeline end to end

Every number of Part V, in evaluation order, with inputs on the left. Table 39 is the audit itself: read top to bottom it re-executes the entire closure pipeline, each row consuming only rows above it, and its final row returns the observed relic density to within rounding — the round-trip check-mark that the chain closes:

Table 39: End-to-end audit. Inputs of the analytic scaffold, which this appendix audits and which is *not* the calculation of record (Section 25): $\chi_{\text{QCD}}^{1/4} = 75.5(5)$ MeV, $b = 4.08(20)$, $T_{\text{QCD}} = 0.150(15)$ GeV, $g_* = g_{*s} = 73$, $f_K = 7.1 \times 10^{11}$ GeV, target $\omega = 0.1215(18)$, $M_{\text{Pl}} = 1.221 \times 10^{19}$ GeV, $s_0 = 2891 \text{ cm}^{-3}$, $\rho_c/h^2 = 1.0537 \times 10^{-5} \text{ GeV cm}^{-3}$. Every input is printed so that the chain can be recomputed independently rather than taken on trust.

Step	Expression	Value
Susceptibility	$\sqrt{\chi_{\text{QCD}}}$	$5.70(8) \times 10^{-3} \text{ GeV}^2$
Normalization (estimate)	Eq. (25.1), $f = 10^{12} \text{ GeV}$, $\theta_i = 1$	$\Omega h^2 = 0.090$
Normalization (integrated)	Eq. (12.2), same inputs	$\Omega h^2 = 0.135$ (adopted)
cross-check	community formula, 0.12 at $9 \times 10^{11} \text{ GeV}$	same statement
Decay constant	calibrated; cap loading $\Xi_{\text{cap}} = 43$	$f_K = 7.1 \times 10^{11} \text{ GeV}$
Mass	$\sqrt{\chi_{\text{QCD}}}/f_K$	$m_K = 8.1 \times 10^{-15} \text{ GeV}$ $= 8.1 \mu\text{eV}$
Frequency	$241.80 \text{ MHz} \times m[\mu\text{eV}]$	$\nu_K = 1.95 \text{ GHz}$ (one line; window $1.59\text{--}2.13 \text{ GHz}$)
Misalignment angle	scaffold: $(0.1215/0.135)^{1/2}(10^{12}/f_K)^{0.583}$	1.16 rad ; record: 1.22 harmonic, 1.18 cosine
Onset	$T_{\text{osc}}^{6.08} = m_K T_{\text{QCD}}^{4.08} M_{\text{Pl}}/42.6$	$T_{\text{osc}} \simeq 1.0 \text{ GeV}$; $t_{\text{osc}} \simeq 0.3 \mu\text{s}$
Onset mass	$(0.15/1.0)^{4.08}$	$m_K(T_{\text{osc}}) \simeq 4.4 \times 10^{-4} m_K(0)$
Round trip	integration forward at $\theta_i = 1.18$	$\Omega h^2 = 0.1215$ ✓
Coupling bracket	$8/3 - 1.92(4)$	$0.75(4)$
Coupling	$0.75/(137.036 \cdot 2\pi f_K)$	$g_{K\gamma\gamma} = 1.23 \times 10^{-15} \text{ GeV}^{-1}$
Lifetime	$64\pi/(g_{K\gamma\gamma}^2 m_K^3)$	$1.7 \times 10^{50} \text{ s}$

E.2 Claims and falsifier registry

The master theorem-versus-conjecture ledger is Table 13. The registered falsifiers, restated for the record: **F-Kaxion-1c** (DFSZ-depth exclusion of the unique window $\nu_K \in [1.59, 2.13]$ GHz, Eq. 36.1) and **F-Kaxion-1d** (a confirmed compact-angle-like line at any frequency outside that window: the prediction is unique, so a wrong location falsifies); the **count test** (a second in-window line established at the same grade as the first refutes the count, T-Kaxion-5, and names a light registered mode) and the **address**, **F-Kaxion-1e** (a confirmed in-window line off the DFSZ locus at $> 5\sigma$, which falsifies $E/N = 8/3$ and the theorem-level sheet); **F-Kaxion-1a** and **1b** — the partner tone and the channel ratio — are retired with the doublet (Part VI), their identifiers kept unassigned, and the count test replaces them at zero hardware cost wherever a scan crosses the window. Constitutional registration: **GG-A1-Constitution P7** and its kill criterion. This paper’s cohort-global identifier families, under the program’s Prefix-Suffix- N convention: **T-Kaxion- N** (theorems, Parts II–VI), **F-Kaxion-1c–1e** (falsifiers, above; 1a/1b retired), **O-Kaxion-1–69** (objections, Part XIV), **PS-Kaxion-1–8** (problems solved, §72.1), and **I-Kaxion-1–10** (structural innovations, §72.2). The count test and the address test the parents’ registered accounting and the DFSZ placement independently of the calibration; (c) tests the calibrated placement and the ladder as fenced; (d) tests the placement, or, at the randomized-angle mass, the pinning conjecture. Numerical decision rules, pre-declared: the count test fires on a second in-window line established at the same grade as the first (a persistent excess with the halo line shape, surviving the same vetoes); the address fires on a confirmed line off the DFSZ locus at $> 5\sigma$; F-Kaxion-1c fires on $g_{90} \leq 0.8 g_{\text{DFSZ}}$ across the window; F-Kaxion-1d fires on a confirmed compact-angle-like line outside it.

E.3 Cross-paper inheritance table

The Kaxion is not free-standing: every number it predicts is inherited from a theorem proved upstream in the **GG-Theory** cohort, and this paper’s task is to *deploy* those theorems, not to re-prove them. Table 40 names, for each headline quantity, the parent paper and the specific result it descends from — the ledger a reader needs to audit provenance, and the reading list a student needs to go deeper.

Table 40: What each Kaxion number inherits, and from where. “This paper” marks the deployment step where the inherited theorem meets a standard-physics input to produce the observable. A row with no companion named in the middle column would be a claim this paper has not earned, and there is none.

Quantity	Parent / source	Inherited result
Bulk geometry, integers $(1, 3, 5; 7, 35)$	Arisaka (2026, A2-GG6)	6D bulk with scale + phase; the OGN-5 ledger fixing $K = 35$
Two-form B_2 , $2/5$ baryonic trace ratio	Arisaka (2026, A3-BI6)	BI-6 anomaly identity; the locked vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$
Boundary law, current package	Arisaka (2026, A5-GDOS)	GDOS projection permitting the portals
Attractor selection, single patch	Arisaka (2026, A6-GRIP)	GRIP cascade births/selects/stabilizes the Kaxion as the GGA
5:2 current locking	Arisaka (2026, A7-Qi4)	four-current identity fixing $(L, L - N) = (5, 2)$
$E/N = 8/3$, $\alpha^{-1} = 2^J$	Arisaka (2026, A9-SM)	SU(5) hypercharge normalization; the DFSZ line
Target $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$	Arisaka (2026, A10-Cosmos)	π -locked inventory; the abundance the relic must saturate
Ground-state selection	Arisaka (2026, A11-Origin)	GGA: the Kaxion is the unique deepest vacuum
Lock-cluster cousin $M_{\text{lock}} = 3.85$ TeV	Arisaka (2026, P1-KK)	the same $K = 35$ read at the TeV lock
Mass $m_K f_K = \sqrt{\chi_{\text{QCD}}}$	this paper + lattice QCD	the pNGB relation, with f_K calibrated once

E.4 Named-theorem-and-conjecture deck

Table 41 is the honest ledger of what the paper proves and what it fences. It separates the theorem-level claims (forced by geometry and standard physics, with no continuous freedom) from the conjecture-fenced ones (each resting on a single, named, reader-checkable computation), and gives each its status. What the discovery figures depend on is worth stating exactly rather than waving away: the coupling sheet is theorem-level, the count derived-conditional, and the address 1.95 GHz carries the calibration (O14) and C-Kaxion-b; retired rows are kept as the record of the withdrawn doublet configuration. The figures are drawn at the benchmark address, so they inherit that fence.

Table 41: The claims of the paper, graded. T = theorem-level; D* = derived-conditional (the derivation is complete, its premises rest on named open computations); C = conjecture-fenced (one named computation from promotion, see Appendix E.5); A = ansatz (the cap-loading ladder, §26.3). Every claim the paper makes appears once, with its grade and with what it would take to promote it. T is theorem-level; D* is a completed derivation whose premises are named open computations; C is fenced by a single named computation, and the computation is named rather than gestured at. The grades are assigned before the discussion Parts argue for the result, so a reader can hold this table open and check that no grade improved along the way.

ID	Statement	Grade	Rests on
T-Kaxion-1	A forced compact-angle pseudoscalar exists and is light	T	shift symmetry; dualization (App. C.1)
T-Kaxion-2	$m_K f_K = \sqrt{\chi_{\text{QCD}}}$ (gluonic portal)	T	PQ mechanism + lattice (App. C.3)
T-Kaxion-3	$E/N = 8/3$: the Kaxion sits on the DFSZ line	T	SU(5) normalization (App. C.3)
T-Kaxion-4	Baryonic ledger entry $-2/5$ of X_4 ; no channel below M_{lock}	T	anomaly ledger (App. C.4)
T-Kaxion-5	One light QCD-coupled angle: the count	D*	mode analysis (GG-A3-BI6, self-duality of the Green–Schwarz tensor; O13/O14), conditional on five registered exposures (§9.4)
C-Kaxion-a	(retired) the $1/K$ split of the withdrawn doublet	—	retired with the count (Part VI)
C-Kaxion-b	$\Xi_{\text{cap}} = K + N + L = 43$ (absolute mass; the ladder is an ansatz)	C	a defining computation for the cap loading, not yet formulated (§26.4)
C-Kaxion-c	$\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$	C	closed-packet conjecture, Conj. VIII / T-Cosmos-17 (Arisaka (2026, A10-Cosmos))
C-Kaxion-d	Inflation-era pinning, $m_{\text{pin}} \gg H_1$	C	the angle’s potential on GG-A10-Cosmos’s primary inflation (§13.5)
MC	Misalignment completeness: the condensate is the sole saturator of the π -locked share	C	GG-A10-Cosmos’s conjecture: no second cold component in the ledger, no late entropy injection (§15.3)
Ladder ansatz	The window is the discrete cap ladder $\Xi_{\text{cap}} \in \{35, \dots, 47\}$ around the calibrated benchmark	A	Ansatz clauses (ii)–(iii) (§26.3); retired by a line off every rung
F-Kaxion-1c–1e	The surviving master falsifiers (1a/1b retired with the doublet)	—	the above, read as kill criteria

E.5 Promotion roadmap for the open items

Every conjecture in the deck is fenced deliberately, and each is promotable to a theorem by *one* named, finite computation — not by new physics or new data. Table 42 lists them in priority order, with the computation that would settle each and the consequence of settling it. This is the paper’s own to-do list, and the natural set of problems for a student entering the program.

Table 42: The open items, the single computation that promotes each, and the payoff. Ordered by sharpness. Three columns: the open item, the single computation that would close it, and what closing it buys. The ordering is by sharpness rather than by importance, because a vague open problem cannot be worked on and a sharp one can. Each row is a piece of work somebody could start tomorrow, which is the standard the table is meant to meet.

Open item	The one computation that promotes it	Payoff
$\Xi_{\text{cap}} = 43$ (C-Kaxion-b)	Formulate a defining computation for the cap loading on the admissible class of Arisaka (2026, A3-BI6) (which carries no orbifold cap): the manifold, the class and the integrand; then evaluate it, together with the orientation computation of O14	Collapses the seven-value window to a point, or retires the ladder ansatz; fixes the absolute mass
Orientation convention (O14)	Derive the anti-self-dual sector's orientation convention and cap normalization	Decides which chiral norm the remnant carries; on present readings neither reproduces the calibrated f_K (Appendix C.1), so the computation either supplies the missing factor or retires the calibration
π -inventory (C-Kaxion-c)	Settle the closed-packet conjecture (Conj. VIII, Lemma VIII) in Arisaka (2026, A10-Cosmos)	Turns $\Omega_m = 1/\pi$ into a theorem; leaves the declared epoch input as a calibration
Misalignment completeness (MC)	Show, in Arisaka (2026, A10-Cosmos) 's history, that the compact-angle condensate is the sole saturator of the inventory: no second cold component in the ledger and no late entropy injection	Turns the closure from a consistency statement into a derivation of the angle; a measured second dark component, or a found line with Ω_K short of the inventory, refutes it
Inflation-era pinning (C-Kaxion-d)	Compute the compact angle's effective mass on Arisaka (2026, A10-Cosmos) 's primary-inflation background, and show it switches off below the lock crossing	Turns the single-patch history from assumed to derived; or, if $m_{\text{pin}} \lesssim H_1$, moves the line to $\gtrsim 6$ GHz (misalignment alone) and out of the window
Nucleon-spin sheet	Extend the KSVZ-class chiral computation to other completion classes	Bounds the completion dependence of the archived spin-port record
Massive-tower inflow (O13)	Carry the anomaly-inflow bookkeeping through the massive baryonic tower	Closes the registered gap between the termination and the tower's spectrum

F Retired readings of earlier revisions: the record

Earlier revisions of this paper carried three readings that the present revision does not: a *doublet* — two light QCD-coupled tones split by $\Delta\nu/\nu = 1/K = 1/35$ — as the paper’s first fingerprint; a *channel ratio* $P_B/P_\gamma = 25/4$ between an electromagnetic port and a baryonic (spin) port as its second, with a *harmonic-booster* instrument concept designed to read it; and a *four-peak* coincidence template built from the two. All three rested on a premise the anomaly companion never supplied — a second light QCD-coupled angle — and all three are withdrawn by that companion’s registered mode analysis (its section on the self-duality of the Green–Schwarz tensor): the anti-self-dual tensor leaves exactly one light pseudoscalar, nothing is eaten, the boundary condition that terminates $U(1)_B$ leaves no baryonic channel below M_{lock} , and the count is one (Part VI). The main text of this paper states the surviving claim only — a single line in the fenced window, on the DFSZ address, alone — and carries none of the withdrawn apparatus. This section collects the record of what was withdrawn, in the words of the revisions that carried it, so that a reader comparing versions can see the change rather than discover it. Nothing here is a claim of the present paper. The derivation courses behind the withdrawn readings are kept, marked as archived, in Appendices F.6 and F.7 and in the archived two-port chain and inversion step of §F.9–F.10 below; the live courses of Appendices C.4 and D.2 carry only their withdrawal notes.

F.1 The doublet and the $1/K$ split

As the withdrawal was recorded. Earlier revisions of this paper carried both directions as light: a doublet geared at $1/K$, with a twin-peak split $\Delta\nu/\nu = 1/35$ as the paper’s first fingerprint and a matched-channel power ratio $25/4$ as its second; a later revision kept that algebra as the ungauged sector’s spectrum and had the $U(1)_B$ vector eat the tensor’s shift, leaving the holonomy mode as the single line with the partner absorbed at $\sim 10^{13}$ GeV. Both readings are withdrawn, and for one reason: the premise of a second light QCD-coupled angle was never supplied by the anomaly companion, and its registered mode analysis (the self-duality section of GG-A3-B16, O13/O14) settles the count at one — the anti-self-dual tensor leaves exactly one light pseudoscalar per chirality, only one chirality’s zero mode is present on the physical class, and the $U(1)_B$ vector is terminated by a Dirichlet boundary condition that eats nothing. There is no light partner tone, no $1/35$ split, no $25/4$ channel ratio, and no longitudinal Z'_B at 10^{13} GeV; the Z'_B carries one mass, $M_{\text{lock}} \approx 3.85$ TeV, and the Kaxion is the tensor remnant with its full X_4 couplings. What survives of the doublet apparatus is its cheapest measurement, reframed as a null test: a paired re-binning of any scan crossing the window costs nothing and tests the count (Part VI).

What Part VI replaced. Earlier revisions gave this Part to a twin peak: a doublet split by $\Delta\nu/\nu = 1/K = 1/35$, read from a two-field gearing of the tensor’s dual with the colored Wilson line, with a matched-channel power ratio $(5/2)^2 = 25/4$ as the second fingerprint; an intermediate revision kept the doublet as the ungauged sector’s spectrum and retired it dynamically, with the $U(1)_B$ vector eating one combination. The apparatus is withdrawn — premise, algebra, doctrine box and all — by the anomaly companion’s registered mode analysis (GG-A3-B16, self-duality of the

Green–Schwarz tensor; O13/O14): the second light angle was never there, nothing is eaten, and the count is one. The withdrawal is recorded rather than erased, because the replacement claim is *stronger* in the only way that matters to an instrument: a single line is a cleaner null hypothesis than a doublet, and the doublet apparatus survives in exactly one form — as the count test below, the cheapest discriminating measurement the paper asks for.

F.2 The channel ratio and the spin ports

[The text below is the spin-port section of Part VIII as earlier revisions carried it, moved here unchanged except for its heading levels. Its target — the 25/4 channel ratio — is withdrawn; the KSVZ-class spin couplings it quotes stand, and Part VIII states them without the instrument.]

Principle — and the withdrawal of its target. Concept B was designed to read the second lock of earlier revisions: an electromagnetic port and a *spin* port — nuclear or electron — in the same condensate, with the deconvolved power ratio targeted at 25/4. That target is withdrawn with the doublet (Part VI): the boundary condition that terminates $U(1)_B$ leaves no baryonic channel below M_{lock} , so there is no gauged ratio for a matched pair of receivers to read. What survives at the chiral level are ordinary nucleon- and electron-*spin* couplings on the KSVZ-class sheet, and the sheet must be printed with its verdict: for a KSVZ-class heavy-quark completion in the $(\mathbf{3}, \mathbf{1}, \frac{2}{3})$ representation, $C_p = -0.47 \pm 0.03$ and $C_n = -0.02 \pm 0.03$ (Grilli di Cortona et al., 2016), with the electron coupling loop-suppressed at $C_e \sim 3 \times 10^{-4}$ — and at $f_K = 7.1 \times 10^{11}$ GeV every one of these ports is out of reach of the instrument classes below on any useful timescale. The electron-magnon port has no target at this f_K . The spin ports are therefore *archived*: the engineering that follows is recorded because it is sound technology this program may yet need — and because retiring a concept without its record would hide the change — not because this paper still supplies it a measurement. The spin port couples to the nucleon channel (NMR-class, CASPER-style (CASPER Collaboration, 2021)) or the electron channel (ferromagnetic-magnon class); resonance requires matching the Larmor frequency to ν_K .

The Larmor arithmetic. The spin ports are frequency-rigid, $\nu_L = \gamma B$ with $\gamma_p = 42.58$ MHz/T (proton) and $\gamma_e = 28.0$ GHz/T (electron). At the Kaxion window,

$$\nu_K = 1.95 \text{ GHz} : \quad B_e = (70 \pm 7) \text{ mT}, \quad B_p = 45.8 \text{ T (impractical)}, \quad (\text{F.1})$$

so the practical spin port is the *electron* channel: a (70 ± 7) mT ferromagnetic-magnon hybrid operated alongside the Concept-A cavity — a sub-tesla magnet, not an exotic one. The direct nucleon port is out of Larmor reach at the window; and with the 25/4 target withdrawn (Section 32), what these stages would have measured no longer exists at laboratory frequencies — the record below stands as engineering, not as a measurement program.

F.2.1 Non-Larmor nucleon readout: the principle

The 45.8 T wall is a statement about one readout *mechanism*, not about the coupling. A magnetic-resonance port is frequency-rigid because it ties the readout frequency to the bare gyromagnetic relation $\nu_L = (\gamma_N/2\pi)B_0$; reaching ν_K then demands an impractical B_0 . *Non-Larmor* readout means breaking that tie — supplying the resonance frequency by some route other than a laboratory static field, or reading an observable that is not spin precession at all. Two routes follow from the two ways the gauged- $U(1)_B$ portal reduces onto ordinary matter.

The axial (spin) reduction. Chirally reduced, the $a F_B \tilde{F}_B$ operator induces an effective coupling to nucleon spin — an “axion wind” $\sim (\partial_\mu a) \bar{N} \gamma^\mu \gamma_5 N$ — the same observable CASPER-class instruments read (CASPER Collaboration, 2021). To make it resonant at ν_K without a 45.8 T applied field, one supplies the field *internally* (Concept B') or up-converts the response (Concept B'').

The vector (charge) reduction. $U(1)_B$ gauges baryon *number*, a scalar charge proportional to nucleon count, not to spin. Its most faithful readout is therefore not precession at all but a coupling to baryon density — an oscillating, composition-dependent force on bulk matter at ν_K , read mechanically with no magnet and no spins (Concept B'''). This route reads the very quantity $U(1)_B$ gauges, and is the conceptual complement of the spin ports.

We record all three. Each was developed as a candidate baryonic port for the retired matched-ratio test; each is fenced by the same model-level input — the chiral reduction that converts a portal coupling into an instrument response — and each now stands as archived technology awaiting a framework that supplies it a target.

F.2.2 Concept B': internal-field (hyperfine-enhanced) nuclear magnon resonance

The lead candidate. In a magnetically ordered crystal a nucleus does not feel the applied field; it feels the *hyperfine field* B_{hf} of the ordered electronic moments, which routinely reaches tens to hundreds of tesla. The nuclear resonance is thereby pushed into the gigahertz band at zero or small applied field,

$$\nu_n = \frac{\gamma_N}{2\pi} (B_{\text{hf}} + B_0), \quad (\text{F.2})$$

and, decisively, the electronic system *enhances* both the rf drive on the nucleus and its radiated signal by the hyperfine enhancement factor $\eta \sim B_{\text{hf}}/B_a \sim 10^2\text{--}10^4$ (with B_a the anisotropy field) (Turov & Petrov, 1972); the electron-spin analog of an enhanced magnon port is the quantum-limited QUAX ferromagnetic haloscope (Crescini et al., 2020). A survey of ordered nuclei brackets the window:

$$^{55}\text{Mn} (\text{MnF}_2) : \nu_n \simeq 0.66 \text{ GHz}; \quad ^{159}\text{Tb} : \sim 3 \text{ GHz}; \quad ^{165}\text{Ho} (B_{\text{hf}} \sim 7\text{--}8 \times 10^2 \text{ T}) : \sim 6\text{--}7 \text{ GHz},$$

so material engineering plus a sub-tesla trim field B_0 places an enhanced transition on $\nu_K = 1.95 \text{ GHz}$ — the resonance frequency is supplied by the crystal, not by a 45.8 T magnet. Three quantitative advantages make this the natural baryonic port.

(i) *Near-unit thermal polarization, for free.* At the resonance frequency the nuclear Zeeman gap is $h\nu_n/k_B = 93.6 \text{ mK}$ at 1.95 GHz; in a dilution refrigerator at $T = 10 \text{ mK}$ the equilibrium

polarization is

$$p = \tanh \frac{h\nu_n}{2k_B T} = \tanh(4.68) = 0.9998, \quad (\text{F.3})$$

so the ensemble is essentially fully polarized with *no* dynamic nuclear polarization — the “frozen core” a conventional NMR search must engineer is here the thermodynamic ground state.

(ii) *A dense, coherent spin bath.* A 1 cm^3 crystal carries $N \sim 2 \times 10^{22}$ resonant nuclei; the enhanced transverse magnetization couples to a tuned rf resonator (or hybridizes directly with the Concept-A microwave mode), read by a quantum-limited Josephson parametric amplifier or, preferably, a single-microwave-photon counter at 1.95 GHz.

(iii) *Built-in matching.* Because the port is itself a resonant, polarized, $Q \sim Q_a$ receiver in the same cold volume as the photon cavity, its instrument factor is calibrated against the *same* condensate the photon port sees — the matched configuration any two-channel ratio measurement would require, recorded here for whatever future framework supplies one.

F.2.3 Concept B'': heterodyne (up-converted) nuclear readout

A bridge technique using ordinary materials and fields. Hold a polarized nuclear ensemble at a modest field so its bare transition lies at $\nu_0 \sim$ tens of MHz, and apply a strong rf “pump” at ω_p that dresses the spins; the Kaxion at ω_K then appears not as a direct resonance but as a *sideband* at $\omega_p \pm \omega_K$. Frequency-conversion (heterodyne) readout already extends axion–photon haloscopes beyond their bare cavity band; the spin analog lets a low-field nuclear sample respond to ν_K at a static field of a few tesla rather than 45.8 T. The conversion efficiency and the pump-induced noise must be calibrated — so B'' is a less clean ratio instrument than B' — but it requires no engineered hyperfine material and is the natural pathfinder while B' crystals are developed.

F.2.4 Concept B''': the baryon-number force readout

The conceptually cleanest port, reading the vector charge directly. The Kaxion’s gauged- $U(1)_B$ coupling makes a dense “source” mass a baryon-number charge whose Kaxion-modulated field exerts an oscillating, equivalence-principle-violating acceleration on a “test” mass, with amplitude set by the baryon-number-to-mass ratio N_B/m — which differs between materials through nuclear binding and neutron excess. The signal is a force at ν_K , detected with an optomechanical or levitated resonator tuned to the line (Carney et al., 2021) — no static field, no spin polarization, no Larmor condition. Two honest limitations fix its place in the program. First, the $U(1)_B$ gauge boson is heavy (the leptophobic TeV vector of the companion P1-KK), so any *long-range* baryonic force is Yukawa-killed; the readable operator is the local, anomaly-induced modulation of nucleon number, and its laboratory normalization carries the same chiral fence as B'/B''. Second, a mechanical resonator at 1.95 GHz is at the frontier of optomechanical-crystal technology, so the concept is most natural as a broadband search or awaits gigahertz-class optomechanics. We therefore register B''' as the magnet-free complement to B' rather than the window’s first instrument. Table 43 compares the three concepts — coupling read, resonance route, technological maturity, and the role each plays in the staged program — and its reading order is its recommendation: B' first, B'' as the up-conversion alternative, B''' as the long-horizon complement.

Table 43: The baryonic (nucleon) port at the window $\nu_K = 1.95$ GHz: two Larmor reference rows and the three non-Larmor concepts. “Reads” names the channel each port accesses; the *absolute* normalization of every nucleon row carries the KSVZ-class sheet of Part VIII, and the ratio these ports were designed to test is withdrawn (Part VI) — the table stands as the archived engineering record.

Port	Static field	Resonance set by	Reads	Status / readout
Direct proton NMR (Larmor)	45.8 T	applied field $\gamma_p B_0$	nucleon spin (g_{KBB})	impractical at the window (reference row)
Electron magnon (Larmor)	70 mT	applied field $\gamma_e B_0$	electron spin (g_{ae}^{eff})	mature, QUAX-class (Crescini et al., 2020); reads the J_e lock, not J_B
B' internal-field nuclear	$\lesssim 1$ T trim	hyperfine $B_{\text{hf}} \sim 10\text{--}10^2$ T	nucleon spin (g_{KBB})	lead ; $p = 0.9998$ at 10 mK; enhanced rf + JPA / photon counter
B'' heterodyne nuclear	few T	pump sideband $\omega_p \pm \omega_K$	nucleon spin (g_{KBB})	pathfinder; ordinary NMR material; conversion calibration
B''' baryon-number force	none	mechanical mode $= \nu_K$	baryon number ($U(1)_B$ charge)	magnet/spin-free; GHz optomechanics frontier; or broadband

What gates the nucleon port. All three concepts share one gate and one virtue. The gate is the absolute magnitude: the integration time any nucleon port needs is set by the chiral-level spin coupling, and the sheet above prices it out of reach at this f_K ; a theory deliverable that should precede metal. The virtue is that the harmonic-booster observable is a *ratio* in a shared condensate: ρ_{loc} , the halo velocity distribution, and the coherence time cancel, so the systematic budget reduces to the *relative* calibration of the two co-located ports — the one quantity an experiment controls. The sequence the record keeps: the electron port was to demonstrate the matched-ratio method first; B' was the baryonic-port R&D target once its chiral calibration was in hand; B'' and B''' are the pathfinder and the magnet-free complement.

F.2.5 Feasibility band: signal, integration time, and the chiral coefficient

The one number the nucleon port waits on can be isolated and its impact bounded, so that “awaiting a chiral computation” becomes a quantitative band rather than an open shrug. Write the effective Kaxion–nucleon-spin coupling in the standard form

$$g_{aNN} = C_N \frac{m_N}{f_K} = C_N \times 1.3 \times 10^{-12} \text{ GeV}^{-1}, \quad (\text{F.4})$$

with $C_p = -0.47 \pm 0.03$ and $C_n = -0.02 \pm 0.03$ on the KSVZ-class sheet, $m_N = 0.94$ GeV, and $f_K = 7.1 \times 10^{11}$ GeV. The *order* is pinned by the $1/f_K$ suppression carried by every Kaxion portal; the dimensionless coefficients are the chiral matrix elements of the KSVZ-class sheet above —

$C_p = -0.47 \pm 0.03$, $C_n = -0.02 \pm 0.03$, $C_e \sim 3 \times 10^{-4}$ (loop) (Grilli di Cortona et al., 2016) — no longer an $\mathcal{O}(1)$ unknown but a computed set, and a small one.

A resonant spin haloscope accumulates significance as a radiometer, $\text{SNR} \propto g_{aNN}^2 \sqrt{t_{\text{int}}}$, so the integration time to a fixed significance scales as

$$t_{\text{int}} \propto g_{aNN}^{-4} \propto c_N^{-4}. \quad (\text{F.5})$$

At the Concept-B' benchmark — $N \simeq 2 \times 10^{22}$ polarized nuclei ($p \simeq 0.9998$, Eq. F.3), hyperfine enhancement $\eta \sim 10^3$, and a quantum-limited readout co-located with an *already detected* photon line — a spin-port confirmation campaign would be comparable to the photon discovery run only for a coupling coefficient of order one, and lengthens as the fourth power as the coefficient falls (F.5): at the sheet's $|C_p| \simeq 0.47$ the time cost is a factor ~ 20 , at $|C_n| \simeq 0.02$ it is a factor $\sim 10^6$, and the loop-level electron port is beyond any of it. The proton-spin port is the one entry within sight of feasibility, and only for a signal already found; nothing in this Part's identification program rides on it.

One fact keeps the record worth carrying: a spin port would ride on the photon detection — a retrofit to an instrument that has already found the line — so its integration time would be spent confirming a second channel of a known signal, not searching blind. Should a future framework supply the port a locked target, the engineering above is where it would start; this paper supplies none.

F.3 The harmonic-booster module of the KAXIO design

[The booster-module paragraph and the booster engineering section of Part X as earlier revisions carried them, moved here unchanged except for heading levels. The KAXIO design of Part X now carries the two sub-arrays only.]

The harmonic-booster module (archived). The spin port — Part VIII's electron channel (Concept B), archived with the coupling sheet of Part VIII: the computed electron entry is loop-level, $C_e \sim 3 \times 10^{-4}$, and at this f_K the port has no target — was to be realized as two electron-magnon hybrid resonators — ferrimagnetic spheres strongly coupled to microwave cavities, the architecture demonstrated by the QUAX collaboration (Crescini et al., 2020) — biased at $B_e = (70 \pm 7)$ mT so the magnon mode sits at ν_+ (first hybrid) and ν_- (second). The module lives in a magnetically shielded side cryostat at 10 mK: a mu-metal enclosure plus an active compensation coil suppress the main magnet's stray field to < 0.1 mT at the module (a 10^{-3} perturbation on B_e , servo-corrected). The booster's engineering requirement — the fenced input of Part IX — is a per-tone line sensitivity $\text{SNR}_B = 5\sqrt{t/5\text{d}}$; the pre-construction R&D program (one year, bench-top, before the main build) exists precisely to demonstrate this figure, and the archived concept stands or falls with it, declared here so the record is honest (here and throughout, booster exposure times count *live* days; the run plan's duty budget converts them to calendar time).

This section is carried as a record. The booster was designed for the withdrawn 25/4 ratio test, and the KSVZ-class sheet of Part VIII leaves it no target at this f_K ; the γ -only baseline is the plan,

and nothing below gates the run. The engineering case is kept at full depth — the coupling budget, the linewidth budget, the shield design, and the noise staging from the first-generation chain to the quantum limit — because the subsystem has no running precedent at exactly this operating point and the case may serve a future framework that supplies a target.

Why the spin-port signal is weak, and why that is not fundamental. The middle row of Figure 22 shows a week-one booster signal of only $\text{SNR}_{\text{line}} \simeq 5.9$ per tone against the γ ports' peak-bin $\text{SNR} \sim 3.3 \times 10^2$. Three factors, in order of importance. (i) *The per-spin coupling is minute:* the condensate drives each electron spin through the axion-wind interaction with an effective Rabi rate suppressed by $1/f_K$; only the collective enhancement of $N_s \sim 10^{21}$ polarized spins, coherently coupled to one cavity mode, lifts the deposited power to detectability. (ii) *The magnon linewidth is broad:* the hybrid mode inherits the ferrimagnet's damping, $\kappa_m/2\pi \simeq 1$ MHz — five hundred times the 2 kHz signal line — so the readout integrates receiver noise over a bandwidth vastly wider than the signal. (iii) *The first-generation amplifier adds noise:* a flux-pumped JPA contributes $n_{\text{add}} \simeq 1.5$ quanta above the vacuum. None of these is a floor of nature. At 10 mK and 1.95 GHz the thermal occupation of the mode is $\bar{n} = [e^{h\nu/k_B T} - 1]^{-1} \simeq 8 \times 10^{-5}$: *the bath is quantum-empty*, and the entire noise budget is readout. Replacing the linear amplifier with a single-microwave-photon counter — demonstrated in the transmon architecture, dark-count limited rather than SQL-limited — removes the n_{add} penalty and the wideband integration penalty together; the R&D program's quantitative goal is the conservative $3\times$ SNR of Table 44, which, one month after the year-two readout insertion, turns the booster spectra into the bottom row of Figure 22: an 11σ peak bin, per tone, in the spin channel.

Shield design. The module needs 70 mT, stable and homogeneous, sitting beside a 9.4 T solenoid. The stray field at the side-cryostat position (~ 1.5 m off-axis) is of order 10 mT with strong gradients; the shield is therefore three-staged: (i) a passive three-layer mu-metal enclosure (attenuation $\sim 10^3$); (ii) a superconducting niobium can inside it, freezing the residual flux at cooldown; (iii) an active compensation coil servoed on a fluxgate/NMR probe pair, holding the bias at $\Delta B_e = \pm 0.1$ mT ($\pm 10^{-3}$ relative — a ± 2 MHz Kittel shift, tracked and absorbed by the hybrid's tuning servo). The same probe pair writes the field record used in the systematic table. Magnet quenches and ramps are the stress case: the superconducting can saturates gracefully and the module simply pauses — a duty-cycle entry, not a calibration risk.

Cross-couplings. The module sees the cavities only through the DAQ: microwave isolation between the booster lines and the sub-array lines exceeds 100 dB (separate cryostats, separate lines, circulator chains), so a common tone in all four streams cannot be an electronic artifact of one chain — which is precisely why the archived four-peak coincidence carried identification weight, and why the two-array exclusivity check inherits the same isolation.

Table 44: Booster coupling and noise budget (per hybrid; two hybrids per module). Anchored where possible to demonstrated magnon-hybrid performance (Crescini et al., 2020); entries marked *req.* are design requirements gated by the R&D year. The staging row is the quantitative content of the “quantum-limit goal” of Figure 22.

Parameter	Value	Basis / note
ferrimagnet	YIG spheres, $\varnothing 2 \text{ mm} \times 10$	demonstrated material class
spin density	$2.1 \times 10^{28} \text{ m}^{-3}$	YIG; $N_s \simeq 8.8 \times 10^{19}$ per sphere, 8.8×10^{20} per hybrid
bias field B_e	70.6 mT (ν_+), 68.6 mT (ν_-)	Kittel mode on each tone
bias homogeneity	$\Delta B/B < 10^{-4}$ over array	keeps inhomogeneous broadening subdominant
single-sphere coupling $g_{cm}/2\pi$	$\simeq 40 \text{ MHz}$	demonstrated class; array scales as $\sqrt{N_{\text{sph}}}$
magnon linewidth $\kappa_m/2\pi$	$\simeq 1.0 \text{ MHz}$	Gilbert damping + inhomogeneity
cavity linewidth $\kappa_c/2\pi$	$\simeq 20 \text{ kHz}$	$Q_L = 10^5$
cooperativity $g_{cm}^2/\kappa_c\kappa_m$	$\sim 10^5$	strongly hybridized (level repulsion resolved)
readout, first generation	JPA, $n_{\text{add}} \simeq 1.5$	week-one row of Fig. 22
readout, upgrade	photon counter, dark-count limited	<i>req.</i> : $3\times$ SNR (bottom row)
sensitivity, first generation	$\text{SNR}_B = 5\sqrt{t/5 \text{ d}}$ per tone	<i>req.</i> ; R&D-year exit criterion
sensitivity, upgraded	$3\times$ first generation	<i>req.</i> ; year-two insertion

F.4 The two-lock simulation: the channel-ratio panel and the four-peak week

[The channel-ratio panel of Part IX and the two-lock week-one figure of Part X as earlier revisions carried them, moved here unchanged except for heading levels.]

Earlier revisions gave this section a fourth verdict: the convergence of a deconvolved two-port power ratio $\hat{R} = P_B/P_\gamma$ to a locked $25/4$, simulated for an electromagnetic port and an electron-magnon port camped on the same condensate. The panel is withdrawn with its target (Part VI): the terminated $U(1)_B$ carries no channel below M_{lock} , the surviving spin couplings sit on the KSVZ-class sheet of Part VIII, and the electron-magnon port has no target at this f_K . What survives of the panel’s machinery is worth naming, because none of it was wasted: the two-port deconvolution and calibration chain of §F.9 below is generic instrument analysis, valid for any future two-channel measurement; and the statistical error separation it formalized,

$$\left. \frac{\sigma_{\hat{R}}}{R} \right|_{\text{stat}} = \sqrt{\frac{1}{\text{SNR}_B^2} + \frac{1}{\text{SNR}_\gamma^2}}, \quad \left. \frac{\sigma_{\hat{R}}}{R} \right|_{\text{sys}} = \sigma_{\text{cal}}, \quad (\text{F.6})$$

carries over unchanged to any ratio a later framework may lock. The discovery week of this paper rests on the three verdicts that remain — the spectra, the modulations, and the count test — and loses nothing it can currently defend.

Week one: the line — and the archived panels. Figure 22 is the archived two-lock simulation, kept for the record; the surviving content is its γ -port arithmetic, which the single-line plan of Part X

recomputed with the matched filter’s effective width: a camped sub-array saturates immediately — peak-bin SNR near 2.7×10^2 per week in the live plan (3.3×10^2 as this archived figure drew it; stat. only, the $\pm 24\%$ amplitude systematic scaling it), line shape fully resolved — and with both sub-arrays stacked on one rung the exposure doubles. The archived booster rows (middle and bottom) show what the retired configuration would have added — the visual record of the spin channel’s faintness is a readout property, not a floor of nature (§F.3). F-Kaxion-1a, as this configuration defined it, was to be settled inside the first week; it is retired with the doublet (Part VI), and the count test took its place.

F.5 Primer (archived): the withdrawn 1/35 split — and why it could not renormalize away

[Archived record. This primer explains the doublet of earlier revisions. The doublet is withdrawn: the registered mode analysis settles the count at one light angle (Part VI), so there is no split to explain. The primer is kept because it is the record of how the split was argued — the flux-cell overlap it teaches fed that reading and is not present in the present construction (§9.3, Appendix C.2) — and because retiring a fingerprint without its record would hide the change.]

The bulk hands the boundary *two* pseudoscalar directions, not one: the dual of the two-form and the holonomy of the Wilson line. Think of them as two clocks geared together, and the gearing ratio is not a coupling constant — it is the *geometric overlap* of one flux cell with the throat, which is one part in $K = 35$ by construction (Appendix F.7). Two geared oscillators of nearly common frequency always beat as a doublet, and the fractional splitting is the gearing: $\Delta\nu/\nu = 1/35 = 2.857\%$, i.e. a 56 MHz partner beside a 1.95 GHz line. As for quantum corrections: what renormalizes is dynamics — couplings, masses, kinetic terms’ overall scale; what cannot renormalize is topology — how many times one angle winds around another. The gearing is a winding datum, so radiative corrections enter only at relative order $(5/8)K^{-2} \simeq 5 \times 10^{-4}$, some fifty times below the splitting itself (Section 31). That was the primer as earlier revisions wrote it. The present reading is simpler and stricter: the second light field was never supplied — the colored holonomy is lifted at the circle’s curvature scale, and the terminated $U(1)_B$ eats nothing — so there is no partner to gear, no split to protect, and the partner search became the count test: a *missing* partner is the prediction, and a second line anywhere in the fence would refute the count (T-Kaxion-5, Part VI).

F.6 Why did earlier revisions expect twin peaks split by 1/35? (archived)

[Archived record. The doublet this essay explains is withdrawn: the registered mode analysis leaves exactly one light QCD-coupled angle (Part VI), and the essay’s premise — two light fields — is the premise the anomaly companion never supplied. It is kept as the record of what was withdrawn and of why it once seemed forced.]

Why does the Kaxion show twin peaks? Why do they split by exactly $\Delta\nu/\nu = 1/K = 1/35$? Are the two peaks identical in height (coupling) and in width?

The doublet is the *fine structure* of the Kaxion, and it exists because the forced light sector is *two* fields, not one. Dualizing the Green–Schwarz two-form gives one periodic pseudoscalar a_{GS} ; the holonomy of the gauge field around the phase circle S^1_θ gives a second, a_θ . Both are shift-protected zero-modes, and the shift symmetry that protects one protects the other, so both stay light (§9.3). The full low-energy spectrum is three states: two light zero-modes (a_{GS}, a_θ) and one heavy cap mode lifted to throat scale. The Kaxion is the QCD-massive light eigenstate; the doublet is the fine structure of the two light zero-modes seen through the QCD potential.

Why both are massive (and why a generic axion has no doublet). Each zero-mode couples to $G\tilde{G}$, but through its own geometric route — a_{GS} via the Bl-6 inflow, a_θ via the Wilson-line holonomy of the color subgroup of $\text{SU}(5)$. A throat reflection symmetry exchanging $a_{\text{GS}} \leftrightarrow a_\theta$ forbids the off-diagonal potential term, so the QCD-induced curvature is diagonal and equal,

$$V = \frac{1}{2}\mu^2 (a_{\text{GS}}^2 + a_\theta^2), \quad \mu^2 = \frac{\chi_{\text{QCD}}}{f_K^2}. \quad (\text{F.7})$$

Contrast the generic case: if both coupled to QCD through the same combination ($a_{\text{GS}} + a_\theta$), one would get *one* massive state and one exactly massless would-be Goldstone — a single peak, no doublet. It is precisely the exchange symmetry that promotes the massless partner to a second massive tone. No Peccei–Quinn construction has this.

Why the split is exactly $1/K$. With the potential diagonal, the only off-diagonal structure is kinetic. Reducing the $|H_3|^2$ term, the Chern–Simons improvement produces an overlap between the two zero-mode profiles,

$$\mathcal{L}_{\text{kin}} = \frac{1}{2} \left[(\partial a_{\text{GS}})^2 + (\partial a_\theta)^2 + 2\epsilon \partial a_{\text{GS}} \cdot \partial a_\theta \right], \quad \epsilon = \frac{1}{K} = \frac{1}{35}, \quad (\text{F.8})$$

where ϵ is the dimensionless kinetic gearing — the normalized single-transit overlap of the two profiles across a throat of depth K (computed geometrically in Appendix F.7). Because $V \propto \mathcal{K}$ and the metric is $\begin{pmatrix} 1 & \epsilon \\ \epsilon & 1 \end{pmatrix}$, the normal modes are the symmetric and antisymmetric combinations $u_\pm = (a_{\text{GS}} \pm a_\theta)/\sqrt{2}$, with kinetic coefficients $1 \pm \epsilon$. Canonical normalization ($\hat{a}_\pm = \sqrt{1 \pm \epsilon} u_\pm$) gives the physical masses

$$m^2(u_\pm) = \frac{\mu^2}{1 \pm \epsilon} \implies \nu_\mp = \frac{\nu_0}{\sqrt{1 \pm \epsilon}}, \quad \nu_0 = \frac{\mu}{2\pi}, \quad (\text{F.9})$$

and, to leading order in the small gearing,

$$\frac{\Delta\nu}{\nu_0} = (1 - \epsilon)^{-1/2} - (1 + \epsilon)^{-1/2} = \epsilon + \mathcal{O}(\epsilon^3) = \frac{1}{K} = \frac{1}{35} = 2.857\%, \quad (\text{F.10})$$

i.e. $\Delta\nu = 1.95 \text{ GHz}/35 \simeq 56 \text{ MHz}$, the denominator being the ungeared ν_0 and not either tone; the tone-to-tone ratio is the $\sqrt{18/17}$ of the earlier revisions’ tone-ratio relation. Two features make this a clean fingerprint. It is a pure frequency *ratio*, hence scale-free: it depends on ϵ alone, not on μ or f_K , so even if the absolute mass shifts within the cap uncertainty (§B.5.3), the split stays $1/35$

exactly. And it is an integer read-out: the throat depth $K = 35$ that also sets the TeV lock cluster and the whole hierarchy (§9.6).

The one assumption, honestly. The value $1/K$ is the *symmetric-limit* split. If the exchange symmetry is slightly broken (curvature asymmetry $\delta \equiv |\omega_1^2 - \omega_2^2|/\mu^2$), the split becomes $\Delta\nu/\nu = \sqrt{\epsilon^2 + \delta^2}$, so $1/K$ is the value when the routes are symmetric and a lower bound otherwise. The doublet’s *existence* is structural; its exact *ratio* $1/K$ rests on the equal-curvature assumption — the one model-level input behind the first fingerprint’s precise value, flagged as such in the ledger.

Are the peaks identical? Width, yes; height, no. The linewidth is not intrinsic; it is Doppler, set by the halo velocity dispersion, $\delta\nu/\nu \simeq \langle v^2 \rangle / c^2 \sim 10^{-6}$ ($Q_a \sim 10^6$, $\simeq 2$ kHz at 1.95 GHz). Both tones are excitations of the *same* cold condensate seen through the *same* halo, so both carry the identical boosted-Maxwellian line shape and the identical fractional width; the absolute widths differ only by the 2.857% that separates the centers. The hierarchy

$$\underbrace{\nu/Q_a \simeq 2 \text{ kHz}}_{\text{condensate line}} \ll \underbrace{\nu/Q_L \simeq 20 \text{ kHz}}_{\text{receiver bandwidth}} \ll \underbrace{\nu/35 \simeq 56 \text{ MHz}}_{\text{doublet split}} \quad (\text{F.11})$$

means the two tones are spectacularly resolved (separated by $\sim 3 \times 10^4$ linewidths) but a single-mode cavity tuned to one is *blind* to the other. The *height* is a different story. The frequency split is parameter-free, but the relative amplitude of the two tones is not: both peaks are bright (both DFSZ-coupled), but the observed power in each, $P_{\pm} \propto g_{\pm}^2 \rho_{\pm}$, carries the primordial energy partition between the two clocks at the lock/saddle crossing. The tone-power ratio κ is thus a free parameter, and the search of Part VIII is designed to be efficient even for κ far from unity. The identification test is “is there a partner at exactly 2.857%?”, not “are the two equal?”

The picture in one line. Two identical pendulums (equal QCD curvature, enforced by the exchange symmetry) geared by a weak spring (the kinetic overlap $\epsilon = 1/K$) have two normal modes — in-phase and out-of-phase — split by exactly the coupling strength; the coupling strength here is one over the throat depth. Both tones ride the same carrier, so their shape and width are shared; only how hard the Big Bang first plucked each clock is left free, so the two may ring at different volumes.

F.7 The flux-cell measure: $2\pi/K$ and the single-transit gearing $1/K$

[Archived record. The flux-cell measure and the single-transit overlap computed below fed the earlier two-field reading; on the intended class no scalar–scalar bilinear appears, so the overlap is not present in the present construction (§9.3), where it serves only as the most generous bound on a residual mixing (Appendix C.2). Their earlier application — a light doublet split by $1/K$ — is withdrawn with the doublet (Part VI): the second light field the split required was never supplied. The overlap was computed as geometry; the split it once predicted is gone.]

The compact-angle cylinder of the canonical metric is the rectangle $(r, \theta) \in [0, K] \times [0, 2\pi)$ with *unit* measure in both directions ($G_{rr} = G_{\theta\theta} = 1$): total cell area $2\pi K$, with no modulus to rescale it (GG-A2-GG6 T2). Two ratios follow by pure geometry. (i) *Winding normalization*. One full angular winding at fixed radial depth occupies the fraction $2\pi/(2\pi K) = 1/K$ of the cell; equivalently, the flux quantum spreads over K radial units, so any single-winding density is normalized by $2\pi/K$. (ii) *Single-transit gearing*. The tensor-shadow zero mode is uniform on the cell (it descends from the cohomology class, App. C.1); the holonomy echo lives on the angular fiber and samples the radial direction only through one transit of the throat. Their normalized overlap is therefore one radial unit in K :

$$\varepsilon = \frac{\langle \psi_{\text{GS}}, \psi_{\theta} \rangle_{\text{transit}}}{\|\psi_{\text{GS}}\| \|\psi_{\theta}\|} = \frac{1}{K} = \frac{1}{35}, \quad (\text{F.12})$$

the gearing the earlier reading used in the rotation of Eq. (9.4); the present construction does not (§9.3). The same counting with the L electromagnetic rungs of the throat gives the projection ratios of the series ($L/K = 1/7$ where a channel threads L of K units); we record it for cohort consistency, noting that after the exact anomaly matching of Part V no such projection enters the photon coupling.

The overlap integral, in full. Because the single-transit overlap fed the earlier reading and bounds any residual mixing (Appendix C.2), it is recorded in full rather than by the one-line counting above. The two light profiles on the cell $(r, \theta) \in [0, K] \times [0, 2\pi)$ are the tensor shadow, uniform because it descends from a cohomology class,

$$\psi_{\text{GS}}(r, \theta) = \frac{1}{\sqrt{2\pi K}}, \quad (\text{F.13})$$

and the holonomy echo, which is uniform on the angular fiber but localized to a single transit of the radial throat by the Wilson line that generates it,

$$\psi_{\theta}(r, \theta) = \frac{1}{\sqrt{2\pi}} u(r), \quad \int_0^K u(r)^2 dr = 1, \quad (\text{F.14})$$

with $u(r)$ the normalized single-transit density (unit-normalized because the holonomy threads the throat exactly once). The kinetic overlap that appears in the reduced Lagrangian is

$$\varepsilon = \langle \psi_{\text{GS}}, \psi_{\theta} \rangle = \int_0^{2\pi} d\theta \int_0^K dr \frac{1}{\sqrt{2\pi K}} \frac{1}{\sqrt{2\pi}} u(r) = \frac{1}{K} \int_0^K u(r) dr. \quad (\text{F.15})$$

For a transit that deposits its weight uniformly across the throat, $u(r) = 1/\sqrt{K}$ and $\int_0^K u dr = \sqrt{K}$, so $\varepsilon = \sqrt{K}/K = 1/\sqrt{K}$; for a transit sharply localized at unit depth, $\int_0^K u dr \rightarrow 1$ and $\varepsilon \rightarrow 1/K$. The physically correct case is the second: the Chern–Simons improvement term $B_2 \wedge d\theta \wedge (\dots)$ couples the shadow to the echo only through the *boundary* of a single radial cell, so the overlap samples one radial unit in K , giving

$$\varepsilon = \frac{1}{K} = \frac{1}{35} = 2.857\%, \quad (\text{F.16})$$

independent of the warp coefficients — a topological ratio, not a dynamical one.

From the overlap to the split, step by step. With the QCD potential diagonal and equal on the two clocks (the throat exchange symmetry, $V = \frac{1}{2}\mu^2(a_{\text{GS}}^2 + a_\theta^2)$, $\mu^2 = \chi_{\text{QCD}}/f_K^2$) and the kinetic metric $(\frac{1}{\varepsilon} \frac{\varepsilon}{1})$, the normal modes are the symmetric and antisymmetric combinations $u_\pm = (a_{\text{GS}} \pm a_\theta)/\sqrt{2}$, with kinetic coefficients $1 \pm \varepsilon$. Canonically normalizing $\hat{a}_\pm = \sqrt{1 \pm \varepsilon} u_\pm$ divides the common curvature μ^2 by $1 \pm \varepsilon$, so

$$m_\pm^2 = \frac{\mu^2}{1 \pm \varepsilon}, \quad \nu_\pm = \frac{\nu_0}{\sqrt{1 \pm \varepsilon}}, \quad \nu_0 = \frac{\mu}{2\pi}, \quad (\text{F.17})$$

and expanding to first order in the small gearing,

$$\frac{\Delta\nu}{\nu} = (1 - \varepsilon)^{-1/2} - (1 + \varepsilon)^{-1/2} = \varepsilon + \mathcal{O}(\varepsilon^3) = \frac{1}{K} = \frac{1}{35}. \quad (\text{F.18})$$

The result is a pure frequency *ratio*: it depends on ε alone, not on μ , f_K , or Ξ_{cap} , so within the withdrawn two-field reading the split stayed exactly $1/35$ wherever the absolute mass sat; higher transits enter at $\mathcal{O}(\varepsilon^2) = 1/K^2 \simeq 8 \times 10^{-4}$. That reading is withdrawn (Part VI); what the computation delivered was the overlap $\varepsilon = 1/K$ itself, which fed the earlier gearing and is not present in the present construction (§9.3); it survives only as the most generous bound on a residual mixing (Appendix C.2).

F.8 The archived four-peak template and its matched filter

[Archived record. The four-peak pattern below belongs to the withdrawn two-lock configuration — two tones, two channels. The count (T-Kaxion-5) leaves one line, and the surviving instrument reads it with the count test (Parts VI and X). The filter algebra is kept because matched-filter statistics transfer to the single-line template unchanged, and because the record should show what was withdrawn.]

Earlier revisions asserted that the Kaxion presents not one line but a structured *four-peak* pattern, and that a matched filter over that pattern is what converts a marginal single-line excess into a decisive identification. This course builds that template and its statistics from the two ingredients those revisions derived: the doublet spacing (Appendix F.7) and the two-channel ratio (Appendices C.4, D.2).

Step 1 — why four peaks, on the withdrawn reading. The light sector was taken to be *two* tones (the symmetric and antisymmetric eigenstates $\hat{a}_\pm$, split by $\Delta\nu/\nu = 1/35$), and each tone couples through *two* portals (electromagnetic $g_{K\gamma\gamma}$ and baryonic g_{KBB} , in the locked ratio 5:2). A two-channel receiver therefore records a 2×2 pattern — four peaks — whose frequencies and relative powers are fixed by the geometry:

$$\{\nu_-, \nu_+\} = \nu_K \{(1 + \varepsilon)^{-1/2}, (1 - \varepsilon)^{-1/2}\}, \quad \varepsilon = \frac{1}{35}; \quad \frac{P_\pm^B}{P_\pm^\gamma} = \left(\frac{g_{KBB}}{g_{K\gamma\gamma}} \right)^2 = \frac{25}{4}. \quad (\text{F.19})$$

The template has only two free amplitudes — the overall signal power and the primordial tone-power ratio $\kappa = P_+/P_-$ (the one number the geometry leaves open, Appendix F.6); everything else is locked. This is what makes the pattern, rather than any single peak, the observable.

Step 2 — the matched filter. Given a recorded spectrum $S(\nu) = \mathcal{T}(\nu; \vartheta) + n(\nu)$ with Gaussian noise of one-sided density $N_0 = k_B T_{\text{sys}}$ and the template family $\mathcal{T}(\nu; \vartheta)$ of (F.19) (parameters $\vartheta = \text{center}, \kappa, \text{amplitude}$), the optimal (Neyman–Pearson) statistic is the matched filter — the noise-weighted cross-correlation of data with template,

$$\Lambda = \frac{\int \mathcal{T}(\nu) S(\nu) d\nu}{\sqrt{\int \mathcal{T}(\nu)^2 d\nu}} \bigg/ \sqrt{N_0}, \quad \langle \Lambda \rangle_{\text{signal}} = \text{SNR}_{\text{tot}}, \quad (\text{F.20})$$

whose expectation on signal is the total signal-to-noise. Because the four peaks add *in quadrature* under the matched filter, the combined significance is

$$\text{SNR}_{\text{tot}}^2 = \sum_{c \in \{\gamma, B\}} \sum_{s \in \{+, -\}} \text{SNR}_{c,s}^2, \quad (\text{F.21})$$

so a network of four marginal 2σ peaks combines to $\sqrt{4} \times 2 = 4\sigma$ when the pattern is correct, and — crucially — collapses toward zero when it is not, because the filter is anchored to the locked spacing and ratio. The radiometer growth $\text{SNR} \propto \sqrt{t}$ (Appendix D.1) then applies to the combined statistic.

Step 3 — coincidence, and why the pattern is unforgeable. A random noise fluctuation can mimic one peak; mimicking the *locked* four-peak pattern requires four independent fluctuations at pre-declared frequencies and in a pre-declared power ratio. If each single-bin false-alarm rate is R_{FA} and the coincidence window is τ_w , the accidental four-fold (or, conservatively, two-fold across the doublet) rate is suppressed as a power of $R_{\text{FA}}\tau_w$,

$$R_{\text{acc}} = R_{\text{FA}}^2 \tau_w \quad (\text{doublet}), \quad R_{\text{acc}} \propto R_{\text{FA}}^4 \quad (\text{full pattern}), \quad (\text{F.22})$$

the estimate used in the archived analysis. The two identification observables were engineered to be astrophysics-free: the doublet spacing is a pure frequency *ratio* (no amplitude enters), and the channel ratio $\hat{R} \rightarrow 25/4$ cancels the local density, the velocity distribution, and the magnet field identically (Appendix D.2). A four-peak template match at the locked spacing and ratio would have been a reading of the gauge vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$ in a laboratory; the withdrawal removes the pattern, not the discipline, and the surviving count test inherits the same pre-registered, astrophysics-free character.

F.9 The archived two-port ratio estimator and dual-mode remark

[Archived record, moved here from the transfer-function course: the two-port estimator and the dual-mode remark belong to the withdrawn booster configuration.]

Two ports and the unbiased ratio (archived). The booster concept of Part VIII recorded two spectra S_B, S_γ through two independent chains. Define the deconvolved estimator

$$\hat{R} = \frac{\int W(\nu) [S_B(\nu) - \hat{S}_{\text{noise}}^B] / \mathcal{G}_B |\chi_B|^2 d\nu}{\int W(\nu) [S_\gamma(\nu) - \hat{S}_{\text{noise}}^\gamma] / \mathcal{G}_\gamma |\chi_\gamma|^2 d\nu}, \quad (\text{F.23})$$

with W any common spectral window covering the line. Because the condensate factor $\rho_{\text{loc}} \mathcal{T}(\nu)$ is common to both numerator and denominator, every astrophysical unknown — the local density, the velocity distribution, the tone ratio κ , even the day's halo weather — cancels identically, leaving

$$\langle \hat{R} \rangle = \left(\frac{g_{KBB}}{g_{K\gamma\gamma}} \right)^2 = \frac{25}{4}, \quad \sigma_{\hat{R}} \simeq \frac{25}{4} \sqrt{2} \left(\text{SNR}_B^{-1} \oplus \text{SNR}_\gamma^{-1} \oplus \sigma_{\text{cal}} \right), \quad (\text{F.24})$$

the error budget being statistical noise in each port plus the calibration-chain systematic σ_{cal} (the prefactor $\sqrt{2}$ is deliberate conservatism over the pure statistical limit, which is $1/\text{SNR}_B$ when $\text{SNR}_\gamma \gg \text{SNR}_B$) — the quantity that set the $\pm 20\% \rightarrow \pm 5\%$ staging of the retired falsifier F-Kaxion-1b. The archived power-ratio lock is (F.24) read at infinite signal-to-noise.

Dual-mode integrity (archived). For the coincidence resonator of the earlier configuration, the only new question is whether two modes spaced by $\Delta\nu/\nu = 1/35$ interfere with one another. They do not, parametrically: at $Q_L = 10^5$ each mode's width is $\nu/Q_L \simeq 20$ kHz against a spacing of 56 MHz — a separation of 2.8×10^3 linewidths — so cross-talk and mode-pulling corrections enter at $(2Q_L \Delta\nu/\nu)^{-2} \sim 10^{-8}$, far below calibration floor. The pair behaves as two independent single-port receivers that happen to share a condensate, which is exactly the statistical assumption behind the coincidence rate $R_{\text{acc}} = R_{\text{FA}}^2 \tau_w$ of the archived analysis.

F.10 The archived inversion step of the baryonic-ledger course

[Archived record, moved here from the baryonic-ledger course. The inversion was a mistake: dualization preserves the coefficient vector, so the ledger ratio is 2:5 for source and response alike, and no reciprocal arises. The step is kept as the record of how the 5:2 readout entered.]

Step 4, as earlier revisions printed it — source becomes response: why 2:5 was read as 5:2. At the level of the two-form *source* couplings (C.15), baryon number stands to electromagnetism as 2 : 5: the two-form is sourced more strongly by the (rank-five) hypercharge sector than by the (rank-two-effective) baryon sector. But the observable is the coupling of the *dual* field — the propagating pseudoscalar a_K — not of B_2 . Dualizing $H_3 = dB_2 \rightarrow \star da_K$ exchanges the roles of equation of motion and Bianchi identity (Appendix C.1), and under that exchange each sector's *response* coupling is the reciprocal of its *source* normalization — the standard p -form/dual swap. Hence the ratio the detector reads is inverted,

$$\frac{g_{KBB}}{g_{K\gamma\gamma}} = \left(\frac{2}{5} \right)^{-1} = \frac{5}{2}, \quad (\text{F.25})$$

and a matched receiver, which reads power $\propto g^2$, sees

$$\boxed{\frac{P_B}{P_\gamma} = \left(\frac{g_{KBB}}{g_{K\gamma\gamma}}\right)^2 = \frac{25}{4} = 6.25}. \quad (\text{F.26})$$

Baryon number, written faintly into the ledger as 2, is read loudly by the Kaxion as 5: it was, on that reading, the *stronger* of the two portals, which is what made the two-channel ratio of earlier revisions a decisive, astrophysics-free test — until the count (T-Kaxion-5) removed the light vector it needed (Part VI).

The same inversion, as the on-ramp essay printed it. This was the step that made *baryon* the stronger channel, and it deserves printing in four lines. (i) The anomaly factorization fixes the two-form’s *source* couplings by the vector: at source level baryon stands to electromagnetism as 2 : 5. (ii) Dualizing $H_3 \rightarrow da_K$ exchanges the roles of field and source (Appendix C.1), so the pseudoscalar’s *response* coupling to each sector inverts that sector’s source normalization — the standard p -form/dual swap. (iii) Hence

$$\frac{g_{KBB}}{g_{K\gamma\gamma}} = \frac{5}{2}. \quad (\text{F.27})$$

(iv) A matched receiver reads power — the square — so

$$\frac{P_B}{P_\gamma} = \left(\frac{g_{KBB}}{g_{K\gamma\gamma}}\right)^2 = \frac{25}{4} = 6.25. \quad (\text{F.28})$$

The integer pair is $(L, L - N) = (5, 2)$. The same 2/5 is locked by the four-current identity of GG-A7-Qi4 into the informational current, which is why the same number reappears as a bioenergetic conversion constant elsewhere in the cohort: one ratio in two costumes.

G Notation

Table 45: The notation of this paper at a glance, in two columns of symbol and meaning. The table opens the Notation appendix and is meant to be read beside the paper rather than before it: it gives each symbol its definition and, where the definition belongs to a companion, the theorem of that companion which fixes it. Symbols introduced and discharged inside a single proof are left out, and the glossary later in this appendix carries the fuller treatment for the ones that recur.

$\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S_\theta^1$	bulk manifold (GG-A2-GG6 T1)	$(M, N, L; J, K)$	OGN-5 = (1, 3, 5; 7, 35)
$r \in [0, K]$	scale coordinate, $r = \ln(E_{\text{GUT}}/\mu)$	$\theta_{\text{GS}}, \theta_W$	tensor shadow; holonomy echo
B_2, H_3	Green–Schwarz tensor, field strength	a_K, a_H	light/heavy eigenstates; β mixing
X_4	BI-6 ledger four-form	f_K, m_K	decay constant; mass
GDOS, TGG, CCI	boundary law; total stress; seam law	θ_i	misalignment angle $\simeq 1.1$ rad
GRIP, GFF, GEP, CEP	runtime; free energy and its parts	χ_{QCD}	topological susceptibility
DGI, GSSB, GA, GGA	cascade: top, saddles, attractors; GGA the unique deepest (ground state)	$g_{K\gamma\gamma}$	photon coupling; the baryonic ledger entry opens no light channel
GRIP-In, GRIP-Out	runtime regimes: fixed-landscape descent; landscape reorganization		
$M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K}$	lock scale 3.85 TeV	Ξ_{cap}	cap index (bench. 43)
$\Omega_{\text{DM}} = 1/\pi - 1/2\pi^2$	π -locked target	ν_K	the unique Kaxion window, Eq. 36.1




Figure 22: Archived: the retired two-lock configuration's simulated week, kept for the record (the doublet and ratio it draws are withdrawn — Part VI; the γ -port arithmetic was recomputed for the single-line plan with the matched filter's effective width, 2.7×10^2 against the 3.3×10^2 drawn here). *Top:* the morphology as the earlier revision drew it — the single line at ν_K beside the retired pair, and the three widths that govern what a receiver sees. *Below:* the four peaks (fixed-seed simulation, KAXIO parameters of Table 29). Left column is the lower tone ν_- , right column the higher tone ν_+ , drawn at a representative amplitude ratio $\kappa \simeq 0.7$ — the tone partition is not parameter-free (Appendix F.6), so the per-tone design values quoted below are those of the reference tone ν_+ . *Top row:* the electromagnetic ports after one week — the reference tone ν_+ camps at a peak-bin SNR $\sim 3.3 \times 10^2$ (stat. only; $\pm 24\%$ syst.), ν_- at κ times that; shapes fully resolved, frequency ratio read at 10^{-4} . *Middle row:* the harmonic-booster ports after the same week with first-generation readout (SNR_{line} $\simeq 5.9$ for ν_+): faint, because the per-spin coupling is minute and the first-generation chain adds ~ 1.5 quanta of amplifier noise over a magnon linewidth $500\times$ the signal line (§F.3) — yet already arriving in four-fold coincidence at two locked frequencies. *Bottom row:* the same booster ports one month after the year-two insertion of the quantum-limited readout that the R&D program targets (photon counting; $3\times$ SNR, Table 44): SNR_{line} $\simeq 37$ for ν_+ , peak bin $\simeq 11\sigma$ — the noise floor is not fundamental, since at 10 mK the thermal magnon occupation is $\bar{n} \sim 8 \times 10^{-5}$: the bath is quantum-empty, and the floor is set by the readout alone. Statistical noise $\sigma = 1/\text{bin}$ by construction throughout.

Part XIV

Common Objections and Responses

The strongest objections a theorist, a cosmologist, a haloscope experimentalist, or a philosopher of physics is most likely to raise on first reading — answered, with pointers into the paper

A prediction this sharp invites hard objections, and it should. This Part collects the sixty-nine strongest — those a theorist would raise against the framework, those a cosmologist would raise against the relic pipeline, those an experimentalist would raise against the detection, and, closing the ledger, those a philosopher of physics would raise against the enterprise itself — grouped by theme, each with an extended response and a pointer to where in the paper the full treatment lives. They are not strawmen; they are the questions the author has been asked, sharpened. Where the honest answer is “not yet,” we say so, and point to the named open items of Part XI that record exactly which claims are theorems, which are derived, and which rest on the named conjectures of the deck or on the one declared calibration. The final section deserves a word of motivation: a theory that claims to derive the constants of nature from axioms is making moves in metaphysics whether it acknowledges them or not, and it owes its readers the same explicitness there that it demands of itself in the laboratory. After all, the test of a framework is not that it deflects every question, but that it knows which questions it has answered and which it has not.

H Methodology and the derivation chain

Objection 1 (O-Kaxion-1) (the axion-model historian)

“You added an axion by hand, like every axion model since 1977.”

Response. No field was added, and the point is structural rather than rhetorical. For the GG-6 fermion content the reducible anomaly cancels if and only if a single Green–Schwarz two-form B_2 is present (the rank-one criterion of GG-A3-BI6), so deleting the Kaxion deletes anomaly freedom and with it the consistency of the quantum theory. Every installed axion model *chooses* to add a field to solve a problem; here the field is a theorem of consistency, present before the dark-matter question is asked. The burden of proof is correspondingly reversed: it is the *absence* of the Kaxion, not its presence, that would require new physics — the honest comparison (Part IV) is a forced sector against an installed one.

Objection 2 (O-Kaxion-2) (the numerology skeptic)

“It is numerology — integers fitted after the fact.”

Response. The charge would be fair if the integers were fitted; they are not. The ledger (1, 3, 5; 7, 35) is fixed in GG-A1-Constitution–GG-A3-BI6 by anomaly freedom and the Diophantine

closure, before any dark-matter observable is computed, and it is the *same* ledger that fixes the Standard-Model and cosmological papers. Numerology has the freedom to choose a new combination for each new datum; this ledger has none, and every Kaxion observable is a gear ratio of those five integers, pre-registered. The test of numerology is whether it can be wrong — a line outside the window retires the calibrated placement and the ladder that fences it, a second line inside it refutes the count and names a registered mode, so this can.

Objection 3 (O-Kaxion-3) (the precision auditor)

“Your Standard-Model papers claimed per-mille accuracy; a 30%-wide window here is backsliding.”

Response. A different epistemic class, declared as such rather than smuggled. The results of GG-A9-SM are algebraic ledger locks and carry per-mille accuracy honestly; the Kaxion’s absolute placement, by contrast, rests on a calibrated normalization (O14) whose geometric residual is a discrete ladder of loadings, and the paper prints the window that ladder spans rather than a percentage it cannot derive. At fixed f_K the line carries $\pm 1.3\%$ from the lattice. Crucially, the dimensionless fingerprints that do the identifying are algebraic and carry *no* error at all; the window lives only in the placement.

Objection 4 (O-Kaxion-4) (the neatness skeptic)

“ $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ is suspiciously neat.”

Response. Neatness is not the test; falsifiability is. The value 0.2676 is derived in GG-A10-Cosmos from the boundary closure of the cosmic inventory, not chosen to match data, and it agrees with the measured dark-matter density at the percent level — a prediction with an error bar. A number that is both simple and falsifiable is the signature of a compression, not of fitting; the way to discredit it is to measure Ω_{DM} away from 0.268, which the data do not.

Objection 5 (O-Kaxion-5) (the fudge skeptic)

“The normalization has now been corrected twice — first GG-A10-Cosmos’s, then your own — which looks like fudging to the answer.”

Response. Both corrections were forced, and neither moved the line. The first replaced an engineering normalization of GG-A10-Cosmos by the instantaneous-onset estimate; the second, here, replaced that estimate by a direct integration of the equation of motion, which is larger by 1.49 (Part V). Under the calibration convention this paper now states, neither correction touches the frequency: the benchmark is the row at which a generic angle saturates the inventory, and what moved was the angle, from 1.4 to 1.1 rad — both generic. What would have been fudging is holding the angle fixed and letting the normalization move the line; the paper says in print which quantity absorbs a correction and why.

Objection 6 (O-Kaxion-6) (the self-citation skeptic)

“Twenty-two unpublished papers citing one another is not evidence; it is a closed loop.”

Response. Correct — and the paper is built so that nothing depends on the loop being believed. The cohort citations are *provenance*, not authority: they say where a theorem is proved, not

why the reader should accept it, and Part II re-derives from scratch every ingredient the Kaxion actually uses, so this paper is checkable standing alone. Decisively, the exit from the loop is not textual at all: the ledger of falsifiers (Appendix E.2) hands the verdict to instruments that owe the cohort nothing. Twenty-two papers agreeing with each other prove nothing; a haloscope finding one line inside the window and none beside it, or failing to, proves everything that matters here. The loop closes on hardware, which is the only closure physics has ever accepted.

Objection 7 (O-Kaxion-7) (the certification skeptic)

“You call things ‘theorems’ that nobody outside this program has checked.”

Response. ‘Theorem’ in this paper is a claim about deductive structure — stated premises, explicit derivation, checkable steps — not a claim about social certification. The master ledger (Table 13) grades every claim precisely so the reader can check it: which rows are proved, which are derived through Standard-Model inputs, which rest on the named conjectures or on the one declared calibration. Outside checking is welcome; but the paper’s decisive review was designed to be harsher than any reader’s: four pre-registered kill conditions, executed by experimental collaborations that owe this paper nothing (§36.9). A reader can reject an argument; only a haloscope can reject this theory.

Objection 8 (O-Kaxion-8) (the methods conservative)

“Pre-registration is a social-science ritual; physics has never needed it.”

Response. Physics needed it in 1974 and did not have the word: the J/ψ was persuasive overnight precisely because the charm quark’s mass window had been staked in advance. What pre-registration buys here is not rhetoric but arithmetic: the look-elsewhere penalty shrinks from the $2.2\times$ in integration time a blind gigahertz scan pays (local 7.4σ for a global 5σ) to $2.0\times$ for the declared window (local 7.0σ) and $1.8\times$ on the benchmark rung (local 6.7σ) — no one-bin collapse is claimed (§37.1) — and the February-2026 episode — a 5.1σ local excess deflated to 3.5σ global for want of a prior stake (§37.4) — is the live demonstration. In a field whose history is a graveyard of floating anomalies, writing the number down first is not ritual; it is the difference between an anomaly and an answer.

Objection 9 (O-Kaxion-9) (the combinatorics skeptic)

“With five integers and freedom to combine them, some formula was always going to land at an accessible frequency.”

Response. Count the freedom honestly. The integers were not chosen for this paper: $(1, 3, 5; 7, 35)$ is the terminal output of the OGN-5 cascade, fixed in GG-A2-GG6 and spent — before any dark-matter question — on α^{-1} , $\sin^2 \theta_W$, the lock scale, and the electroweak scale (Table 2). The Kaxion re-uses that frozen ledger through one new geometric object, the cap index, whose admissible values form a seven-element set (Table 11) — and the entire set spans barely a factor 1.3 in frequency. The reachable target space is not ‘somewhere on the dial’; it is 0.42 of an octave, published, with a marked row. A combinatoric search that could land anywhere proves nothing by landing; a cascade that can only land in 0.42 of an octave is refuted by silence there — which is exactly what F-Kaxion-1c schedules.

Objection 10 (O-Kaxion-10) (the cynic)**“How convenient: the window sits exactly where no experiment has yet reached.”**

Response. Invert the chronology. The window is where it is because $\sqrt{\chi_{\text{QCD}}}/f_K$ puts it there; that the field’s at-or-near-DFSZ frontier currently sits at $5.4\,\mu\text{eV}$ — the frontier at DFSZ depth proper is IBS-CAPP’s $4.91\,\mu\text{eV}$, and the gap below the window is therefore wider still — just below it, is a fact about the field’s scan history, not about the theory’s choices — and it is a *temporary* fact with a printed expiry: the scan fronts enter the window from below within roughly two years, the upper edge closes by $\sim 2030\text{--}2032$ (Table 16), and the camping arithmetic of §37 makes the benchmark testable in *minutes* by instruments that already exist. A convenient prediction hides beyond reach; this one has published the dates on which its convenience expires.

Objection 11 (O-Kaxion-11) (the ordering skeptic)**“Why should anyone read C1 before the eleven parent papers are digested and accepted?”**

Response. Because the dependency runs through hardware, not through belief. Part II states, self-contained, the three structures the Kaxion needs; the reader may treat them as this paper’s axioms and judge everything downstream on its own coherence. More pragmatically: the experimental program is worth executing under any attitude to the parents — the count-test re-binning costs nothing and tests a sharp hypothesis; a camped week at 1.95 GHz is good physics as a DFSZ probe regardless of provenance; and a DFSZ-depth line standing alone in the fenced window would force the parents onto everyone’s desk overnight, while a swept-and-silent window closes the case without anyone reading them at all (§3). The reading order resolves itself either way.

Objection 12 (O-Kaxion-12) (the tone critic)**“The ‘Search Card’ and the ‘bets’ rhetoric read like marketing, not physics.”**

Response. The register is deliberate, and it is load-bearing. A prediction that wants to be executed by busy collaborations must be *transmissible*: the Search Card exists because the alternative — asking an experimentalist to extract four falsifiers, a line shape, and a protocol from two hundred pages — is how sharp predictions die of friction. As for the wager language: it is the falsification postulate spoken plainly. Every rhetorical flourish in this paper is attached to a number, a date, or a kill criterion; strip the prose and the ledger remains (Appendix E.2). We would rather be accused of salesmanship for a refutable product than of scholarly reticence about an irrefutable one.

I The geometry and its protection**Objection 13 (O-Kaxion-13) (the extra-dimension skeptic)****“A sixth dimension is excluded by collider and gravity experiments.”**

Response. Those bounds are real, but they constrain *spatial* extra dimensions with gravitating Kaluza–Klein towers, and that is not what GG-6 has. The two extra coordinates are non-spatial:

I_r is the renormalization-scale direction and S_θ^1 a dimensionless phase circle (GG-A2-GG6 T2), neither carrying a spatial KK graviton. The only accessible tower is the radial lock cluster at $M_{\text{lock}} \approx 3.85 \text{ TeV}$ — a collider *prediction* (P1-KK), not a liability — while the angular tower sits unobservably at E_{GUT} . The objection mistakes a non-spatial geometry for a large spatial one.

Objection 14 (O-Kaxion-14) (the two-clocks skeptic)

“Two compact-angle directions means two dark-matter species — are you double counting?”

Response. One light field. The circle’s holonomies are lifted at the circle’s curvature scale (Section 9.3) and are not dark matter; the tensor remnant is the single condensate. The invariants established in Part III — a single $m_K f_K$, a single energy density, a single angular amplitude — are the explicit proof that nothing is counted twice, and a detector sees one population. What is not settled by theorem is whether the parents’ five registered exposures (nine modes) are all heavy (Section 9.4); that is the count, and the count test is what decides it.

Objection 15 (O-Kaxion-15) (the renormalization skeptic)

“Kinetic mixings renormalize — does anything in the count depend on a geometric mixing surviving quantum corrections?”

Response. No. Earlier revisions leaned on a kinetic overlap between the tensor remnant and the holonomy (“the gearing $1/K$ ”); on the intended class no such bilinear appears in the quadratic action, and the count does not need it (Section 9.3). The count rests on the lift of the holonomies and on the single inflow, neither of which is a coupling constant with a renormalization-group flow; the overlap integral is recorded in Appendix F with the doublet it once fed.

Objection 16 (O-Kaxion-16) (the quantum-gravity skeptic)

“Quantum gravity violates global symmetries; your shift symmetry will not survive Planck-scale physics.”

Response. There is no global symmetry to violate. The Kaxion’s shift symmetry is not the approximate global $U(1)_{\text{PQ}}$ of a scalar potential but the exact gauge invariance of the two-form B_2 together with the large gauge invariance of the Wilson loop, and gravitational operators cannot break gauge redundancies. The only breaking is the topological one the theory itself dictates (the gluonic portal). This is the standard quality argument for the model-independent string axion, here holding for a field the theory could not have omitted — the quality problem is structurally absent, not merely small.

Objection 17 (O-Kaxion-17) (the ontological literalist)

“‘Non-spatial dimensions’ is a word game: either they are dimensions or they are not.”

Response. They are dimensions in the only sense differential geometry requires — coordinates on which fields depend, with a metric and connection — while carrying labels that are not positions: one is a renormalization-group scale, the other a phase. The literalist demand that ‘dimension’ mean ‘place you could walk’ was already obsolete in ordinary quantum mechanics,

whose wavefunctions live on configuration space. What matters physically is not the label but the consequences: the scale coordinate produces the exponential hierarchies (Table 36 traces the usage), the phase coordinate produces winding, holonomy, and hence the Kaxion itself (Theorem 8.2). Call them what one likes; the single line inside the window and the count are their signature, and both are measurable.

Objection 18 (O-Kaxion-18) (the hierarchy skeptic)

“Towers at 10^{16} GeV, locks at TeV, zero-modes at μeV — thirty-one orders of magnitude from ‘one geometry’ is fine-tuning by exponential.”

Response. An exponential is only a tune if its argument is chosen; here the arguments are the ledger integers, and two of the three scales are two *readings* of the same throat — $e^{-2\pi}$ per unit depth and depth $K = 35$ set the lock scale, and the residual $\bar{\theta}$ carries $e^{-K}/2^J$ (Section 9.6) — while the third, the μeV mass, is no exponential at all but QCD’s susceptibility over a decay constant calibrated once. The hierarchy is not a set of knobs but one integer exponentiated and one scale declared — which is precisely how nature already behaves elsewhere: the proton-to-Planck hierarchy is QCD’s own dimensional transmutation, an exponential of a fixed coupling. The skeptic’s real question is whether the *same* K shows up at all three depths. It does, testably: the TeV lock cluster (GG-P1-KK) reads $K = 35$ aloud, and the μeV window’s ladder reads $K + \Sigma$ under the ansatz; either experiment can convict the geometry of inconsistency.

Objection 19 (O-Kaxion-19) (the string theorist)

“The Green–Schwarz mechanism belongs to ten-dimensional string theory; you cannot borrow it piecemeal in six dimensions.”

Response. Green–Schwarz is not proprietary to strings: it is the general statement that a reducible anomaly polynomial can be canceled by a two-form with a modified Bianchi identity, and its consistency conditions are checkable in any dimension. GG-A3-B16 does exactly that check in six dimensions — the factorization, the coefficient vector, and the trace ratio $2/5$ are derived, not imported (Section 6.3). If anything, the borrowing runs the other way: the string-theoretic axion literature is cited (Section 27) as the precedent that two-form duals generically yield axion-like fields; what the present framework adds is that the two-form is *forced*, its coefficients are fixed, and the resulting pseudoscalar therefore has no dial. The mechanism’s pedigree is welcome; its numbers here are its own.

Objection 20 (O-Kaxion-20) (the counting skeptic)

“Why exactly two pseudoscalar directions on the boundary — not one, not three?”

Response. Because the bulk offers the boundary exactly these angles: the dual of the three-form flux through the throat (the tensor’s shadow) and the holonomies of the rank-six group around the compact phase circle — six of them, one per Cartan direction — and Section 9 builds them. A different count would require a different topology or a different gauge group, and $\mathcal{M}_4 \times I_r \times S_\theta^1$ with the forced group supplies neither; the tensor cannot be absent, because the anomaly forbids it. What became of the holonomies is the count’s content (Part VI): they are expected to lift at the scale the coset zero mode registers, none of the six computed, the

terminated $U(1)_B$ eats nothing, and the tensor’s dual is the one angle that descends light. The count of light lines is one, conditionally on the parents’ registered exposures; a doublet at the window would refute it and name which registered mode is light.

Objection 21 (O-Kaxion-21) (the discrete-dial skeptic)

“A cap index chosen from seven discrete options is just a smaller dial.”

Response. Three replies, in ascending strength. First, a seven-value discrete set spanning a factor 1.3 is different in kind from a continuous parameter spanning ten orders: the whole set is falsifiable at once, on a schedule (Table 31), which no dial-equipped model can offer. Second, the set is not flat: $\Xi_{\text{cap}} = K+N+L = 43$ carries a stated selection argument (the cap receives the UV integers, the boundary keeps J), registered with its status honestly marked as the paper’s sharpest open item (Section 26.3). Third, the discreteness is itself testable, at the grade it carries: the ladder is an ansatz (Section 26.3), and at the ± 25 MHz the lattice leaves on the line the rungs are resolved, so a line confirmed *between* rungs — at, say, 1.87 GHz — would retire the ladder while leaving the Kaxion, its address and its count standing. Dials cannot be caught between positions; ladders can, and this one has said in advance what being caught would cost.

J Cosmology and the relic pipeline

Objection 22 (O-Kaxion-22) (the initial-conditions skeptic)

“The misalignment angle is a tuned initial condition.”

Response. Three replies, in increasing force. It is a dial, and the paper says so: the bulk normalization is calibrated so that the benchmark row saturates the inventory at a generic angle, and the abundance then returns $\theta_i \simeq 1.1$ rad — the generic $\mathcal{O}(1)$ value every standard axion-dark-matter scenario assumes, no more tuned than that. What is not a dial is what follows: the ladder, the count, and the DFSZ address are fixed by the integers, and a confirmed line whose absolute rate implied a tuned angle would strain the closure visibly. And the alternative history is not hidden but analyzed within the series (Part IV): the post-inflationary randomized-angle history moves the line to $\gtrsim 6$ GHz, outside the window, and is excluded within the series only conditionally — on the inflation-era pinning the paper names as Conjecture C-Kaxion-d.

Objection 23 (O-Kaxion-23) (the rival-candidates advocate)

“The rivals’ exclusions are conditional on an unproven framework.”

Response. Conceded plainly: the exclusions of Part IV are structural exclusions *within* GG-Theory, each resting on a constitutional clause, not model-independent experimental bounds — and we have framed them that way throughout. But the framework does not hide behind the caveat. It is falsifiable, and sharply — a Kaxion null across the window retires the calibrated placement and its ladder, not the framework, and sends the search to $\gtrsim 6$ GHz; a compact-angle line elsewhere retires the same objects — and it offers the converse hostages: a confirmed

thermal WIMP, a Peccei–Quinn axion off the closed window, a keV sterile-neutrino dark matter, or a primordial-black-hole-dominated halo would each refute the dark-sector closure. The theory stakes a falsifiable claim against the rest of the field, and that claim is testable from both sides.

Objection 24 (O-Kaxion-24) (the inflation historian)

“The two-epoch descent, with a low-curvature lock crossing setting the angle, is exotic.”

Response. The primary inflation is standard and GUT-adjacent; what is unusual is the second, low-curvature lock crossing of the two-basin descent, at which the compact angle is released — and it is GG-A10-Cosmos’s, and it is load-bearing in the favorable direction. It is what makes a single, patch-independent θ_i meaningful, *provided* the angle was pinned through the primary inflation that precedes it — the conditional this paper names as C-Kaxion-d (Section 13.5); the lock crossing itself is prompt and imprints no isocurvature at all. The unusual epoch is not a liability bolted on; it is where the release happens, and the pinning is the one thing about it that remains to be computed.

Objection 25 (O-Kaxion-25) (the fine-tuning cosmologist)

“An initial angle of 1.1 radians that happens to yield today’s abundance is anthropic selection wearing a derivation’s clothes.”

Response. The direction of inference matters. Anthropic reasoning *assumes* the abundance and *excuses* the angle; this paper takes the abundance target independently ($\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$, derived-conditional in GG-A10-Cosmos under its closed-packet conjecture, tested against the CMB to 1.5%), calibrates the benchmark so that a generic angle saturates it, and reads the angle back (Section 15.3). The check is weaker than a derivation and we say so: had the benchmark required $\theta_i = 0.03$ or 3.10 — the tuned edges — the calibration would have been suspect and we would have said so; it returned 1.1, a generic value. The right description is not that the angle was selected to fit us, but that the geometry left exactly one number to the universe and the universe filled it unremarkably.

Objection 26 (O-Kaxion-26) (the two-epoch skeptic)

“Your isocurvature safety hangs on a low lock-crossing scale H_2 ; shift it and the story collapses.”

Response. The objection aims at the wrong epoch, and earlier revisions of this paper did too. The lock crossing is prompt ($N_2 = 0$) and imprints no isocurvature whatever its H_2 ; the isocurvature budget is decided by the primary inflation, whose $H_1 \simeq 1.5 \times 10^{13}$ GeV (from $r = 0.0038$ and $A_s = 2.1 \times 10^{-9}$, as GG-A10-Cosmos now prints) sits six orders *above* the ceiling $H \lesssim 2.3 \times 10^7$ GeV that a light angle could tolerate (Section 13.5). The safety therefore hangs not on H_2 but on the angle having been heavy during primary inflation — Conjecture C-Kaxion-d — and the paper says so rather than quoting headroom on a scale that does not enter. Fragility is when a conclusion sits near its cliff; this one sits on a named conjecture, which is a different thing, and the conjecture has a named killer.

Objection 27 (O-Kaxion-27) (the substructure specialist)

“Post-inflationary axion scenarios predict miniclusters and streams; where is your small-scale story?”

Response. The Kaxion’s history is single-patch conditionally — no Peccei–Quinn scalar exists to restore thermally, reheating sits a factor 12 below f_K , and de Sitter randomization during primary inflation is foreclosed by the inflation-era pinning of Conjecture C-Kaxion-d (Section 13.5) — so the violent minicluster spectrum of post-inflationary scenarios is absent under the conjecture, not by theorem. What survives is the ordinary CDM hierarchy down to the Jeans scale $\lambda_J \sim 3 \times 10^{-4}$ pc, plus the standard tidal-stream phenomenology any cold particle inherits; Part IV’s fences state the observational consequences (and the contingency section addresses what a detected minicluster population would mean: trouble). We register the asymmetry plainly: small-scale smoothness at the minicluster scale is a *prediction* here, and a confirmed Kaxion minicluster spectrum would falsify the single-patch clause of the framework.

Objection 28 (O-Kaxion-28) (the survey cosmologist)

“DESI and its successors increasingly favor evolving dark energy; your π -locked static inventory is on the wrong side of the data.”

Response. The π -locks fix the *present* inventory ($\Omega_m = 1/\pi$, tested to 1.5%) and the framework’s attractor requires $w = -1$ exactly — a two-sided, pre-registered stake that the cohort’s own kill-test ledger carries (the dark-energy falsifier of GG-A9-SM’s timeline). If the evolving- w hints harden into a converged detection, the GRIP attractor is refuted and this paper’s abundance target loses its derived-conditional status: we say so without hedging, because that is what sharing a ledger means. But the present evidential state is hints with shifting significance across analysis choices, against a locked prediction that has survived every published inventory to date. The right posture is the one the falsification postulate assigns: watch the surveys, keep the stake on the table, and let the convergence — either way — be the verdict.

Objection 29 (O-Kaxion-29) (the multi-component advocate)

“Why should dark matter be exactly one species? Multi-component dark sectors are generic in every UV framework.”

Response. Generic in frameworks with room for them. Here the constitution’s anti-tuning clause forbids sectors the consistency conditions do not force, and the consistency conditions force exactly one dark candidate (Part II) — the same parsimony that deletes the WIMP deletes the second component. This is not aesthetics; it is exposed: the misalignment-completeness conjecture (the Kaxion saturates the whole π -locked share) is named, flagged, and carried as one of the named fences of the deck (Tables 13 and 41), and it fails visibly if the laboratory line is confirmed but the implied local density falls short of ρ_{loc} . Multi-component advocates and this paper thus disagree about a measurable ratio, not a taste — and KAXIO’s absolute-rate calibration is the instrument that measures it.

Objection 30 (O-Kaxion-30) (the concordance defender)

“ Λ CDM already fits the CMB and structure perfectly with a featureless CDM fluid; the Kaxion adds nothing cosmology needs.”

Response. Exactly — and that is the design. On every scale where Λ CDM is tested, the Kaxion *is* Λ CDM’s cold fluid (Part III derives the credentials rather than assuming them), so cosmology’s successes are inherited, not threatened. What the Kaxion adds is everything the fluid description must leave blank: *which* particle, *what* mass, *why* that abundance — with the abundance itself upgraded from a fitted parameter to $1/\pi - 1/(2\pi^2)$, a number the CMB now checks rather than supplies. Concordance cosmology is the question’s premise, not its answer; a fluid with the right equation of state tells us how the scaffolding behaved, and only a laboratory line at a frequency declared in advance can tell us what it was made of.

K Couplings, locks, and the strong-CP assignment

Objection 31 (O-Kaxion-31) (the DFSZ-degeneracy skeptic)

“Exactly DFSZ means experimentally indistinguishable from DFSZ.”

Response. Indistinguishable in *reach*, maximally distinguishable in *identification* — which is the point. The coupling is identical because $E/N = 8/3$ places the Kaxion on the DFSZ line (a convenience, since every published DFSZ-sensitivity statement applies verbatim). But DFSZ predicts no mass while the Kaxion predicts one window; and DFSZ has nothing to say about the rest of the band, while the Kaxion’s count does — one line, and silence everywhere else the fence allows (Part VI). A signal passing the window test, with the rest of the window interrogated and silent, is not “a DFSZ axion” but a reading of the cap loading under the ladder ansatz and of the count that the boundary condition terminating the bulk’s baryonic sector leaves behind.

Objection 32 (O-Kaxion-32) (the QCD-axion loyalist)

“Strong CP is already solved by the QCD axion; why a new mechanism?”

Response. The objection assumes the two are competitors for one job; they are not. GG-A1-Constitution (Problem P7) assigns the strong-CP resolution to the Kaxion itself: the same forced pseudoscalar that is the dark matter relaxes $\bar{\theta} \rightarrow 0$. There is therefore no residual strong-CP problem left for a separate Peccei–Quinn axion to solve, and — since there is no Peccei–Quinn scalar in the spectrum — no field from which to build one. The Kaxion is the strong-CP solution the theory compels, not an alternative installed beside it.

Objection 33 (O-Kaxion-33) (the collider skeptic)

“Gauged $U(1)_B$ requires a TeV Z'_B ; where is it?”

Response. It is not missing; it is a separate, pre-registered prediction. Gauged $U(1)_B$ requires a vector, and the sister paper P1-KK places it in the radial lock cluster near $M_{\text{lock}} \approx 3.85$ TeV — as a Dirichlet-terminated mode that endpoint-localized quarks do not produce at tree level (conditional on GG-A9-SM’s cascade-depth lemma), so the collider exposure is the lock cluster

itself rather than a Z'_B resonance in dijets — at the derived coupling $g_B = 0.565$ (GG-A3-BI6’s terminated-seam capacity theorem). The cohort is deliberately cross-exposed: the same gauge factor whose ledger entry the Kaxion’s coupling sheet carries gives P1-KK its Z'_B , so a collider search and a haloscope search probe one structure.

Objection 34 (O-Kaxion-34) (the rigor checker)

“Vafa–Witten applies to vector-like theories at $\theta = 0$; invoking it for your relaxation argument overreaches.”

Response. The argument uses Vafa–Witten only where it is licensed: for the pure-gluon functional at the point the Kaxion dynamics has already reached, to certify that the relaxed vacuum is the true minimum. The relaxation itself is the standard axion mechanism — the potential $\chi(T)(1 - \cos \theta_{\text{eff}})$ generated by the anomaly, minimized at zero — and needs no theorem beyond the existence of the QCD-generated potential that lattice computation exhibits directly (Appendix C.3). The gauged sector adds no new θ -source: $U(1)_B F B \tilde{F}_B$ is separately relaxed by the same field through the locked portal. Where assumptions enter — the chiral reduction to nucleon-level couplings — they are fenced explicitly as the model-level input (the chiral reduction of §73.4), which is the honest boundary of the claim.

Objection 35 (O-Kaxion-35) (the algebra auditor)

“A trace ratio of 2/5 in the anomaly becoming 5:2 in the couplings looks like sign-and-inversion gymnastics.”

Response. The algebra the objector distrusts was sound — the inversion is a change of basis, written out in Appendices C.1 and C.4 — but the objection is now answered at a more radical level: the readout it targets is withdrawn (Part VI). The anomaly fixes the *source* ratio 2/5 (GG-A3-BI6’s coefficient vector), and that entry stands; the inversion into a light-channel response ratio presupposed a light baryonic vector for the condensate to drive, and the boundary condition that terminates $U(1)_B$ leaves none below M_{lock} . Nothing gymnastic survives because nothing light survives; what remains is the ledger entry, and its laboratory-scale consequence is an absence, not a ratio.

Objection 36 (O-Kaxion-36) (the chiral perturbation theorist)

“Quoting the photon bracket as 0.75(4) hides the model dependence of the -1.92 chiral term.”

Response. The $-1.92(4)$ is not this paper’s estimate; it is the community-standard next-to-leading-order chiral computation with lattice inputs, the same number every DFSZ analysis on every exclusion plot uses, and its quoted error *is* the model dependence as currently bounded (Section 27). The paper’s own contribution to the bracket — $E/N = 8/3$ — is exact group theory with zero uncertainty (Theorem 27.1), so the bracket’s $\pm 5\%$ is entirely inherited, shared with the field, and shrinks with the same lattice program that sharpens everyone’s DFSZ line. If the chiral term shifts by more than its stated error, the DFSZ band moves for the whole community; the Kaxion moves with it, still locked to $E/N = 8/3$, still at the same fractional offset — the identification content survives untouched.

Objection 37 (O-Kaxion-37) (the loop skeptic)

“The magnon port reads an electron coupling that in DFSZ-like models arises only at loop level; your tree-level $25/4$ is optimistic.”

Response. The objection lands, and harder than its author intends. On the computed KSVZ-class sheet the electron coupling is indeed loop-level, $C_e \sim 3 \times 10^{-4}$ (Part VIII) — which is why no magnon port is staged: at this f_K it has no reachable target, and the $25/4$ readout it was designed to serve is withdrawn (Appendix F). The engineering is recorded there because it is sound technology; nothing in the surviving program rests on it.

Objection 38 (O-Kaxion-38) (the topological-defect expert)

“Two coupled pseudoscalars generically produce domain walls; $N_{\text{DW}} = 1$ by assertion does not remove them.”

Response. $N_{\text{DW}} = 1$ is the unit color anomaly of the single portal in the period convention Theorem 24.1 adopts (the descent with the anomaly companion’s κ and $t = 5$ made explicit is registered open, and could return $N_{\text{DW}} = 5$ in another convention) — and in either case the single-patch clause, if it holds, makes walls moot, because walls bounded by strings collapse harmlessly *if* strings form at all. Here they do not, conditionally: defect networks require the field to decorrelate across horizons, and the single-patch clause forecloses that route — no Peccei–Quinn restoration (no scalar exists to restore), reheating a factor 12 below the decay constant, and, conditionally, the inflation-era pinning of C-Kaxion-d (Section 13.5). The lifted holonomy adds no loophole, because it is integrated out at the scale it is lifted, long before any cosmological era that could paint it non-uniformly. The defect story is thus a chain with one named conditional, and it is testable at the margin: a confirmed minicluster or wall-crossing signature would falsify the same clauses (see the substructure objection above).

L Experimental practice

Objection 39 (O-Kaxion-39) (the scan-cost realist)

“DFSZ is the hardest target; reaching it will take years.”

Response. The DFSZ premium is real — reaching the DFSZ line costs the standard factor over a KSVZ scan — but it is a premium the community is already paying. IBS-CAPP, with DFSZ depth already demonstrated above 1 GHz and a published roadmap of 1–8 GHz at that depth, and ADMX-EFR, with its eighteen-cavity array at 2–4 GHz, are the programs whose stated reach contains this window (Section 36). Because the Kaxion sits *on* the DFSZ line, those scans need no sensitivity breakthrough: they either find the line or exclude the window. The cost is integration time on funded instruments, not new physics.

Objection 40 (O-Kaxion-40) (the null-result skeptic)

“A null result is not a real test.”

Response. On the contrary, the null is the cleanest possible test. A full DFSZ-depth scan of [1.59, 2.13] GHz that finds no *line* retires the calibrated placement and its ladder (F-Kaxion-1c)

— the Kaxion is then owed at $\gtrsim 6$ GHz on the unpinned branch, or the relic is not this sector’s; the framework is not retired, and this paper does not claim that it would be. The prediction is pre-registered down to the frequency, the depth, and the count; a falsifier that specific is the mark of a real prediction, not an escape hatch.

Objection 41 (O-Kaxion-41) (the spectroscopist)

“Could the withdrawn 56 MHz doublet split (Appendix F) have been resolved?”

Response. Trivially — the split was $\sim 3 \times 10^4$ condensate linewidths, far wider than any instrumental resolution at 1.95 GHz — and the question is kept because its answer explains the surviving test. A single-mode cavity is blind to a tone that far detuned, so partner interrogation was always a deliberate act; that act survives the withdrawal as the count test: split the sub-arrays, interrogate the other rungs, and accept a candidate only if the window is otherwise silent (Part VI).

Objection 42 (O-Kaxion-42) (the background hunter)

“Backgrounds (interference, dark photons, thermal) could fake a signal.”

Response. A background — a radio-frequency line, a thermal fluctuation, a dark photon — has no reason to sit inside a window staked in advance, to carry the halo’s thermal shape, or to breathe with the sidereal day and the year. The Kaxion must do all of them at once, each pre-registered with its error bars (Part IX). No background conspires to satisfy the full set, so the stack is a near-background-free discriminator — the single largest advantage the geometry hands the experimentalist, and the feature that turns the famously laborious axion “candidate vetting” into a one-shot physics statement.

Objection 43 (O-Kaxion-43) (the quantum-readout skeptic)

“Single-photon counting at 1.95 GHz is immature.”

Response. Single-microwave-photon counting at gigahertz frequencies is no longer speculative — transmon-qubit counters have been demonstrated. At 1.95 GHz, however, the thermal occupation at a cavity temperature of 40 mK is $n_{\text{th}} \simeq 0.1$, so the counter is not dark-count-limited there; the gain over a quantum-limited amplifier is at most the factor $\sim (1 + 2n_{\text{th}})^{-1}$ in noise, worthwhile but not decisive, and the paper’s baselines are computed for the amplifier. And the fallback is proven: a quantum-limited Josephson parametric amplifier retains full discovery capability if the counter underperforms, so the program does not stake discovery on an immature readout.

Objection 44 (O-Kaxion-44) (the halo-fraction questioner)

“The haloscope assumes the Kaxion is 100% of the local dark matter.”

Response. It is 100% by the misalignment-completeness conjecture, named as such: the π -locked closure makes the Kaxion saturate Ω_{DM} — the whole inventory, not a fraction — so the density entering the signal is the full local $\rho_{\text{loc}} \simeq 0.45 \text{ GeV/cm}^3$, with no dilution by a second component. And even if one doubted the local density, the two verdicts that carry the geometry — the window and the count — are frequency data, independent of ρ_{loc} by construction; only the coupling verdict (F-Kaxion-1c) depends on it, and its definition states the density it assumes (Part VI).

Objection 45 (O-Kaxion-45) (the veteran operator)

“Camping does not integrate like $\sqrt{t}$ in the real world: gain drifts, $1/f$ noise, and RFI put a floor under everything.”

Response. That floor is engineering, and the paper budgets it rather than waving at it. The radiometer regime is protected by working at 100 Hz bins where $t\Delta\nu \gg 1$ throughout; gain drift is removed by the continuous calibration tone and enters the budget only as its measured residual; RFI is handled by the same environmental-antenna veto the field already runs; and the mirrored-bin correlations of JPA receivers are noted with their measured few-percent size (Table 26). The honest statement is quantitative: real camps carry 10–100 $\times$ duty and calibration overheads over the idealized times, and every quoted timescale in Parts VII–X says so explicitly. Multiply the 5σ times of Table 17 by a pessimistic hundred and the conclusion — minutes to hours, not years — is untouched; that robustness to the operator’s own skepticism is why the design leans on locked *ratios* rather than absolute rates wherever a verdict is at stake.

Objection 46 (O-Kaxion-46) (the semantics skeptic)

“A 544 MHz window is a quarter of an octave; ‘camping’ on that is just scanning with better publicity.”

Response. The distinction is statistical, not semantic. A scan pays the look-elsewhere penalty of its full sweep and dwells per position only long enough for its target depth; the pre-registered band collapses the trials factor to the declared window (§37.1) and concentrates the entire time budget inside it — which is why the window costs of order one hundred days of single-cavity live time (hours for an array-class instrument) at depths a blind GHz-scan buys with years. And the window is not the prediction’s whole content: the discrete ladder gives seven marked rungs, the benchmark rung is ± 25 MHz wide, and the address and the count ride along at no cost (Section 27). Call it a short scan if preferred — 0.42 of an octave; what cannot be renamed is the factor the arithmetic delivers, and the fact that the band’s center was printed before the data.

Objection 47 (O-Kaxion-47) (the local-density skeptic)

“If the local halo density is half the standard value — a dark disk, a void, a stream — your minutes become days and your kill dates slip.”

Response. Scale every rate by the skeptic’s factor of two and read the tables again: the camped 5σ times quadruple, from seconds to seconds and from minutes to minutes (Table 17; $t \propto \rho^{-2}$), the rung exclusions of Figure 20 rise from $g_{\text{DFSZ}}/28$ to $g_{\text{DFSZ}}/20$ and the window sweep from $g_{\text{DFSZ}}/4.3$ to $g_{\text{DFSZ}}/3.1$, and the F-Kaxion-1c dates of a scanning campaign move by months to years, since a scan’s coverage time scales the same way. The density enters as $\sqrt{\rho}$ in amplitude and is already carried at $\pm 22\%$ as the budget’s dominant systematic (Table 20); a factor-two local deficit is a 2.3σ surprise to the halo community before it is a problem for this paper. And the identification observables are engineered to be immune: the window, the address, and the count are frequency-and-existence data in which ρ_{loc} never appears, so even the skeptic’s worst halo delays the verdict without ever blurring it.

Objection 48 (O-Kaxion-48) (the burned skeptic)

“The 1.036 GHz affair shows how five-sigma candidates evaporate; promising a count verdict ‘within a day’ is bravado.”

Response. The 1.036 GHz candidate evaporated *through* the vetting protocol, not despite it — persistence, field-off, cross-instrument — and that protocol is precisely what §37.3 adopts, with one addition the field has never had: a pre-registered fence and count. The day-one claim is arithmetic, not confidence: at camped signal-to-noise a rung’s cavity bandwidth reaches decisive depth in seconds to minutes of live time, and its full ± 25 MHz in nine hours for the KAXIO array or ten days for a single CAPP-MAX-class cavity (Tables 17 and 29), so the calendar cost of a rung’s count verdict is re-tune and settle, not integrate and hope. What took the 1.036 GHz team months was establishing that a lone excess was *nothing*; establishing whether a fenced window holds a second line is a different and far faster epistemic task. Proving an open-ended negative is slow; interrogating a printed fence is not.

Objection 49 (O-Kaxion-49) (the blinding skeptic)

“A sealed frequency offset cannot blind an analysis whose target is public; everyone knows where to look.”

Response. The blinding is not meant to hide the target — nothing could, and nothing should. It is meant to break the one feedback loop that matters: the analyst’s ability to steer thresholds, vetoes, and bin boundaries by watching the science band respond. With the offset sealed, every such choice is made against data whose absolute frequency is unknown to the chooser by up to 200 Hz — a tenth of a linewidth, preserving all internal relations while decoupling tuning decisions from the answer (§64). This is exactly the sealed-parameter pattern of LIGO’s blind injections and the muon $g - 2$ clock offset, both of which guarded analyses whose targets were equally public. Knowing where to look was never the danger; knowing what you want to see while you set the cuts is.

Objection 50 (O-Kaxion-50) (the operations manager)

“An 85% duty cycle for a full year on a dilution-refrigerator experiment is a fantasy; real haloscopes live near half that.”

Response. Real *scanning* haloscopes live near half that, because their duty is consumed by the scan itself: re-tune, re-couple, re-calibrate at every step, with each mechanical touch risking a thermal cycle. A camp deletes most of that overhead by construction — the tuning system sweeps for eight weeks in two bands, moves six times for the rung camps and once for the cavity-set exchange (Table 31), and the rest of the year is a camp; the magnet is a persistent-mode MRI solenoid whose product line holds field for years in hospitals with far worse infrastructure than the Dark Wave Laboratory; and the cryogenic plant is the EFR baseline’s, engineered for continuous Generation-2 operation. The 85% is a requirement with a budget line behind it (the duty-cycle row of Table 33), and its failure mode is graceful: duty lost is exposure stretched, never a verdict changed — the year becomes fourteen months, and the falsifiers do not care.

M Instruments and astrophysics

Objection 51 (O-Kaxion-51) (the magnet engineer)

“The nucleon port needs 45.8 T — impossible.”

Response. True for a free-spin Larmor port, and we say so: proton resonance at 1.95 GHz would need 45.8 T, incompatible with a quiet 10 mK haloscope. But the wall is about one readout mechanism, not the coupling. The non-Larmor internal-field (hyperfine-enhanced) concept reaches the window at a sub-tesla trim field by letting the crystal’s own hyperfine field — tens to hundreds of tesla — supply the resonance, with the nuclear ensemble essentially fully polarized at 10 mK; the electron-magnon port at 70 mT was the mature first stage (Part VIII). The objection is answered by the readout, not evaded — though the answer is now archival: with the 25/4 readout withdrawn and the KSVZ-class sheet out of reach at this f_K , both spin ports are records, not proposals (Part VIII).

Objection 52 (O-Kaxion-52) (the untestability skeptic)

“You computed the baryonic coupling astrophysically negligible, so it is untestable.”

Response. When earlier revisions first answered this objection, the reply was the matched-channel ratio; that readout is withdrawn (Part VI), and the honest answer is now different. The absolute dark-baryon coupling is indeed too weak to move a galaxy — we computed it (§73.4) — and too weak, on the computed sheet, for any current laboratory port to read. The baryonic sector is tested where its strength actually lives: at the collider, through the Z'_B at M_{lock} that the same gauge factor requires (P1-KK). Untestable is the wrong word; the test moved to the machine that can afford it.

Objection 53 (O-Kaxion-53) (the strong-gravity observer)

“EHT and LIGO see pure Kerr; does that not constrain the Kaxion?”

Response. No — and the consistency is a feature. As cold dark matter the Kaxion makes no horizon-scale imprint; the companion black-hole paper (Arisaka, 2026, C3-BH) shows the predicted deviations from Kerr are suppressed by $(\ell_{\text{lock}}/r_s)^2$, far below any Event Horizon Telescope shadow or LIGO ringdown measurement. The absence of deviations is therefore exactly what the Kaxion predicts; it is a consistency, not a constraint, and a confirmed horizon-scale deviation would embarrass the framework rather than support it.

Objection 54 (O-Kaxion-54) (the window skeptic)

“The window may be wrong if Ξ_{cap} shifts.”

Response. The prediction is robust to exactly this. The window is bracketed by announced programs from both edges, so a shift of the benchmark within the discrete Ξ_{cap} range moves a *marked row* inside a pre-committed window, not the window’s existence; and the DFSZ depth and the count are fixed regardless of where the center lands. Decisively, a confirmed compact-angle-like line at *any* other frequency also falsifies the calibrated placement (F-Kaxion-1d), so a wrong location is fatal to the placement, while a second line inside the fence refutes the count — the prediction cannot be quietly relocated to save itself.

Objection 55 (O-Kaxion-55) (the magnon realist)

“QUAX-class magnon experiments sit orders of magnitude above the sensitivity your booster requires; the gap is not an ‘R&D year’, it is a chasm.”

Response. The comparison misplaces the gap. QUAX searches for the *electron* coupling of a QCD axion whose signal power is set by $g_{aee}^2 \rho / m_a$ at unknown mass — a blind, deep, scan-limited problem. The booster’s task was categorically easier on three axes: the frequency is known (no scan), the trigger is a coincidence with a $\sim 3000\sigma$ electromagnetic line (no blind discovery threshold), and the observable is a *ratio* in which most calibration burden cancels (Eq. (F.6)). The requirement it was to meet is $\text{SNR}_B = 5$ per five live days at a known bin — and the paper never claimed it as demonstrated: it was fenced as a design requirement, gated behind a bench year. The point is now archival: the port is withdrawn with its target (Part VI; the record is Appendix F), and the photon-only configuration — always the fallback — is the baseline (Table 33). A chasm one refuses to price is a fantasy; a gap with a measured width, a gate, and a fallback was called engineering, and the record keeps the price.

Objection 56 (O-Kaxion-56) (the magnetics engineer)

“Holding 70 mT stable and homogeneous beside a 9.4 T solenoid will be a permanent commissioning nightmare.”

Response. (An engineering question about the module recorded in Appendix F; answered for the record.) The numbers are less dramatic than the adjacency sounds. At the side-cryostat station the stray field is of order 10 mT with gradients set by meters of standoff; three staged layers — passive mu-metal ($\sim 10^3$), a superconducting can that freezes the residual at cooldown, and an active loop servoed on an NMR probe — deliver the required ± 0.1 mT, which is a 10^{-3} trim on the bias coil, well inside routine magnet-lab practice (§F.3). The demanding specification is homogeneity across the sphere array ($\Delta B/B < 10^{-4}$), and that is a coil-design problem solved daily by MRI shim engineering at far larger volumes. The stress cases — ramps and quenches — cost duty, not calibration: the can saturates gracefully, the module pauses, and the field record from the probe pair decides when data resumes. Commissioning will be work; nightmares require an unbounded unknown, and this system has none.

Objection 57 (O-Kaxion-57) (the observer)

“The JWST ‘early massive galaxies’ story is contested — dust, IMF assumptions, calibration — and you lean part of the case on it.”

Response. The paper leans on JWST asymmetrically, and the asymmetry is the point. The *qualitative* direction — the data demand more early structure, not less — is what survives every proposed systematic (dust and IMF revisions move masses, not the sign of the surprise), and that direction bounds away precisely the candidates that suppress small-scale power, the Kaxion’s rivals (Part IV). The one *quantitative* head-start lever that survives the termination — the collapse threshold its cosmology companion carries as a working value — is quoted at model level, and nothing downstream — no falsifier, no laboratory number — depends on them. If the early-galaxy tension dissolves entirely, the Kaxion loses a supporting contrast and keeps its

entire case; if it hardens, the warm and fuzzy alternatives lose outright. An argument used only where it is one-sided is not a dependence; it is a free roll.

Objection 58 (O-Kaxion-58) (the simulator)

“A 3% power-spectrum head start is smaller than the systematic spread between galaxy-formation simulations; the structure argument overreaches.”

Response. Agreed — which is why the head start is presented as a *contrast*, not a measurement. Against the simulation spread, a percent-level power boost of the kind earlier readings carried would indeed be invisible, which is one reason this paper no longer carries one; against a fuzzy-field *suppression* of small-scale power by orders of magnitude at the JWST-relevant masses, it is the difference between the right side and the wrong side of a cliff (Part IV states both fences). The paper’s structure claims are all of that cliff-edge kind: cold-not-fuzzy by sixteen orders in mass, cuspy-not-cored by thirty-two orders in cross-section at the very least — discriminations that no simulation systematic can blur. Where percent-level cosmology is genuinely at stake (the π -locked inventory), the test is the CMB inventory itself, not simulations; where simulations rule, the paper claims only signs. Matching the precision of a claim to the precision of its arena is not overreach; it is the opposite.

N Foundations and philosophy

Objection 59 (O-Kaxion-59) (the empiricist philosopher)

“Why should the universe be describable by axioms at all? Physics is an empirical science, not a branch of mathematics.”

Response. The axioms here are not a metaphysical wager that reality is mathematical; they are a compression claim about the regularities experiment has already delivered, stated sharply enough to be wrong. ‘Physics from axioms’ means: here are four structural commitments and seven postulates; if they are the right compression, the consequences of Table 1 follow with no further freedom — and each consequence is an empirical instrument’s business, not a philosopher’s (Table 1 indexes clause to duty). The empiricist’s own criterion is thereby honored more strictly than usual: most theories meet the data at a fitted interface, while this one has removed the interface and left nothing between axiom and antenna but arithmetic, declared Standard-Model inputs, one declared calibration, and the named conjectures of the deck — each flagged where it enters. Whether the universe ‘is’ axiomatic is a question the paper never needs answered; whether *this* axiom set compresses it is answered at 1.95 GHz.

Objection 60 (O-Kaxion-60) (the underdetermination theorist)

“Consistency arguments underdetermine reality: many self-consistent structures exist, so deriving ‘the’ world from consistency is rationalism, not physics.”

Response. Granted freely — consistency alone underdetermines, and the framework never claims otherwise. Its claim has the classical two-premise form: consistency *given the constitution*,

and the constitution is a contingent, empirical bet whose justification is nothing deeper than the falsifiable consequences it forces. Other self-consistent structures surely exist, and the concession should be made in full: confirmation attaches, in the first instance, to the predictive sheet — the frequencies, the locks, the window — and the axioms behind it are preferred over unexhibited rivals that might host the same sheet on grounds of risk and parsimony, not of proved uniqueness. The rationalist’s sin is deriving necessity from armchair coherence and then insulating it from data; the discipline here is the inverse — the deductive chain exists precisely to expose a single arbitrary-looking choice (the axioms) to maximal empirical risk, concentrated into kill criteria with dates (Appendix E.2). Underdetermination is escaped the only way it ever is in science: not by proving uniqueness, but by being the candidate that volunteered for execution.

Objection 61 (O-Kaxion-61) (the modal metaphysician)

“If every constant is derived, nothing about physical law is contingent — your universe leaves no room for things to have been otherwise.”

Response. The contingency has not vanished; it has been relocated and concentrated. Within the framework, given the axioms, the constants are indeed necessary — that is the content of ‘no dials.’ But the axiom set itself, and behind it the brute fact that anything satisfies it, remain as contingent as anything in Hume: the theory tells us what *must* follow if the world is a GG-6 world, and stays silent on why there is a GG-6 world at all (the capstone GG-A11-Origin addresses which questions of that kind are even well-posed). One may find it strange that the modal joint sits between axioms and everything, rather than being smeared across twenty-six parameters — but strangeness is not an objection, and the arrangement has an empirical virtue the smeared alternative lacks: a single joint can be stress-tested. Every falsifier in this paper is an attempt to break it.

Objection 62 (O-Kaxion-62) (the Popperian)

“A framework that derives the Standard Model, the cosmos, and dark matter at once is exactly the ‘explains everything’ holism Popper warned against.”

Response. Popper’s target was theories whose explanatory reach came with *immunity* — psychoanalysis could accommodate any behavior, Marxist history any event. Reach with exposure is the opposite pole, and it is where this framework lives: the same integer ledger that explains much can be killed in any one place — a second line inside the fence, a lock cluster absent at 3.85 TeV, a converged $w \neq -1$, a line outside the window — and a kill at any *theorem-level lock* is a kill everywhere, because the locks are not compartmentalized: one integer ledger carries them all (Section 3 traces the shared spine; App. E.2 lists this paper’s four kill conditions). Below theorem level the degradation is graded, and the ledger says so in advance: a benchmark can move within its declared discrete window, and a named conjecture can die taking only its named dependents. Popper himself drew the line exactly here: what matters is not how much a theory explains but what it forbids. This construction’s prohibitions are concrete: a second light line inside the fence, a line outside it, an empty lock cluster at

3.85 TeV, a CMB inventory off the π -locked fractions — each with a stated depth and, where an instrument exists, a date.

Objection 63 (O-Kaxion-63) (the realist)

“OGGEP makes observation constitutive of the world — that is idealism. Does the bulk exist when no one projects it?”

Response. OGGEP concerns what is *physical content*, not what *exists*: it says the lawful description available to any finite observer is the boundary projection, exactly as thermodynamics says the available description of 10^{23} molecules is a handful of state variables — without anyone concluding that molecules exist only when coarse-grained. The bulk in this framework has the standing of a gauge orbit’s interior or a coordinate chart in general relativity: lawful machinery whose *invariant* content — and only that — is physical; it is the structure whose consistency requirements force the observables. The stance, here and in O-Kaxion-65, is structural realism about the projection-invariant content, and ‘existence unobserved’ is a question the formalism neither needs nor answers. What the formalism does answer, distinctively, is a realist’s dream disguised as an idealist’s premise: the projection is so tightly lawful that the unseen structure leaves fingerprints — a window, an address, a count — on a laboratory bench. Idealism was never that falsifiable.

Objection 64 (O-Kaxion-64) (the anthropic critic)

“‘The universe occupies the deepest vacuum’ is a selection principle, and selection principles are anthropics in disguise.”

Response. Anthropic selection conditions on our existence within an unexplained ensemble; the GRIP selection is a dynamical statement within one system — descent to the minimum of a derived free-energy functional, no more anthropic than a ball rolling to the bottom of a bowl (Section 10.4). The disguise runs the other way: an anthropic argument could never tell you the depth of the bowl, while this one must — the same functional that selects the vacuum fixes the attractor’s curvature and hence numbers this paper spends (θ_i ’s single-patch coherence; the stability that makes the Kaxion permanent). And where a selection principle would be untestable, this one is over-determined: if the vacuum selection is wrong, the electroweak minimum, the dark-energy equation of state, and the Kaxion’s stability all fail together, each measurably. A principle that risks that much has left the anthropic genre entirely. One concession is owed and given: why *this* functional — why these axioms — is a selection question the dynamics cannot answer about itself; it is left open explicitly, and the capstone GG-A11-Origin maps which versions of that question are well-posed.

Objection 65 (O-Kaxion-65) (the nominalist)

“Integers are our bookkeeping, not nature’s furniture. ‘ $K = 35$ is in the world’ confuses the map with the territory.”

Response. The claim is more modest, and harder to escape: whatever the territory is, it supports *exact discreteness* — windings, anomaly coefficients, cycle counts — and exact discreteness is integer-valued in every admissible map. When the same ‘bookkeeping entry’ surfaces as

a TeV mass ratio at a collider and as a fenced window in a cavity — two instruments, two continents, no shared calibration — the nominalist must either credit the territory with the structure or credit an extraordinary conspiracy of maps (the double-reading is the explicit program of Parts VI–X). Physics has been here before: no one now thinks electric charge quantization is a notational preference. The claim, as in O-Kaxion-63, is structural realism: the projection-invariant structure — exact discreteness read twice — is located *in the world*, while the metaphysics of number is left open either way. If the fence and the count are confirmed, that structure will have been measured, not merely written; if it fails, the framework loses empirically — and the deeper dispute about what numbers *are* stands exactly where philosophy left it.

Objection 66 (O-Kaxion-66) (the teleologist)

“Suppose you find it. With the constants derived and the dark matter identified, what would remain for physics — and what was the point?”

Response. What ends is a particular genre — parameter archaeology, the fitting of numbers nature was hiding — and its end would be a graduation, not a funeral. Everything above the parameters remains: the Kaxion identified is an instrument aimed at the halo (velocity structure, streams, the dark sky as astronomy rather than inference — Part X’s spectrometer deliverables are the first page of that book); the framework confirmed is a machine for questions we could not previously pose sharply, from the horizon’s information budget (GG-C3-BH) to why observers find themselves at this cosmic epoch at all. And if the point of the enterprise was ever ‘more constants to fit,’ physics chose its point badly. The point was always to know what the world is made of and why it holds together; a theory that answers a piece of that has not exhausted physics — it has paid physics’ oldest invoice.

Objection 67 (O-Kaxion-67) (the historian of science)

“Mathematical elegance has misled physics repeatedly — $SU(5)$ ’s proton decay, the string landscape. Why is π -locking different?”

Response. The historian’s own cases carry the answer. $SU(5)$ was beautiful *and prohibitive* — it forbade a stable proton, experiment looked, and the beauty died on schedule: that is the system working, and it is exactly the wager this cohort re-enters with the proton watch among its own falsifiers. The landscape is the true cautionary tale, and its failure was not elegance but the *loss of prohibition* — 10^{500} vacua forbid nothing. π -locking is structurally on the $SU(5)$ side of that divide: it is one assignment, with no ensemble to retreat into, already tested to 1.5% against the CMB inventory and exposed again through every number in this paper. History’s lesson is not ‘distrust beauty’; it is ‘trust nothing that cannot lose.’ The locks can lose by Thursday of discovery week. That is the difference.

Objection 68 (O-Kaxion-68) (the contingency skeptic)

“Your ‘no dials’ virtue is retrospective luck: had the cascade ended at $K = 36$, you would now be praising 36.”

Response. The counterfactual fails on its own terms, and the honest vocabulary is the philosopher’s own: use-novelty. First, K is not a stopping point that could have drifted by one: 35 is doubly determined — $K = LJ$ and $K = M^2 + N^2 + L^2$ simultaneously — and 36 satisfies neither factorization with the ledger’s smaller entries; the cascade does not ‘end near’ 35, it lands on it or contradicts itself (Section 6.2.3). Second, the electroweak scale: $v = M_{\text{lock}}(2/5)^3 = 246$ GeV cashes the same K against a long-known number, and its status is stated as what it is — a *retrodiction*, since v was known first — with the use-novel point that v is not an input anywhere in the closure: the Diophantine cascade fixes K without it (Table 2 grades the entry accordingly). The pre-registered weight therefore rests on what no instrument has yet read: the window, the DFSZ address, and the single-line count. A skeptic who suspects accommodation dressed as foresight needs only the calendar: the third reading of K is scheduled, and it has not happened.

Objection 69 (O-Kaxion-69) (the Duhemian)

“When a falsifier fires, modus tollens does not tell you which premise died — your kill criteria underdetermine whether the axioms, the lattice input, the benchmark, or a named conjecture takes the bullet.”

Response. Correct in general — and answered here by bookkeeping done in advance: the ledger assigns blame before the data arrive. F-Kaxion-1d and the count test are built from dimensionless data — a location and an existence — that hold whatever the lattice or the halo do; if the count fires, the parents’ registered accounting takes the bullet and the second line names the mode; if F-Kaxion-1d fires, the calibrated placement does — or, at the randomized-angle mass, the pinning conjecture C-Kaxion-d. F-Kaxion-1c strikes the closure sheet exactly as fenced: it retires the calibrated benchmark and the ladder ansatz, and is booked against the calibration (O14) and C-Kaxion-b. The Standard-Model-side inputs — $\chi(T)$, $g_*(T)$, the transport matching — carry stated errors and can shift the *window*; they cannot create or delete the count or the location, so they are never candidates when those fire. And the deniable elements in the chain are exactly those the deck names — the calibration, the ansatz, misalignment completeness (MC), and C-Kaxion-b, c, d — which is precisely why they are named: if denial comes, it will be public, specific, and priced, because the master ledger’s rows state today which claims fall with which premise. The Duhemian problem feeds on holism declared after the fact; a blame table published before the experiment is its antidote.

Part XV

References

Cited works in two groups — the GG-Theory cohort and the external literature — each entry carrying a relevance note and its first citation point

The reference list is organized in two sections. **Section A** lists the twenty-one sister papers of the GG-Theory cohort, in the canonical order of the program (the constitutional root, the GG-6 Trio, the boundary law and its pillars, the three textbook readouts, the capstone, the two phenomenology branches, and the Foundations Branch); every cohort paper is cited in the body of this one, because the Kaxion sits inside a single deductive structure and its verdict propagates through all of it. **Section B** lists the external literature — the ninety-year evidence base for dark matter, the axion’s theoretical ancestry and its QCD inputs, the haloscope instrumentation and the published searches this paper is anchored to, and the halo astrophysics behind the signal model — alphabetically by first author.

Every entry carries a **Relevance** paragraph stating what the work contributes to this paper and where it enters, followed by the location of its first citation. The list contains no entry that is not cited in the text.

All cohort preprints are available at <https://preprints.arisaka-gg.org/>. External references are cited under their published or preprint identifiers as given.

A. Papers by Katsushi Arisaka (the GG-Theory cohort)

Arisaka, K. (2026). 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. The constitutional charter of GG-Theory: it states the four axioms (locality/causality LC, the geometry–gauge equivalence GGEP, its open-system extension OGGEF, and maximum-symmetry/minimum-freedom MSMF) and the seven postulates, and shows that this rulebook is rigid enough to force the six-dimensional bulk, the anomaly identity, the integer ledger, and the open-system boundary law as theorems in the companion papers. For a reader new to the program, this is the right entry point: it explains *why* the theory has no continuous dials and therefore why a dark-matter candidate within it must be forced rather than chosen. Decisively for the present paper, it assigns the strong-CP closure to the Kaxion in Problem P7 and registers the Kaxion’s laboratory line as a constitutional falsifier with a pre-declared kill criterion (the

count and the window of Part VI now carry it) — so the Kaxion is constitutionally obligated, not appended. First cited in §4.1.

Arisaka, K. (2026). GG-Theory with GG-6: Geometry–Gauge Locking at the 6D Bulk (GG-6) at the Heart of Grand Geometric Theory (GG-Theory). Arisaka-GG-A2-GG6.

Preprint, available at <https://preprints.arisaka-gg.org/a2-gg6/>.

doi:10.5281/zenodo.22064298.

Relevance. The bulk paper, and the geometric foundation of everything downstream. It proves the canonical-metric theorem $ds_6^2 = e^{-2r}\eta dx^2 + dr^2 + d\theta^2$ with the two *non-spatial* coordinates — a scale interval I_r and a phase circle S_θ^1 — the dimensional-minimality theorem forcing $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S_\theta^1$, the four-factor bulk gauge group, and the OGN-5 Diophantine cascade (1, 3, 5; 7, 35). The conceptual lesson a newcomer should carry away is that the “extra dimensions” here are not spatial (so they evade the usual collider and gravity bounds) but a renormalization-scale axis and a phase angle. Every geometric ingredient of the Kaxion — the phase circle it lives on, the rank-six holonomies of that circle that the count lifts, and the throat depth $K = 35$ that anchors the cap-loading ladder — is supplied by this paper. First cited in §4.1.

Arisaka, K. (2026). Anomaly Freedom of BI-6: The Modified Bianchi Identity (BI-6) for the Six-Force, Six-Dimensional Bulk (GG-6). Arisaka-GG-A3-BI6.

Preprint, available at <https://preprints.arisaka-gg.org/a3-bi6/>.

doi:10.5281/zenodo.22064300.

Relevance. The anomaly paper, and the reason the Kaxion must exist at all. It proves the rank-one criterion forcing a single anti-self-dual Green–Schwarz two-form B_2 , the modified Bianchi identity $dH_3 = \text{tr } R^2 - \text{Tr } F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Q_i}^2$, and the exact trace-ratio computation $\text{Tr}(B_0^2)/\text{Tr}(Y_0^2) = 2/5$. The pedagogical point is the modern one that anomaly cancellation is not an optional consistency check but an architectural constraint: for the GG-6 fermion content it is satisfiable only if the two-form is present, so deleting the field that becomes the Kaxion would render the quantum theory inconsistent. The Kaxion’s existence and its baryonic ledger entry both descend directly from these two results. First cited in §4.1.

Arisaka, K. (2026). From Euclid to GG-6: The Realization of Euclid’s Elements in Six Dimensions — GU-5, OGN-5, and a Parameter-Free Reconstruction of Modern Physics. Arisaka-GG-A4-Euclid.

Preprint, available at <https://preprints.arisaka-gg.org/a4-euclid/>.

doi:10.5281/zenodo.22064302.

Relevance. Carries the GU-5 reconstruction of the universal constants $\{c, G_N, \hbar; k_B, Q_B\}$ from the bulk geometry, and the historical methodology of the program. The thesis a reader should take from it is that each scientific revolution has absorbed one empirical constant into geometry one dimension higher — special relativity absorbing c , general relativity G_N , quantum mechanics $\hbar$ — and that the dark-sector instance of this pattern is the absorption of dark matter into the non-spatial phase circle. The Kaxion is presented in this paper not as a new particle but as the latest step in a three-century lineage of geometrizing the previously contingent. First cited in §4.1.

Arisaka, K. (2026). GDOS: Geometric Dynamics of Open Systems — 6D GG-6 Bulk Projected to 4D Boundary for Quantum Measurement, General Relativity, and QFT. Arisaka-GG-A5-GDOS.

Preprint, available at <https://preprints.arisaka-gg.org/a5-gdos/>.

doi:10.5281/zenodo.22064304.

Relevance. The boundary-projection paper, which tells one how the six-dimensional bulk is seen by a four-dimensional observer. It supplies the canonical boundary law (T6), the four-currents Ward package (T9.7) by which BI-6 inflow protects the electromagnetic, baryonic, and informational currents, and the TGG/CCI bridge lemmas. For the present work it is the source of the rule that the Kaxion may couple on the boundary only through those conserved currents, which is why its portals come in the locked proportions they do rather than as free couplings; the Kaxion’s boundary effective action and its conserved-current couplings are all governed here. First cited in §4.1.

Arisaka, K. (2026). The GRIP on GDOS — GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind. Arisaka-GG-A6-GRIP.

Preprint, available at <https://preprints.arisaka-gg.org/a6-grip/>.

doi:10.5281/zenodo.22064306.

Relevance. The runtime paper: it defines GRIP and the Morse free-energy landscape $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$ (GFF-ST2) and proves equilibrium uniqueness $[\Pi_*] = \arg \min F_{\text{GFF}}$ (GFF-ST7), together with the two-regime runtime definitions def:GGA, def:GRIPin, def:GRIPout used throughout this paper. The picture a reader should form is of a deterministic descent of a single cost functional whose critical points on the compact-angle (Wilson-line) stratum sit at discrete OGN-5 phases — not a continuum. The Kaxion is born, selected, and stabilized as the Geometric Attractor of this cascade ($\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$), and the runtime origin of the $\mathcal{O}(1)$ misalignment angle — fixed by the observed abundance in the present paper — lies here. First cited in §4.1.

Arisaka, K. (2026). The Qi-4 Pillars — Qi-QUA, Qi-EP, Qi-CON, Qi-HAL: The Quantum-Information Backbone of GDOS. Arisaka-GG-A7-Qi4.

Preprint, available at <https://preprints.arisaka-gg.org/a7-qi4/>.

doi:10.5281/zenodo.22064308.

Relevance. Contains Theorem T-QUA, the Arisaka equation $J_E/Q_B = J_e = J_B = J_{Q_i}$ with $Q_B = (2/5)E_{\text{Ry}}/K$, which locks the four boundary currents into a single operational identity. The takeaway is that the electromagnetic, electron, baryonic, and informational channels are not independent couplings to be measured separately but four faces of one conserved quantity, fixed by the same integers. The Kaxion’s baryonic ledger entry — $-2/5$ of X_4 , the same $(5, 2)$ pair — is the dark-sector reading of this current lock; earlier revisions read it out as a $25/4$ matched-channel power lock, a readout now withdrawn (Part VI). First cited in §4.1.

Arisaka, K. (2026). 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. Promotes the textbook quantum-mechanical postulates — the Born rule, projective update, decoherence, and einselection — to theorems of the boundary completely-positive trace-preserving (CPTP) channel, rather than axioms layered on top of unitary dynamics. For this paper

it underwrites two practical things: the quantum description of the coherent Kaxion condensate as a classical field with definite occupation, and the quantum-limited and sub-quantum-limited (squeezed, photon-counting) readout chains that the detection Parts rely on to beat the standard quantum limit. It is the bridge from the abstract framework to what a real receiver registers. First cited in §4.1.

Arisaka, K. (2026). The Standard Model as Theorem of GG-Theory: Its 28 Parameters as Observations of One Six-Dimensional Geometry. Arisaka-GG-A9-SM.

Preprint, available at <https://preprints.arisaka-gg.org/a9-sm/>.

doi:10.5281/zenodo.22064312.

Relevance. The Standard-Model paper, and a key parent of the present work: it derives the Standard Model's continuous parameters from the integer ledger. The results a reader needs are the dyadic-closure theorem $\alpha^{-1}(M_Z) = 2^J = 128$ with its low-energy companion $\alpha^{-1}(0) = 2^J + N^2 = 137$ (the latter is what enters $g_{K\gamma\gamma}$, Part V), the weak-angle rational partition $\sin^2 \theta_W = 25/108$, the Wilson–Hosotani vacuum-value theorem with the scale chain $E_{\text{Pl}} \rightarrow E_{\text{GUT}} \rightarrow M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K}$, and the four-force uniqueness theorem that forces the extra gauge factors $U(1)_B$ and $U(1)_{Q_i}$. It identifies the Kaxion as the geometric replacement of the QCD axion — there is no separate Peccei–Quinn sector — and carries the static strong-CP bound $|\bar{\theta}| \lesssim \sin(2\pi/L) e^{-K}/2^J$ that the dynamical Kaxion then drives to zero. First cited in §4.1.

Arisaka, K. (2026). The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry. Arisaka-GG-A10-Cosmos.

Preprint, available at <https://preprints.arisaka-gg.org/a10-cosmos/>.

doi:10.5281/zenodo.22064314.

Relevance. The cosmology paper, and the immediate parent of the present work: it derives the cosmic inventory rather than fitting it. The numbers to remember are the π -locked partitions $\Omega_m = 1/\pi$, $\Omega_b = 1/(2\pi^2)$, and $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2) = 0.2676$, which agree with the measured dark-matter density at the percent level and which the Kaxion must saturate. This paper states the compact-angle dark-matter theorem, introduces the compact-angle Kaxion Lagrangian and the misalignment closure, and names the two conjectures — misalignment-completeness and compact-angle mass-uniqueness — that the present paper develops, fences, and exposes to experiment. It also carries the collapse-threshold working value, the one early-structure lever on which Section 21 still draws. First cited in §4.1.

Arisaka, K. (2026). 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 closing capstone of the GG-Theory program, tracing the cohort from the constitutional charter to conscious experience along a single runtime. Crucially for this paper, it defines the *Ground Geometric Attractor* (GGA) and proves T-GGA-as-Ground-State — the ground-state refinement of the GRIP cascade by which the Kaxion is identified as the unique deepest attractor of its stratum (Part II, §10.4), with the universality (every quantum of the

condensate is the same Kaxion) and tunneling-stability consequences that follow. The broader lesson is that the dark sector is not a detached add-on but one link in a single axiomatic chain running from quarks to observers; the present paper is that chain's dark-matter entry. First cited in §4.1.

Arisaka, K. (2026). The GG-6 Lock Cluster: A Non-Spatial Kaluza–Klein Resonance near 3.85 TeV as the Standing LHC Falsifier of GG-Theory. Predictions Branch of the GG-Theory program. Preprint, available at <https://preprints.arisaka-gg.org/p1-kk/>. doi:10.5281/zenodo.22064321.

Relevance. The sister predictions paper, and the second face of the same geometry. It carries the heavy non-spatial tower — the radial lock cluster at $M_{\text{lock}} \approx 3.85$ TeV, including the vector Z'_B of the very baryon number the Kaxion's sheet carries — a Dirichlet-terminated mode, degenerate with the cluster and unproduced by light quarks at tree level, so that the collider exposure is the cluster itself — the sector that gauges the very baryon number whose anomaly entry the Kaxion's coupling sheet carries. The complementarity is the point a reader should note: the same five integers are read at the TeV scale by a collider and at the μeV scale by a haloscope, so a confirmation or refutation in one arena is collateral evidence for the other, and the cohort is deliberately exposed to falsification from both. First cited in §4.1.

Arisaka, K. (2026). Why Three Generations: A Three-Leg Derivation of $N_g = 3$ from BI-6, MSMF, and GDOS Morse Stability. Arisaka-GG-P2-N3. Preprint, available at <https://preprints.arisaka-gg.org/p2-n3/>. doi:10.5281/zenodo.22064323.

Relevance. The sibling Predictions-Branch paper that spends the same OGN-5 ledger on the generation count: $N_g = 3$ forced by BI-6, MSMF, and GDOS Morse stability. It stands beside the Kaxion as a second independent falsifier of the same five integers — the ledger that fixes $K = 35$ here fixes $N_g = 3$ there, so a clean failure of either result wounds the same structure. First cited in §4.1.

Arisaka, K. (2026). 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, available at <https://preprints.arisaka-gg.org/p3-electron/>. doi:10.5281/zenodo.22064325.

Relevance. The third Predictions-Branch sibling: the electron as the unique stable readout of the boundary, its identity and no-exit-gate theorems running on the same GRIP runtime that stabilizes the Kaxion as the Ground Geometric Attractor here. Cited for the branch context of the cohort map — the electron and the Kaxion are the lightest charged and the lightest neutral permanent residents of one landscape. First cited in §4.1.

Arisaka, K. (2026). 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, available at <https://preprints.arisaka-gg.org/p4-real/>.

doi:10.5281/zenodo.22064327.

Relevance. The origin-of-*i* member paper of the Foundations Branch: the two real sources of the boundary’s complex structure — the compact phase circle and the causal clock — with positive frequency read as holomorphy. For the present paper the relation is roster-level: the complex structure in which the boundary’s quantum description — including the coherent account of the Kaxion condensate — is written is derived there, from the same compact phase circle on which the Kaxion lives, rather than postulated. First cited in §4.1.

Arisaka, K. (2026). 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, available at <https://preprints.arisaka-gg.org/p5-cpt/>.

doi:10.5281/zenodo.22064329.

Relevance. The CPT member paper of the Foundations Branch, and the companion with which this paper shares the strong-CP resolution: in its two-fate architecture the weak CP phase is locked by geometry while the one relaxed angle $\bar{\theta}$ is handed to the compact-angle pseudoscalar — its Part VII names the Kaxion as the companion that relaxes what remains. It carries C, P, and T as orientation reversals of the three real bulk factors: the mirrors whose breaking the kaxion sector survives. First cited in §4.1.

Arisaka, K. (2026). 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, available at <https://preprints.arisaka-gg.org/p6-proton/>.

doi:10.5281/zenodo.22064331.

Relevance. The second member of this paper’s triptych, and the one that makes the triptych an argument rather than a list. The electron is matter *on* the circle and the Kaxion is matter *of* it; the proton is matter *guarded by* it, its stability carried by two holonomy ledgers whose joint integrality no local process can undo. Read together the three say that the same one-way descent deposits permanence in three different ways on one geometry, and that the Kaxion’s permanence is the least exotic of them rather than the most. Nothing in the Kaxion’s four numbers is inherited from it. First cited in §10.5.

Arisaka, K. (2026). 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, 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 context for a habit this paper leans on without arguing for it. The Kaxion is light because a compact angle admits no local counterterm — not because a symmetry forbids a mass — and a reader who finds that substitution unfamiliar will find the century of argument behind it here rather than in this paper. It supplies no number and no theorem the Kaxion chain uses. First cited in §10.5.

Arisaka, K. (2026). The Circle That Remembers, the Interval That Forgets: GG-Theory and the Sixteen Problems of Particle Physics and Cosmology. Arisaka-GG-W2-Holonomy.

Preprint, available at <https://preprints.arisaka-gg.org/w2-holonomy/>.
doi:10.5281/zenodo.22064342.

Relevance. The classification of protection mechanisms that this paper occupies one rung of. The compact angle the Kaxion is, whose non-locality is the whole reason the mass stays small, 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 this paper repeats — protection by topology rather than by symmetry — is its subject. First cited in §4.1.

Arisaka, K. (2026). The Hubble Tension in the Light of GG-Theory: The Hubble Constant as One Reading of a Derived Expansion History $H(t)$. Arisaka-GG-C2-Hubble.

Preprint, available at <https://preprints.arisaka-gg.org/c2-hubble/>.
doi:10.5281/zenodo.22064336.

Relevance. The Cosmology-Branch sibling on the expansion side of the same π -locked inventory from which this paper’s abundance target descends: the derived H_0 and the Ω bookkeeping of GG-A10-Cosmos are one ledger, read there for the expansion history and here for the dark-matter share that closes f_K . Cited for the branch context of the cohort map. First cited in §4.1.

Arisaka, K. (2026). 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, available at <https://preprints.arisaka-gg.org/c3-bh/>.
doi:10.5281/zenodo.22064338.

Relevance. The Cosmology-Branch sibling at the strong-gravity end: the exact $A/4$ coefficient from the GDOS reading of the horizon. Cited in the cohort map and at the strong-gravity objection (O-Kaxion-53), where it supplies the computed suppression of horizon-scale deviations that makes the Kaxion’s null imprint on EHT and LIGO observables a consistency rather than a constraint. First cited in §4.1.

B. Other references

Aalbers, J., et al. (LUX-ZEPLIN Collaboration) (2024). Dark matter search results from 4.2 tonne-years of exposure of the LUX-ZEPLIN (LZ) experiment. *Physical Review Letters* (and 2025 updates).

Relevance. The world-leading direct-detection search for weakly interacting massive particles: a multi-tonne liquid-xenon time-projection chamber that now excludes spin-independent WIMP–nucleon cross-sections below $\sim 2 \times 10^{-48} \text{ cm}^2$ near 40 GeV with no signal, and has begun to register solar-neutrino backgrounds — the onset of the irreducible “neutrino fog” at which xenon detectors lose discriminating power. It is the empirical squeeze on the WIMP from below in Part IV’s candidate landscape, and the clearest sign that the once-favored candidate is running out of room. First cited in §19.

Abbott, L. F., & Sikivie, P. (1983). A cosmological bound on the invisible axion. *Physics Letters B*, 120, 133–136.

Relevance. A co-founding misalignment paper, which established the scaling of the relic abundance with the decay constant and the initial angle. That scaling — here written $\Omega_K h^2 \propto \theta_i^2 f_K^{1.165}$ — is what makes the Kaxion’s abundance a sharp function of its two inputs, so that demanding the observed density *closes* the decay constant rather than leaving it free (Part V). The reader should note that what is a constraint to be satisfied for the invisible axion becomes, for the Kaxion, the equation that fixes the mass. First cited in §C.6.

Adams, C. B., et al. (2023). Axion dark matter. *Snowmass 2021 white paper*, arXiv:2203.14923.

Relevance. The community white paper summarizing the present status and the coming decade of axion dark-matter searches, including the haloscope roadmaps and sensitivity projections that Part VII’s experiment map draws on. It is the most current single source for the experimental field and the context in which this paper’s reading of the published roadmaps — that the two programs aimed at this band would reach the required depth in the early 2030s if they proceed — should be read. First cited in §36.2.

Knirck, S., for the ADMX Collaboration (2024). ADMX Extended Frequency Range (EFR): looking to the future. Presented at the 19th Patras Workshop on Axions, WIMPs and WISPs, 16–20 September 2024. <https://agenda.infn.it/event/40078/contributions/240683/>

Relevance. The first announced DFSZ-depth program overlapping the Kaxion window from above (8.3–16.5 μeV), using an eighteen-cavity coherently combined array on a 9.4-tesla magnet at Fermilab. Together with CAPP-class coverage from below it brackets the predicted band [1.59, 2.13] GHz, and as a U.S. Department-of-Energy project it is the institutionally natural partner for a U.S. group (Section 36.5). First cited in §36.2.

Ahn, S., et al. (CAPP Collaboration) (2024). Extensive search for axion dark matter over 1 GHz with CAPP’s main axion experiment at DFSZ sensitivity. *Physical Review X*, 14, 031023.

Relevance. Demonstrates DFSZ-depth scanning near 4.55 μeV with a 12-tesla magnet — the instrument class whose extension to 1.6–2.1 GHz would cover the Kaxion window from below. It establishes IBS-CAPP as the best-situated facility for the window (Section 36.4): the magnet already spans the relevant band, and DFSZ depth is already proven. First cited in §36.2.

Ahn, S., et al. (IBS-CAPP) (2026). Extended haloscope search and exclusion of a candidate signal near 1.036 GHz. *Physical Review Letters*, doi:10.1103/2sn2-h97m; *arXiv:2602.05388* (February 2026).

Relevance. The field’s textbook candidate-validation episode, contemporaneous with this paper: a 5.1σ -local excess at 1.036315 GHz (reduced to 3.5σ globally by the look-elsewhere effect) was subjected to persistence, cross-instrument, and field-off tests, failed them, and was excluded at 90% confidence with DFSZ-adjacent depth. Cited three times over: as the live demonstration of the trials-factor arithmetic that pre-registration removes (§37.1), as the vetting discipline the campaign protocol adopts (§37.3), and as the first faced-and-survived instance of falsifier F-Kaxion-1d — the candidate sat outside the unique window and brighter than the DFSZ line, so its confirmation would have killed the Kaxion outright (§37.4). First cited in §36.4.

Lee, S., et al. (CAPP-8TB) (2020). Axion dark matter search around $6.7\,\mu\text{eV}$. *Physical Review Letters*, 124, 101802.

Relevance. The only modern search ever published inside the predicted window: an 8-tesla solenoid swept 1.60–1.65 GHz — 50 of the window’s 544 MHz — and reached $\simeq 4\times$ the KSVZ coupling, which at these masses is 10.2 times the DFSZ line this candidate sits on. It is the measurement that sets what “the window is not untouched ground, and it is not covered either” means quantitatively (Section 36.4). First cited in §36.4.

Wuensch, W. U., et al. (1989). Results of a laboratory search for cosmic axions and other weakly coupled light particles. *Physical Review D*, 40, 3153.

Relevance. The Rochester–Brookhaven–Fermilab pilot, and with UF one of only two searches to have covered most of the predicted window at all. At some 56 times the DFSZ line it is far above the Kaxion, which is precisely the point: the window’s published floor is a pilot-era limit, not a modern one. First cited in §36.4.

Hagmann, C., Sikivie, P., Sullivan, N. S., & Tanner, D. B. (1990). Results from a search for cosmic axions. *Physical Review D*, 42, 1297(R).

Relevance. The University of Florida pilot, covering the window’s middle at some 28 times the DFSZ line — the deepest of the two historical scans, and still an order of magnitude and a half above the prediction. First cited in §36.4.

Yang, J., Gleason, J. R., Jois, S., Stern, I., Sikivie, P., Sullivan, N. S., & Tanner, D. B. (2020). Search for $5\text{--}9\,\mu\text{eV}$ axions with ADMX four-cavity array. In *Microwave Cavities and Detectors for Axion Research*, Springer Proceedings in Physics 245, 53–62.

Relevance. The only announced instrument designed for the band that contains the window: $5\text{--}9\,\mu\text{eV}$ is 1.21–2.18 GHz, which contains all seven admissible addresses and retires six of them. It is a design and prototype study — it states neither a target sensitivity nor a schedule — which is why this paper attaches no date to F-Kaxion-1c (Section 36.9). First cited in §36.5.

Benakli, K., Goodsell, M. D., & Goudeau, A. (2026). An anomaly-free baryonic completion of the Standard Model. *arXiv:2607.12009*.

Relevance. The orthodox four-dimensional answer to the question this cohort answers geometrically — what does it cost to gauge baryon number? — and therefore the sharpest external contrast available to the whole gauged- B sector. They cancel the anomaly by adding a secluded chiral sector and give the baryonic vector a mass by the Higgs mechanism; the construction here cancels the same anomaly by BI-6 inflow, adds no matter, and removes the zero mode by the Dirichlet clause of its admissible class. Their dark matter is a Majorana state of the secluded sector at 0.1–1 TeV; ours is the compact-angle pseudoscalar at $8.1\,\mu\text{eV}$. Cited in the candidate landscape (Section 19). First cited in §19.

Altenmüller, K., et al. (CAST Collaboration) (2024). New upper limit on the axion–photon coupling with an extended CAST run with a Xe-based Micromegas detector. *Physical Review Letters*, 133, 221005.

Relevance. The strongest helioscope bound on the axion–photon coupling, obtained by pointing a magnet at the Sun and searching for X-rays from axions produced in the solar core. It does not assume axions are the dark matter, so it provides a model-independent context ceiling in the detection map of Part VII — territory adjacent to, but well above, the Kaxion line, illustrating the gap a dedicated dark-matter haloscope must close. First cited in §36.2.

Aprile, E., et al. (XENON Collaboration) (2023). First dark matter search with nuclear recoils from the XENONnT experiment. *Physical Review Letters*, 131, 041003.

Relevance. Independent tonne-scale liquid-xenon corroboration of the WIMP null at comparable depth to LZ. Together the two experiments (with PandaX-4T) make the absence of a WIMP signal robust across independent detectors, which is the basis for treating the WIMP as empirically squeezed rather than merely unconfirmed in Part IV. First cited in §19.

ATLAS and CMS Collaborations (2024). Summary of supersymmetry searches at the LHC (Run 2–3). *ATLAS/CMS public summary plots; Particle Data Group (2024)*.

Relevance. The collider exclusion of weak-scale supersymmetry — gluinos and first-generation squarks above ~ 2 TeV, electroweakinos near 1 TeV — with no superpartner found. Since supersymmetry is the WIMP’s natural ultraviolet home (the lightest neutralino being the canonical WIMP), these searches remove the theoretical motivation that the direct-detection nulls remove empirically. In Part IV this is half of the argument that the WIMP is being squeezed from both ends, and the structural reason GG-Theory (which admits no supersymmetry) forbids a WIMP at all. First cited in §19.

Aybas, D., et al. (2021). Search for axionlike dark matter using solid-state nuclear magnetic resonance. *Physical Review Letters*, 126, 141802.

Relevance. The nuclear-magnetic-resonance search class, sensitive to the axion’s coupling to nucleon spin rather than to photons, in which an oscillating axion field tips polarized nuclei when their Larmor frequency matches the axion mass. It is the instrument family from which the Kaxion’s baryonic-channel readout descends, and the starting point for the non-Larmor (internal-field) nucleon-port concept of Part VIII that overcomes the impractical field a free-spin resonance at 1.95 GHz would require. First cited in §40.

Bai, X., et al. (HAYSTAC Collaboration) (2025). Dark matter axion search with HAYSTAC Phase II. *Physical Review Letters*, 134, 151006.

Relevance. The first haloscope to scan below the standard quantum limit, using a vacuum-squeezed two-quadrature receiver to suppress amplifier noise and accelerate the scan. Its published 16.96–19.46 μeV band lies above the Kaxion’s laboratory line at 8.1 μeV (1.95 GHz), so HAYSTAC is a technology partner rather than a discovery instrument for the Kaxion — but its squeezed-state readout is the template for the quantum-limited identification instruments of Part VIII. First cited in §36.2.

Bertone, G., & Hooper, D. (2018). History of dark matter. *Reviews of Modern Physics*, 90, 045002.

Relevance. The authoritative history of the problem, from the nineteenth-century “dark stars” through Zwicky, Rubin, the rise of the WIMP, and the turn toward axions. It is recommended reading alongside this paper because it makes vivid the central sociological fact Part IV confronts: for four decades the search was organized around one fashionable candidate, and the field is only now returning to its widest question — not where dark matter is, but what it is. First cited in §1.

Bertone, G., Hooper, D., & Silk, J. (2005). Particle dark matter: evidence, candidates and constraints. *Physics Reports*, 405, 279–390.

Relevance. The standard graduate-level review of particle dark matter and the single best place to learn the field’s structure quickly: it organizes the gravitational evidence, the leading candidates (WIMPs, axions, sterile neutrinos, and more), and the detection strategies (direct, indirect, collider) into one framework. Part IV’s candidate landscape can be read as an update of this review through 2026, and its structural-comparison table as a re-scoring of these same candidates by the question Bertone–Hooper–Silk did not ask: is the candidate *forced*? First cited in §1.1.

Bo, Z., et al. (PandaX Collaboration) (2024). Dark matter search results from 1.54 tonne-year exposure of PandaX-4T. *Physical Review Letters*.

Relevance. The third independent multi-tonne xenon program corroborating the WIMP null. The agreement among LZ, XENONnT, and PandaX-4T across different detectors and analysis pipelines is what lets Part IV state the WIMP squeeze as an experimental fact rather than a single-experiment limit. First cited in §19.

Boddy, K. K., & Gluscevic, V. (2018). First cosmological constraint on the effective theory of dark matter–proton interactions. *Physical Review D*, 98, 083510.

Relevance. The effective-field-theory framework for dark-matter–baryon scattering, and the cosmological bounds on its cross-section as a function of velocity dependence. It supplies the language — a velocity- and spin-dependent $\sigma_{\chi b}$ — in which the present paper’s Part XI computation of the Kaxion’s dark–baryon portal magnitude is cast, and the benchmark against which that computed cross-section is found astrophysically negligible. First cited in §22.5.

Boldrini, P. (2022). The cusp–core problem in gas-poor dwarf spheroidal galaxies. *Galaxies*, 10, 5.

Relevance. Establishes the energy floor of the baryonic-feedback solution: supernova-driven coring requires a stellar mass above roughly $10^6 M_\odot$ to supply the energy needed to flatten a cusp, so the gas-poor classical and ultra-faint dwarfs — the most dark-matter-dominated systems — are expected to retain cusps. This identifies the feedback-starved regime as the cleanest test of any dark-sector core mechanism, and it is exactly the regime in which the Kaxion’s (computed-negligible) dark–baryon portal would have had to act. First cited in §22.2.

Borsanyi, S., et al. (2016). Calculation of the axion mass based on high-temperature lattice quantum chromodynamics. *Nature*, 539, 69–71.

Relevance. The benchmark lattice determination of the QCD topological susceptibility $\chi_{\text{QCD}}(T)$ across the temperatures relevant to axion production. It fixes both the zero-temperature normalization that sets $m_K f_K = \sqrt{\chi_{\text{QCD}}}$ and the temperature dependence that produces the $f_K^{1.165}$

exponent in the Kaxion mass and abundance formulae of Part V. For a student it is the concrete demonstration that the once-uncertain axion mass is now a controlled lattice quantity, which is exactly why a sharp frequency prediction is possible at all. First cited in §9.5.

Boyarsky, A., Drewes, M., Lasserre, T., Mertens, S., & Ruchayskiy, O. (2019). Sterile neutrino dark matter. *Progress in Particle and Nuclear Physics*, 104, 1–45.

Relevance. The standard review of keV-scale sterile-neutrino dark matter, the production mechanisms, the unidentified 3.5 keV X-ray line once read as its decay signature, and the laboratory (KATRIN/TRISTAN) probes of the active–sterile mixing. It is the basis for the sterile-neutrino entry in Part IV’s landscape, where the candidate is shown to be disfavored both observationally (the line largely unconfirmed, warm dark matter in tension with structure) and structurally (no spare gauge-singlet in the fixed GG-Theory lepton sector). First cited in §19.

Boylan-Kolchin, M. (2023). Stress-testing Λ CDM with high-redshift galaxy candidates. *Nature Astronomy*, 7, 731–735.

Relevance. The sharpest early statement of the JWST tension: the inferred stellar-mass density of the first massive galaxy candidates approaches, and for the most extreme objects formally exceeds, what is available even if every baryon in the corresponding dark-matter halos turned into stars. It frames the early-structure problem of Part IV quantitatively, and it is the pressure point at which the *direction* of the surprise — a demand for more, not less, early structure — becomes a discriminant that favors cold over suppressing dark matter. First cited in §21.

Brouwer, L., et al. (DMRadio Collaboration) (2022). Projected sensitivity of DMRadio-m³: a search for the QCD axion below 1 μ eV. *Physical Review D*, 106, 103008.

Relevance. The lumped-element program targeting DFSZ sensitivity below 1 μ eV, where cavities become impractically large. While the Kaxion line lies above this band (at 8.1 μ eV), the program illustrates the lumped-element technique and would have hosted the nucleon-spin booster of the withdrawn spin-port concept (Appendix F), showing how the haloscope idea extends to frequencies below the microwave band. First cited in §36.2.

Bullock, J. S., & Boylan-Kolchin, M. (2017). Small-scale challenges to the Λ CDM paradigm. *Annual Review of Astronomy and Astrophysics*, 55, 343–387.

Relevance. The standard review of the small-scale “crises” of cold dark matter — the cusp–core problem, the diversity of rotation curves, missing satellites, and too-big-to-fail — and of the competing baryonic and dark-sector explanations. It is the map of the terrain that Part IV’s astrophysical-significance sections navigate, and the reference a student should read to understand which small-scale puzzles are genuinely about dark-matter microphysics and which are now attributed to ordinary galaxy-formation physics. First cited in §22.1.

Carney, D., Krnjaic, G., Moore, D. C., et al. (2021). Mechanical quantum sensing in the search for dark matter. *Quantum Science and Technology*, 6, 024002.

Relevance. The review of optomechanical and levitated-sensor approaches to detecting dark-matter-induced forces, in which a mechanical resonator registers a tiny oscillating acceleration rather than a spin precession or a photon. It is the platform for the magnet-free, spin-free

baryon-number-force port, Concept B''' of Part VIII — the conceptually cleanest reading of the gauged baryon coupling, even though the present paper shows that coupling to be astrophysically inert. First cited in §F.2.4.

Carniani, S., et al. (2024). A shining cosmic dawn: spectroscopic confirmation of two luminous galaxies at $z \simeq 14$. *Nature* (and arXiv:2405.18485).

Relevance. The spectroscopic confirmation of JADES-GS-z14-0 at redshift 14.32 — under 300 million years after the Big Bang — with some $5 \times 10^8 M_\odot$ already in stars. It is the secure anchor of the JWST early-structure surprise of Part IV: a luminous, massive galaxy existing far earlier than the straightforward expectation, and a datum that cannot be dismissed as a photometric artifact. It makes concrete the demand for an early head start in structure formation that the Kaxion (cold, with an inherited cosmological boost) satisfies and that suppressing candidates worsen. First cited in §21.2.

Carosi, G., et al. (ADMX Collaboration) (2025). Search for axion dark matter from 1.1 to 1.3 GHz with ADMX. *Physical Review Letters*, 135, 191001.

Relevance. Extends the ADMX program upward to $4.5\text{--}5.4\,\mu\text{eV}$ at near-DFSZ depth — the published coverage closest to the lower edge of the Kaxion window — and demonstrates the cavity and readout technology that the announced programs would carry into the window itself. First cited in §36.2.

Clowe, D., Bradač, M., Gonzalez, A. H., et al. (2006). A direct empirical proof of the existence of dark matter. *The Astrophysical Journal Letters*, 648, L109–L113.

Relevance. The Bullet Cluster: a collision of two galaxy clusters in which weak-lensing maps the total mass and X-rays map the hot gas (the dominant *baryonic* component), and the two are spatially separated. Because the bulk of the gravitating mass moved through the collision collisionlessly, ahead of the gas, the observation is the cleanest empirical demonstration that dark matter is real *matter* — collisionless and non-baryonic — rather than a modification of gravity. It is the experiment that makes “what is the particle?” the right question, which is the question the Kaxion answers. First cited in §1.2.

Crescini, N., et al. (QUAX Collaboration) (2020). Axion search with a quantum-limited ferromagnetic haloscope. *Physical Review Letters*, 124, 171801.

Relevance. The mature electron-spin (ferromagnetic-magnon) haloscope, in which the axion’s coupling to electron spin drives a magnon mode in a ferrite sphere read out through a microwave cavity at sub-tesla field and gigahertz frequency. It is the demonstrated technology behind the archived harmonic-booster electron port of Part VIII — kept there as an engineering record, its target withdrawn with the doublet. First cited in §F.2.2.

Crisosto, N., Sikivie, P., Sullivan, N. S., Tanner, D. B., Yang, J., & Rybka, G. (2020). ADMX SLIC: Results from a superconducting LC-circuit investigating cold axions. *Physical Review Letters*, 124, 241101.

Relevance. A lumped-element (LC-circuit) pathfinder that tuned across $0.176\text{--}0.203\,\mu\text{eV}$, demonstrating DFSZ-class lumped-element detection well below the cavity regime. The Kaxion line

itself sits in the cavity band at 1.95 GHz, above this range; SLIC illustrates the complementary lumped-element technique for lighter axions. First cited in §36.2.

Dekel, A., Sarkar, K. C., Birnboim, Y., Mandelker, N., & Li, Z. (2023). Efficient formation of massive galaxies at cosmic dawn by feedback-free starbursts. *Monthly Notices of the Royal Astronomical Society*, 523, 3201–3218.

Relevance. The leading *baryonic* relief of the JWST tension: at the high gas densities of cosmic dawn, star formation can proceed before stellar feedback has time to act, raising the conversion efficiency and producing bright galaxies without any change to cosmology. The present paper’s position (Part IV) is that this astrophysics and the Kaxion’s inherited cosmological head start *combine* rather than compete — and that the essential dark-matter requirement is simply not to *suppress* the early halos, which the Kaxion, being cold, does not. First cited in §21.3.

Dine, M., & Fischler, W. (1983). The not-so-harmless axion. *Physics Letters B*, 120, 137–141.

Relevance. The third founding misalignment paper, which emphasized that an axion with too large a decay constant overcloses the universe — the cosmological danger that bounds the parameter space from above. For the Kaxion this danger is inverted into a virtue: because the abundance is pinned to the derived Ω_{DM} , the same physics that threatens overclosure for a generic axion is what makes the Kaxion’s decay constant, and therefore its mass and frequency, predictive rather than merely permissive. First cited in §C.6.

Fischer, M. S., et al. (2024). Cosmological and idealized simulations of dark-matter halos with velocity-dependent, rare and frequent self-interactions. *Monthly Notices of the Royal Astronomical Society*, 529, 2327.

Relevance. Modern simulations showing that velocity-dependent SIDM, through early core expansion followed by gravothermal core-collapse, can generate the observed rotation-curve diversity. It is the state of the art for the SIDM diversity mechanism, and the basis for one of the present paper’s clean discriminators: the Kaxion, having no dark–dark scattering, predicts *no* gravothermal core-collapse, so a confirmed population of core-collapsed dwarfs would favor SIDM over the Kaxion. First cited in §22.2.

Goodman, C., et al. (ADMX Collaboration) (2025). Axion Dark Matter eXperiment around 3.3 μeV with DFSZ discovery ability. *Physical Review Letters*, 134, 111002.

Relevance. The flagship demonstration that DFSZ-level sensitivity — the community’s hard target, and the depth the Kaxion requires — is achievable in the microelectron-volt decade with a tunable cavity, a strong solenoid, and a quantum-limited amplifier. It is the benchmark against which the Kaxion’s DFSZ-line coupling and the scan-time requirements of Part VIII are calibrated, and the proof of principle that instruments of this class can reach the depth the prediction demands. First cited in §36.2.

Dodelson, S., & Schmidt, F. (2020). *Modern Cosmology*, 2nd ed. Academic Press.

Relevance. The standard graduate textbook on cosmological perturbation theory and structure formation: the growth of density perturbations through radiation and matter domination, the transfer function, and the linear-to-nonlinear transition. It is the reference behind the “why cold

dark matter is the scaffold” argument (Appendix B.5.6) and behind the epoch-dependence of the π -locked fractions (Appendix C.7); a student following the abundance and structure-formation discussion should keep it open. First cited in §C.7.

Gorghetto, M., & Villadoro, G. (2019). Topological susceptibility and QCD axion mass: QED and NNLO corrections. *Journal of High Energy Physics*, 2019(3), 033.

Relevance. Sharpens the temperature-dependent topological susceptibility $\chi(T)$ and the misalignment transport beyond leading order, including electromagnetic and next-to-next-to-leading-order corrections. These are the dominant residual systematics in any axion mass prediction, and the paper is used in Part V to bound the residual of our integrated abundance pipeline against the precision literature — a residual that enters the closure angle, not the mass. First cited in §25.

Green, A. M., & Kavanagh, B. J. (2021). Primordial black holes as a dark matter candidate. *Journal of Physics G*, 48, 043001.

Relevance. The standard review of primordial black holes as dark matter, compiling the lensing, evaporation, accretion, and dynamical constraints that confine them, as a dominant component, to the narrow “asteroid-mass” window $\sim 10^{17}$ – 10^{22} g. It is the basis for the primordial-black-hole entry in Part IV, and a reminder that even the surviving window requires a tuned, enhanced primordial power spectrum — a freedom the fixed inflationary history of GG-Theory does not have. First cited in §19.

Green, M. B., & Schwarz, J. H. (1984). Anomaly cancellation in supersymmetric $D = 10$ gauge theory and superstring theory. *Physics Letters B*, 149, 117–122.

Relevance. The original Green–Schwarz mechanism, in which a gauge and gravitational anomaly is canceled by the inflow from a single antisymmetric tensor field whose Bianchi identity is modified. BI-6 is the constitutional descendant of this mechanism in six dimensions, and the Kaxion is, in the same sense, the descendant of the model-independent string axion that the Green–Schwarz two-form always carries. The reader should recognize that the Kaxion’s existence rests on the same anomaly-inflow logic that launched the first superstring revolution. First cited in §6.3.

Grilli di Cortona, G., Hardy, E., Pardo Vega, J., & Villadoro, G. (2016). The QCD axion, precisely. *Journal of High Energy Physics*, 2016(1), 034.

Relevance. The precision next-to-leading-order chiral computation of the QCD-coupled axion’s properties, and the modern reference for its mass and couplings. Two of its results enter the Kaxion verbatim: the mass relation $m_a = 5.70(7) \mu\text{eV} (10^{12} \text{ GeV}/f_a)$ that, combined with the abundance closure, fixes the Kaxion’s mass, and the model-independent chiral term $-1.92(4)$ in the photon coupling that places the DFSZ line (and hence the Kaxion line) exactly where it is. The paper also fixes the nucleon couplings C_p, C_n that set the chiral input fenced in the Kaxion’s baryonic-channel discussion, and it anchors the cross-validation of our misalignment pipeline. First cited in §24.

Hosotani, Y. (1983). Dynamical mass generation by compact extra dimensions. *Physics Letters B*, 126, 309–313.

Relevance. Introduced the Wilson-line (holonomy) mechanism by which a gauge symmetry on a compact extra dimension is broken not by a Higgs field but by the phase a gauge field accumulates around the circle. In GG-6 the same holonomies that drive the Hosotani breaking of SU(5) are the circle’s lifted angles — the holonomy “echo” of Part II, the sibling directions the count removes — so this paper is the origin of what the Kaxion is *not*, and a clean example of how non-spatial geometry does physical work. First cited in §6.5.

Hu, W., Barkana, R., & Gruzinov, A. (2000). Fuzzy cold dark matter: the wave properties of ultralight particles. *Physical Review Letters*, 85, 1158–1161.

Relevance. Defined the fuzzy-dark-matter regime, in which an ultralight boson ($m \sim 10^{-22}$ eV) has a de Broglie wavelength of kiloparsec scale, so that quantum pressure suppresses small-scale structure and smooths galactic cores. It is the candidate from which the Kaxion is most sharply distinguished: fifteen-plus orders of magnitude heavier, the Kaxion has a meter-scale de Broglie wavelength and behaves as ordinary cold dark matter on every astrophysical scale. The “not-fuzzy” fence of Parts III–IV rests on this contrast, and the recent observational exclusion of the canonical fuzzy window makes the contrast favor the Kaxion. First cited in §16.

Irastorza, I. G., & Redondo, J. (2018). New experimental approaches in the search for axion-like particles. *Progress in Particle and Nuclear Physics*, 102, 89–159.

Relevance. The comprehensive review of the experimental landscape — haloscopes, helioscopes, nuclear-magnetic-resonance searches, dielectric and plasma concepts, and light-shining-through-walls — and the recommended map of who searches where and how. It underlies the detection map of Part VII and the instrument concepts of Part VIII, and it lets a newcomer place the programs that bracket the Kaxion window (CAPP, running; ADMX-EFR, designed) within the full experimental ecosystem. First cited in §36.2.

Kaplinghat, M., Tulin, S., & Yu, H.-B. (2016). Dark matter halos as particle colliders: unified solution to small-scale structure puzzles from dwarfs to clusters. *Physical Review Letters*, 116, 041302.

Relevance. The benchmark demonstration that a single velocity-dependent SIDM cross-section can fit rotation curves from dwarfs (where $\sigma/m \sim$ a few cm^2/g) to clusters (where it must fall to satisfy merger bounds), with the core size set jointly by the self-interaction and the baryon distribution. It is the quantitative SIDM model against which the Kaxion’s parameter-free dark-baryon lever is contrasted in Part IV, and a clear illustration of the price SIDM pays — a tuned function — that the Kaxion does not. First cited in §22.2.

Kim, A., et al. (IBS-CAPP) (2024). Extensive search for axion dark matter over 1 GHz with CAPP’s Main Axion eXperiment. *Physical Review X*, 14, 031023.

Relevance. The flagship CAPP-MAX result: the first axion search at or near DFSZ sensitivity over more than 100 MHz above 1 GHz, achieved at a record DFSZ scan speed of about 3 MHz/day with a 12-tesla magnet, dilution-refrigerator cooling, and Josephson-parametric-amplifier readout. It is the empirical basis for identifying IBS-CAPP as the best-situated instrument for the Kaxion window (Section 36.4): the depth, the magnet band, and the scan speed needed are all demonstrated. First cited in §36.4.

Kim, J. E., & Carosi, G. (2010). Axions and the strong CP problem. *Reviews of Modern Physics*, 82, 557–602.

Relevance. The standard modern review of axion theory and phenomenology, and a good companion to Marsh for the particle-physics (as opposed to cosmological) side. It defines the two canonical benchmark models — KSVZ and DFSZ — whose photon couplings bracket the “axion band,” and against which the Kaxion’s position (exactly on the DFSZ line, but inside a pre-declared window and standing alone) is contrasted in Parts IV and VI. The reader will find here the model-building vocabulary the present paper uses without redefining. First cited in §27.

Kolb, E. W., & Tkachev, I. I. (1993). Axion miniclusters and Bose stars. *Physical Review Letters*, 71, 3051–3054.

Relevance. The founding minicluster paper. In a post-inflationary history the axion field takes independent random values in causally disconnected patches, and the resulting order-unity density inhomogeneities collapse at matter–radiation equality into dense, Earth-mass-scale clumps. This is the small-scale substructure channel by which a post-inflationary Kaxion population could seed early structure (Part IV) — a contingency the present paper analyzes but ultimately fences off, because the series’ single-patch history forecloses it. First cited in §21.10.

Kolb, E. W., & Tkachev, I. I. (1994). Large-amplitude isothermal fluctuations and high-density axion clumps. *Physical Review D*, 49, 5040–5051.

Relevance. Quantifies the nonlinear gravitational collapse of large-amplitude axion inhomogeneities into high-density clumps, giving the characteristic minicluster masses and densities. It is the companion input to the minicluster discussion of Part IV, and a useful illustration for the student of how a field that is perfectly smooth in the pre-inflationary picture can be highly clumped in the post-inflationary one — a distinction that, for the Kaxion, is tied to which cosmological window the prediction lands in. First cited in §21.10.

Labbé, I., et al. (2023). A population of red candidate massive galaxies ~ 600 Myr after the Big Bang. *Nature*, 616, 266–269.

Relevance. The discovery of a population of surprisingly massive galaxy candidates ($M_\star \sim 10^{10}$ – $10^{11} M_\odot$) at $z \sim 7$ – 10 , whose inferred stellar-mass density drives the “impossible early galaxies” tension — the requirement of a near-unity baryon-to-star conversion efficiency. It is the observational starting point of the early-structure discussion of Part IV, and the reason the direction of the JWST surprise (toward more early structure) became a dark-matter discriminant. First cited in §21.2.

Lawson, M., Millar, A. J., Pancaldi, M., Vitagliano, E., & Wilczek, F. (2019). Tunable axion plasma haloscopes. *Physical Review Letters*, 123, 141802.

Relevance. The plasma-haloscope concept, which uses a tunable wire-array metamaterial to decouple the resonant volume from the photon wavelength, enabling large detection volumes at high mass. It is a candidate platform for wide-band instruments in and above the Kaxion window, and an example of the instrumentation now being developed to overcome the volume–frequency tension that limits conventional cavities. First cited in §36.2.

Kolb, E. W., & Turner, M. S. (1990). *The Early Universe*. Addison-Wesley.

Relevance. The classic graduate text on early-universe thermodynamics and relic abundances: the Friedmann background, entropy conservation, freeze-out, and — directly relevant here — the vacuum-misalignment production of a coherent scalar condensate. Its treatment of the axion relic and the effective degrees of freedom $g_*(T)$ is the backbone of the misalignment course (Appendix C.6) and the thermal-history inputs of the closure audit. First cited in §C.7.

Lee, B. W., & Weinberg, S. (1977). Cosmological lower bound on heavy-neutrino masses. *Physical Review Letters*, 39, 165–168.

Relevance. The foundational thermal-relic (freeze-out) calculation behind the “WIMP miracle”: a stable particle that annihilates with a weak-scale cross-section freezes out of equilibrium with roughly the observed dark-matter abundance. Understanding this argument is essential to understanding why the WIMP dominated the field for thirty years — and, by contrast, why the Kaxion’s abundance comes not from freeze-out but from vacuum misalignment, a production mechanism (see Preskill–Wise–Wilczek) that requires no thermal equilibrium at all. First cited in §1.1.

MADMAX Collaboration (2024). The dielectric haloscope MADMAX for 40–400 μeV axion dark matter. arXiv:1611.04549 and updates.

Relevance. The dielectric-stack haloscope concept, which boosts the axion-to-photon conversion using many dielectric interfaces and so reaches masses above the practical cavity regime. It bounds the high-mass side of the experimental landscape in Figure 8, well above the Kaxion corridor, and rounds out the reader’s picture of how different mass decades demand different instrument concepts. First cited in §36.2.

Marsh, D. J. E. (2016). Axion cosmology. *Physics Reports*, 643, 1–79.

Relevance. The comprehensive pedagogical review of axion cosmology and the recommended single source for a graduate student entering this field: it treats the misalignment mechanism, isocurvature constraints, structure formation, miniclusters, the pre- versus post-inflationary histories, and the ultralight (fuzzy) regime in one consistent framework. Nearly every cosmological tool the present paper uses — the abundance integral, the isocurvature bound, the minicluster contingency, the de Broglie/Jeans scales — is laid out here in tutorial form, and reading it makes Parts III and IV substantially easier to follow. First cited in §75.

Nakahara, M. (2003). *Geometry, Topology and Physics*, 2nd ed. Institute of Physics Publishing.

Relevance. The standard physicist’s reference for differential forms, Hodge duality, fiber bundles, characteristic classes, and cohomology — exactly the toolkit behind the geometric courses of this Part. The p -form/dual-scalar equivalence (Appendix C.1), the flux quantization over cycles that forces Ξ_{cap} into an integer (Appendix C.5), and the Hopf fibration behind the π -locked inventory (Appendix C.7) are all developed here at book length. First cited in §C.4.

Navarro, J. F., Frenk, C. S., & White, S. D. M. (1997). A universal density profile from hierarchical clustering. *The Astrophysical Journal*, 490, 493–508.

Relevance. The Navarro–Frenk–White (NFW) profile: cold-dark-matter-only simulations find that halos of every mass relax to a universal, centrally cuspy density profile $\rho \propto r^{-1}$ in the center. This is the theoretical baseline against which the observed cores of dwarf galaxies are compared — the “cusp” of the cusp–core problem of Part IV — and, with the present paper’s computed result that the Kaxion’s dark-baryon portal cannot core a halo, it is precisely the profile the Kaxion predicts for real galaxies. First cited in §22.1.

O’Hare, C. (2020). AxionLimits: cajohare/AxionLimits. *Zenodo* / *GitHub*, <https://cajohare.github.io/AxionLimits/> (data compilation, April 2026 update).

Relevance. The community-standard compilation of published axion exclusion limits, from which the twenty-seven haloscope envelopes of Figure 8 are digitized (lower-envelope decimation; individual collaboration publications cited separately where discussed). The same compilation underlies the landscape figures of the experimental collaborations themselves, so the reader can compare this paper’s map against the field’s directly. First cited in §36.2.

Oman, K. A., et al. (2015). The unexpected diversity of dwarf galaxy rotation curves. *Monthly Notices of the Royal Astronomical Society*, 452, 3650–3665.

Relevance. The modern, sharp form of the cusp–core problem: at a fixed outer (maximum) circular velocity, the inner circular velocity of dwarf galaxies scatters by about a factor of four, from cuspy NFW-like systems to strongly cored ones, a spread that collisionless cold dark matter does not reproduce. This is the diversity that any successful theory of galaxy centers must explain, and the benchmark the cusp–core section of Part IV must confront; the Kaxion’s answer is that the diversity is baryonic, not a dark-matter signature. First cited in §22.1.

Peccei, R. D., & Quinn, H. R. (1977). CP conservation in the presence of pseudoparticles. *Physical Review Letters*, 38, 1440–1443.

Relevance. The original dynamical solution of the strong-CP problem: a new global U(1) symmetry whose spontaneous breaking lets the QCD vacuum angle $\bar{\theta}$ relax dynamically to zero, explaining why the neutron electric dipole moment is so small. Every axion model since builds on this idea. The Kaxion realizes the same relaxation *geometrically* — the compact-angle holonomy plays the role of the postulated global phase — as mandated by GG-A1-Constitution Problem P7, which is why it solves strong CP without adding a Peccei–Quinn scalar and without the attendant quality problem. First cited in §1.1.

Peskin, M. E., & Schroeder, D. V. (1995). *An Introduction to Quantum Field Theory*. Westview Press.

Relevance. The standard first-year graduate quantum-field-theory text: the chiral anomaly, the theta term, spontaneous symmetry breaking, and the pseudo-Nambu–Goldstone mechanism. It is the background for the gluon-portal and mass-relation course (Appendix C.3) and for the anomaly-coefficient computation behind the baryonic ledger entry (Appendix C.4). First cited in §C.4.

Planck Collaboration; Aghanim, N., et al. (2020). Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics*, 641, A6.

Relevance. The precision cosmic-microwave-background determination of the cosmic inventory, including $\Omega_{\text{DM}}h^2 \simeq 0.120$ to about one percent, and the tight bound on dark-matter isocurvature that the Kaxion must respect. This is the measurement the theory’s π -locked target $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ is tested against, and the source of the isocurvature constraint that the low-scale-inflation history of Part III clears by thirty-eight orders of magnitude. For a student it is also the canonical demonstration that dark matter’s *amount* is now known far better than its *nature*. First cited in §1.1.

Preskill, J., Wise, M. B., & Wilczek, F. (1983). Cosmology of the invisible axion. *Physics Letters B*, 120, 127–132.

Relevance. A founding paper of the misalignment mechanism (with Abbott & Sikivie, 1983; Dine & Fischler, 1983), and the key to how a particle so weakly coupled can be all of the dark matter. A periodic field frozen at a random initial angle in the early universe begins to oscillate when the Hubble rate drops below its mass, and the energy in those coherent oscillations redshifts exactly like cold, pressureless matter. This is the production channel by which the Kaxion — never in thermal equilibrium, never relativistic — arrives cold and saturates the π -locked abundance of Parts III and V; the trio of 1983 papers is the conceptual foundation a reader must have to follow the abundance calculation. First cited in §C.6.

Press, W. H., & Schechter, P. (1974). Formation of galaxies and clusters of galaxies by self-similar gravitational condensation. *The Astrophysical Journal*, 187, 425–438.

Relevance. The foundational analytic theory of the halo mass function: the abundance of collapsed objects above a mass M at redshift z follows from the statistics of Gaussian density fluctuations crossing a collapse threshold, with the rare, massive tail falling exponentially. This is the machinery behind the JWST early-structure discussion of Part IV — it is why a one- or two-percent change in the small-scale variance, or a shift in the collapse threshold, produces an order-of-magnitude change in the abundance of the earliest massive galaxies, and therefore why early structure is so sharp a discriminant among dark-matter candidates. First cited in §21.2.

(Representative analyses, 2024.) High-redshift JWST ultraviolet-luminosity-function constraints on warm and fuzzy dark matter (e.g. ApJ and PRD analyses, 2024).

Relevance. The high-redshift ultraviolet luminosity function, populated by JWST, sets one-sided lower bounds on any dark matter that suppresses small-scale structure: $m_{\text{FDM}} \gtrsim 5.6 \times 10^{-22}$ eV for fuzzy and $m_{\text{WDM}} \gtrsim 1.5\text{--}3.2$ keV for warm dark matter. These bounds disfavor exactly the suppressing candidates and are cleared by the Kaxion by some sixteen orders of magnitude — the quantitative core of the “direction matters” argument of Part IV. First cited in §21.4.

Rogers, K. K., & Peiris, H. V. (2021). Strong bound on canonical ultralight axion dark matter from the Lyman- α forest. *Physical Review Letters*, 126, 071302.

Relevance. Uses the small-scale clustering of intergalactic hydrogen (the Lyman- α forest) to exclude fuzzy dark matter below $\sim 2 \times 10^{-20}$ eV, closing the canonical 10^{-22} eV window by about two orders of magnitude. With the JWST bounds it is the observational reason the fuzzy candidate is now disfavored — a candidate GG-Theory also excludes structurally — so the data and the framework again point the same way (Part IV). First cited in §19.

Rubin, V. C., & Ford, W. K. (1970). Rotation of the Andromeda nebula from a spectroscopic survey of emission regions. *The Astrophysical Journal*, 159, 379–403.

Relevance. The flat-rotation-curve evidence that turned dark matter from a cluster-scale anomaly into a galactic universal. Rubin and Ford measured the orbital speeds of gas in Andromeda and found them roughly constant out to large radii, instead of falling as Kepler’s law demands once one leaves the visible disk — the signature of an extended, invisible halo. For a student this is the canonical demonstration that essentially every galaxy is embedded in far more mass than it shows; it anchors the rotation-curve discussion of Part I and frames the cusp–core and diversity puzzles revisited in Part IV. First cited in §1.1.

Saikawa, K., & Shirai, S. (2018). Primordial gravitational waves, precisely: the role of thermodynamics in the Standard Model. *Journal of Cosmology and Astroparticle Physics*, 2018(05), 035.

Relevance. A percent-level tabulation of the Standard-Model relativistic degrees of freedom $g_*(T)$ and $g_{*s}(T)$ across the entire thermal history, including the QCD and electroweak crossovers. This thermal bookkeeping enters the misalignment abundance through the redshifting of the condensate; it is adopted in Part V’s re-anchored pipeline, retiring the transport line of the error budget and removing one source of the historical scatter among axion-mass predictions. First cited in §25.

Sikivie, P. (1983). Experimental tests of the invisible axion. *Physical Review Letters*, 51, 1415–1417.

Relevance. The paper that made the invisible axion detectable: the haloscope and helioscope concepts, in which a strong static magnetic field resonantly converts dark-matter (or solar) axions into microwave photons via the Primakoff effect. This is the detection principle of Part VII and the instrument class — a tunable microwave cavity in a magnet, read by a quantum-limited amplifier — to which the Kaxion’s frequency window is matched. Every haloscope cited below is a descendant of this idea. First cited in §35.

Tremaine, S., & Gunn, J. E. (1979). Dynamical role of light neutral leptons in cosmology. *Physical Review Letters*, 42, 407–410.

Relevance. The phase-space (Tremaine–Gunn) bound: a fermionic dark-matter candidate cannot be packed into a galaxy core more densely than Pauli statistics allow, which sets a lower limit of order a keV on the mass of any fermionic (e.g. sterile-neutrino) dark matter. The argument is part of why warm fermionic candidates are squeezed from below even before laboratory data; it does not constrain the Kaxion, a light *boson* that condenses without any such exclusion limit — a contrast worth holding in mind when reading the candidate exclusions of Part IV. First cited in §19.2.

Trimble, V. (1987). Existence and nature of dark matter in the universe. *Annual Review of Astronomy and Astrophysics*, 25, 425–472.

Relevance. The classic review that consolidated, half a century after Zwicky, the many independent lines of evidence — rotation curves, cluster dynamics, gravitational lensing, and the growth of structure — into the modern consensus that most of the matter in the universe is non-luminous.

A useful first read for a newcomer, because it lays out the evidence *before* any particular candidate is assumed, exactly the candidate-agnostic posture from which Part I begins. First cited in §1.

Tulin, S., & Yu, H.-B. (2018). Dark matter self-interactions and small-scale structure. *Physics Reports*, 730, 1–57.

Relevance. The comprehensive review of self-interacting dark matter (SIDM), the leading particle-physics response to the small-scale puzzles: elastic dark-matter scattering thermalizes halo centers into cores, with a velocity-dependent cross-section σ/m tuned across scales. It is the principal foil for the Kaxion in Part IV: where SIDM introduces a tuned function $\sigma(v)/m$ and produces cores (and, via gravothermal collapse, diversity), the Kaxion has no dark–dark interaction at all and instead carries a baryonic ledger entry whose laboratory-energy remnant the present paper computes to be astrophysically inert. First cited in §22.2.

Turov, E. A., & Petrov, M. P. (1972). *Nuclear Magnetic Resonance in Ferro- and Antiferromagnets*. Halsted Press (Wiley), New York.

Relevance. The standard monograph on nuclear magnetic resonance in magnetically ordered crystals, where the nucleus feels the enormous hyperfine field of the ordered electron spins rather than the applied field, and where the radio-frequency response is amplified by a large enhancement factor. This is the physics behind the internal-field (non-Larmor) nucleon port, Concept B' of Part VIII: it is what lets a nuclear resonance reach the Kaxion frequency of 1.95 GHz at a sub-tesla applied field, sidestepping the 45.8 T a free-spin resonance would require. First cited in §F.2.2.

Vafa, C., & Witten, E. (1984). Parity conservation in quantum chromodynamics. *Physical Review Letters*, 53, 535–536.

Relevance. Proves rigorously that the exact QCD vacuum energy is minimized at vanishing θ -angle. This is the theorem that guarantees a dynamical-relaxation solution actually works: once the Kaxion's gluonic portal (Theorem T-Kaxion-2) makes $\bar{\theta}$ a dynamical field, the field necessarily relaxes to the CP-conserving point rather than to some other value. It is the rigorous backbone beneath the strong-CP closure of Part V, and a reassurance that the Peccei–Quinn idea is not merely plausible but provably lands where it should. First cited in §24.

Weinberg, S. (1978). A new light boson? *Physical Review Letters*, 40, 223–226.

Relevance. Identified (with Wilczek) the light pseudo-Nambu–Goldstone boson that any Peccei–Quinn-style breaking must produce — the axion — and the mass relation $m_a f_a \sim m_\pi f_\pi$ that ties its mass to the QCD scale and inversely to the symmetry-breaking scale. That relation is the direct template for the Kaxion mass formula $m_K^2 f_K^2 \simeq \chi_{\text{QCD}}$ of Part V; the difference is that for the Kaxion the decay constant f_K is not a free symmetry-breaking scale but is closed by the cosmic abundance, turning the mass from a free parameter into a prediction. First cited in §1.1.

Wilczek, F. (1978). Problem of strong P and T invariance in the presence of instantons. *Physical Review Letters*, 40, 279–282.

Relevance. The companion identification of the axion, and the origin of its name. Together with Weinberg's paper it established the kinematics — a periodic pseudoscalar with a model-independent anomalous coupling to gluons and a model-dependent coupling to photons — that

every axion, the Kaxion included, must share. The Kaxion inherits exactly this kinematics while replacing the symmetry-breaking ancestry with a forced compact-angle geometry, which is what lets it carry, in addition, a pre-declared window and a count that an ordinary axion does not. First cited in §1.1.

Zwicky, F. (1933). Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6, 110–127.

Relevance. The founding observation of the dark-matter problem. Applying the virial theorem to the Coma cluster, Zwicky found the dynamical mass needed to bind the galaxies’ velocities exceeds the mass visible in starlight by more than an order of magnitude, and coined the term *dunkle Materie*. The lesson that opens Part I is that dark matter was first inferred not from any property of the matter itself but from a discrepancy between gravity and light — the same epistemic situation we are in ninety years later, and the situation the Kaxion is designed to end by giving the missing mass an identity rather than a weight. First cited in §1.1.

Weinberg, S. (1996). *The Quantum Theory of Fields, Vol. II: Modern Applications*. Cambridge University Press.

Relevance. The authoritative treatment of anomalies, the θ -vacuum, and the Peccei–Quinn mechanism, including the anomaly-matching and current-algebra results used in the photon- and baryon-coupling courses. Its chapters on nonperturbative effects and on broken symmetries underpin the derivations of the $E/N = 8/3$ anomaly ratio and the trace-norm computation of the baryonic ledger entry (Appendices C.3, C.4). First cited in §C.4.