**UCLA/Arisaka-GG-C1-Kaxion**
v20.0 — August 21, 2026 — *Cosmology Branch*

# 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

July 9, 2026

**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 in just the observed amount — about 27 percent of the cosmos. As a bonus, it explains why the strong nuclear force respects mirror symmetry so exactly.

Because nothing about the Kaxion can be tuned, the theory makes hard promises: the precise microwave frequency at which to listen, near 1.95 gigahertz; the faintness of its glow into light; and two fingerprints — every signal is a pair of tones split by one part in thirty-five, in two channels locked five to two. Two experiments now being built to hunt dark matter are aimed almost exactly at this frequency and faintness, and should find the Kaxion — or rule it out — within a few years. 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 fixed by 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 $\pi$: the total matter fraction is $1/\pi$, the ordinary-atom fraction $1/(2\pi^2)$, and their difference,

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

is the dark-matter share — matching what is measured in the Big Bang's afterglow. That is the target the Kaxion must hit. The geometry fixes the one scale that sets how much Kaxion the early universe makes; the observed abundance then fixes the single remaining input — the field's starting tilt — at an utterly ordinary value. The dark-matter density is thus not a coincidence arranged after the fact: with the scale already set by the geometry, the relic has no room left to miss.

**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 long-sought mechanism that 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: $8.1\,\mu$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* — the faintest, hardest target the axion-search community has ever set. Every uncertainty in these numbers traces to laboratory nuclear measurement; the geometry itself contributes none.

**Two fingerprints that cannot be tuned away.** Whatever its exact mass, the geometry stamps the Kaxion with two signatures that carry no adjustable number. First, every Kaxion line is really a close *pair* of tones, split by exactly one part in thirty-five — the spectral memory of how deep the theory's scale dimension runs. Second, it speaks in two channels at once, to ordinary matter and to light, in the fixed ratio five to two. A line with no partner at one part in thirty-five, or a two-channel reading that breaks five-to-two, would refute the Kaxion outright, wherever its mass turned out to lie. These are 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 startlingly early and massive — 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 a modest head start from the same $\pi$-locked ledger. 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 funded experiments — 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 listening for both tones at once, and a small companion magnet reading the second channel, would turn a bare detection into an unambiguous identification, by exposing the one-part-in-thirty-five split and the five-to-two lock. 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 non-spatial angle, the five-to-two lock, its coupling to the strong force, and the coldness of the condensate; the one-part-in-thirty-five split is exact in its symmetric limit. Resting on two named, clearly flagged conjectures: the step that turns the $\pi$-locked inventory into the mass, and the step that pins the laboratory frequency within its narrow allowed window. Both are stated as conjectures wherever they are used, and both are 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 fixed, to the cycle, by five whole numbers and the geometry of a dimension that is not a place but an angle. 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; $\Omega_{\mathrm{DM}} = 1/\pi - 1/(2\pi^2)$; misalignment; haloscope; twin-peak spectrum $\Delta\nu/\nu = 1/35$; 5:2 portal lock; ADMX; HAYSTAC.

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

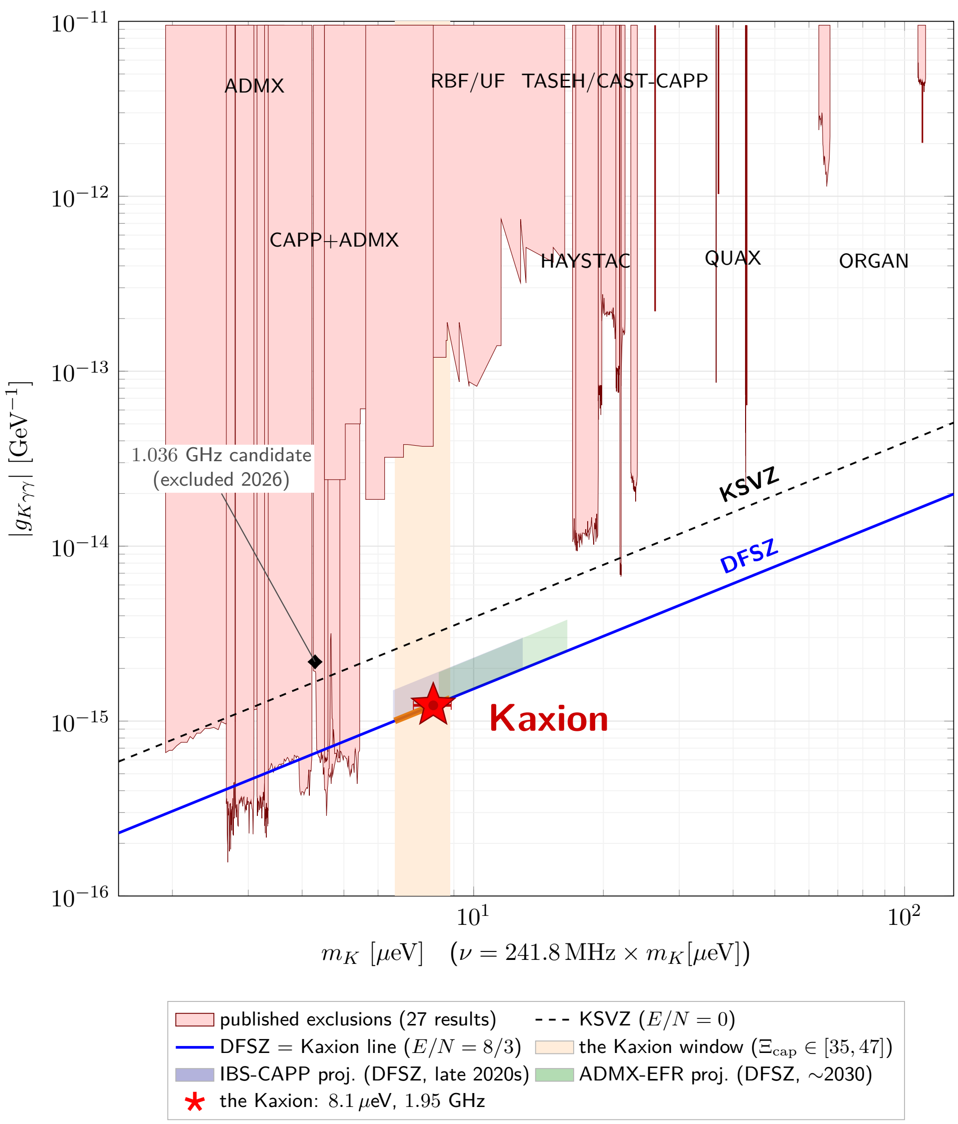


**Figure 1: The Kaxion against the haloscope landscape**: twenty-seven published exclusion envelopes (red) from the ADMX, CAPP+ADMX, RBF/UF, HAYSTAC, TASEH/CAST-CAPP, QUAX, and ORGAN campaigns, the KSVZ and DFSZ model lines, and the one prediction — the Kaxion (red star) at $8.1 \pm 0.8\,\mu$eV on the DFSZ line ($E/N = 8/3$), inside the closed window $\Xi_{\text{cap}} \in [35, 47]$ (orange), every other frequency a mandated silence. The black cross marks the 1.036 GHz candidate excluded in June 2026 — the falsifier F-Kaxion-1d fired and the prediction survived. The projected IBS-CAPP (late 2020s) and ADMX-EFR (∼2030) DFSZ-depth wedges approach the window from below: the window will be confirmed or closed within several years. Full version and sources: Figure 9 in Part VII.

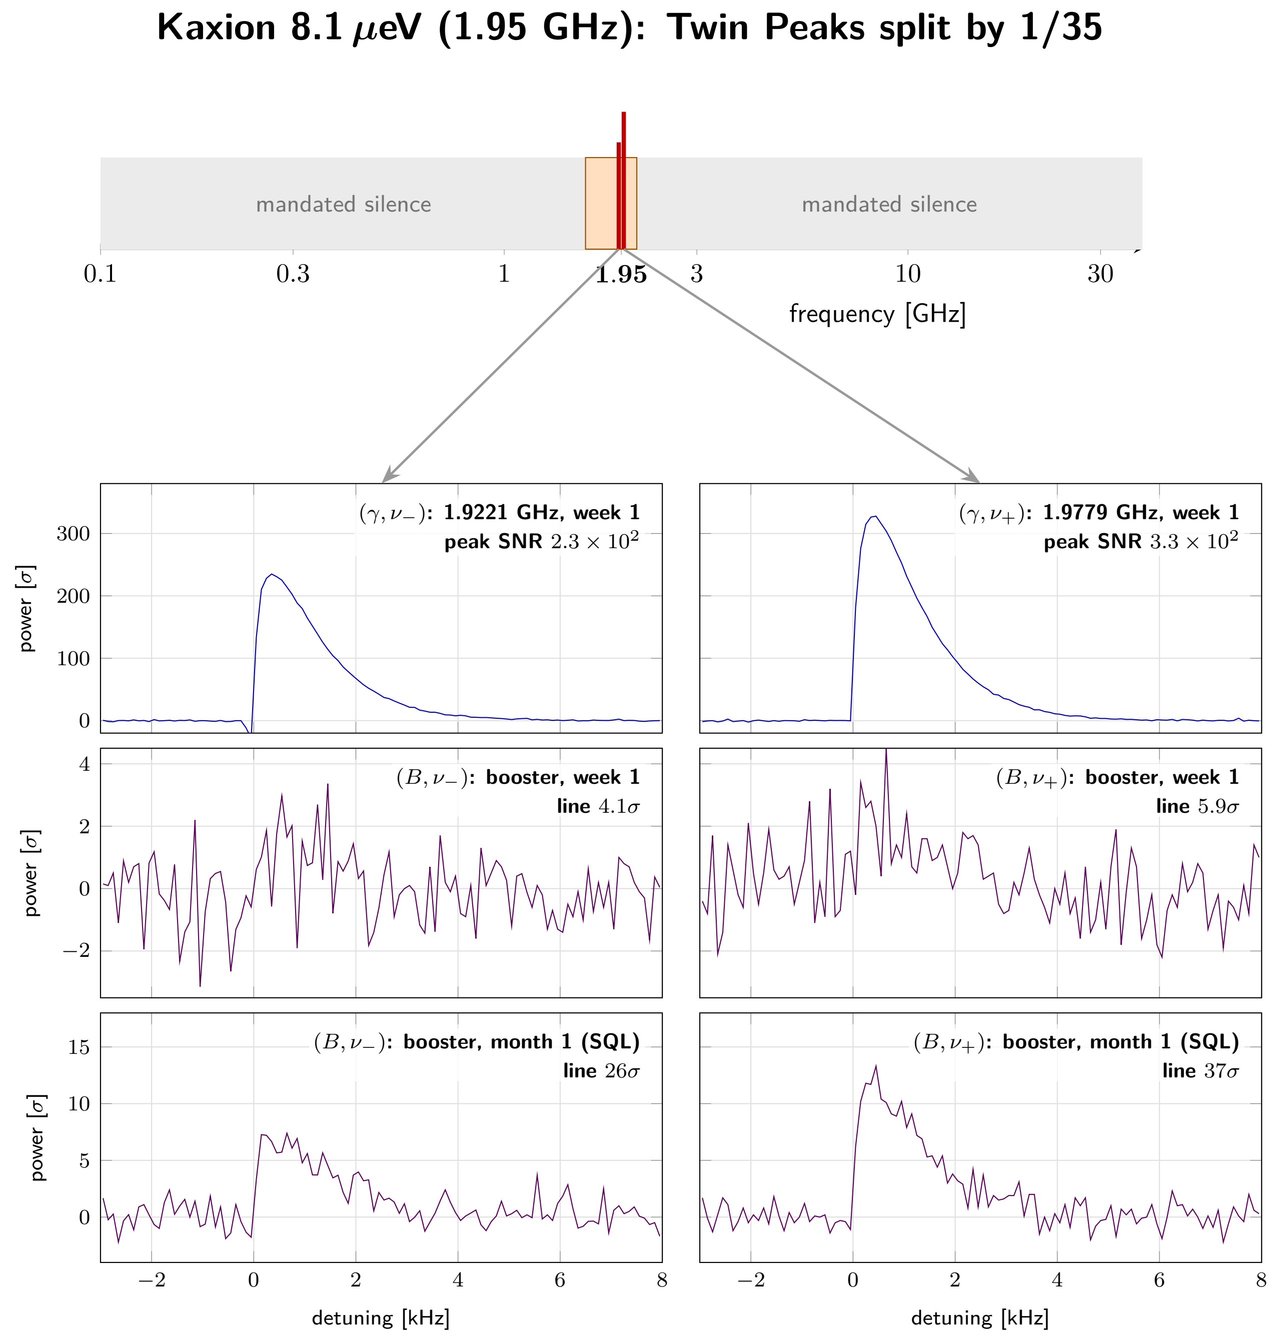


**Figure 2: One line, twice split, three readouts.** The predicted 1.95 GHz line inside its closed window (top; every other frequency a mandated silence), resolved as the 1/35 doublet three ways: at the electromagnetic ports after one week (first row — the reference tone $\nu_+$ camps at a peak-bin SNR $\sim 3.3 \times 10^2$, $\nu_-$ at $\kappa \simeq 0.7$ times that, shapes fully resolved); at the harmonic-booster ports in the same week (second row — faint, $4.1\sigma/5.9\sigma$, yet already a four-fold coincidence at two locked frequencies); and one month after the year-two quantum-limited upgrade (third row — line significance $26\sigma/37\sigma$; the floor is the readout, not the bath). Fixed-seed simulation at KAXIO parameters: Figure 19 in Part X.

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

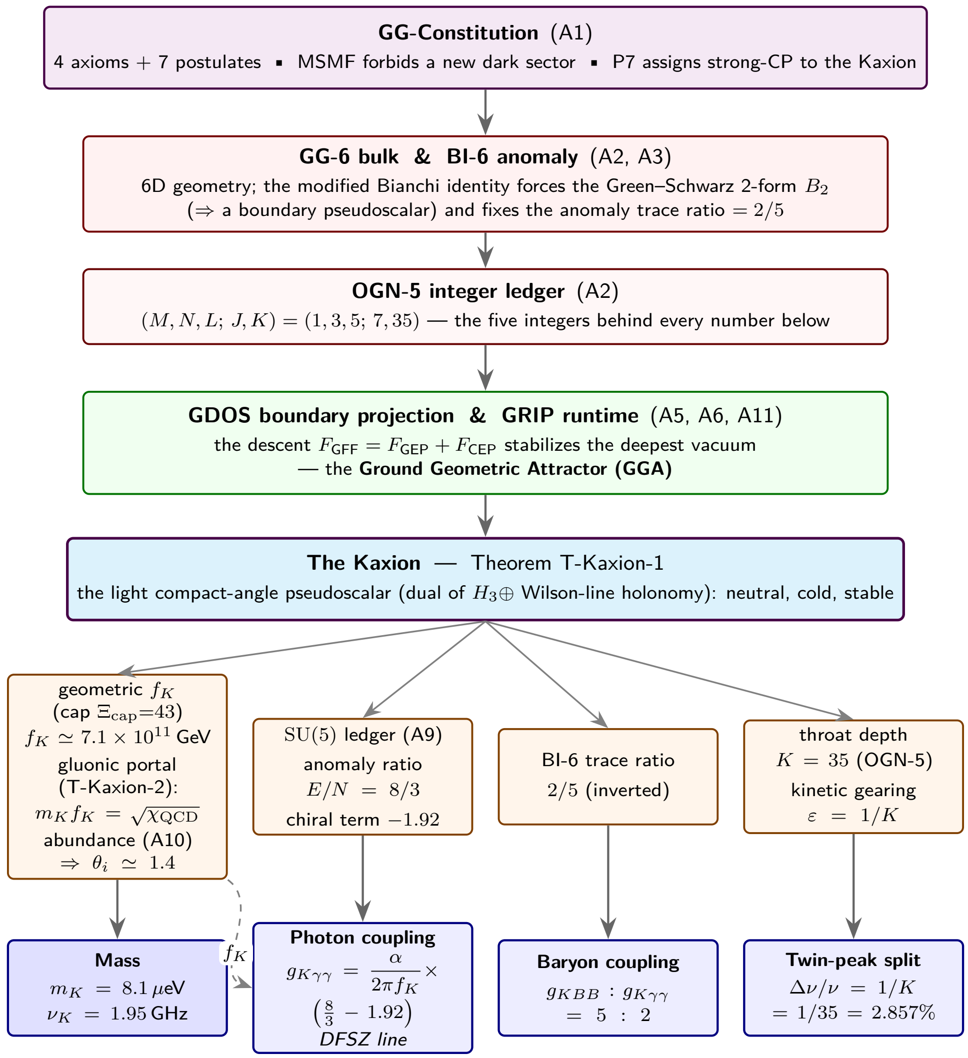


**Figure 3: Derivation flow of the Kaxion: from the constitution to four falsifiable numbers.** Each downward step adds no free parameter: the constitution forces the 6D bulk and the BI-6 anomaly, which carry the five-integer OGN-5 ledger $(1, 3, 5; 7, 35)$; the GDOS boundary law and the GRIP runtime 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 $\pi$-locked abundance and the gluonic portal $m_K f_K = \sqrt{\chi_{\rm QCD}}$; the **photon coupling** from $E/N = 8/3$ on the DFSZ line; the **baryonic companion** from the inverted BI-6 trace ratio, 5:2; and the **twin-peak split** from the throat depth, $\Delta\nu/\nu = 1/K = 1/35$. The dashed link records that the same decay constant $f_K$ fixes both the mass and the photon coupling. Full version: Figure 8 in Part V.

# Ten Numbers,
# and the Measurement That Can Kill Each One

**The paper, stripped of narrative.** Its content is ten 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: *theorem* where the number is inherited or proved, *derived* where a theorem is combined with a Standard-Model input, *benchmark* for the marked row of a discrete window, *closure* where the observed abundance fixes it. Nothing in the last column waits on future theory; all of it waits on data, and two funded experiments — IBS-CAPP in Korea and ADMX-EFR at Fermilab — are built to take it between now and roughly 2030. In full: Table 37, p. 195.

| | **Quantity** | **Value** | **Status** | **Proved** | **Killed by** |
|---|---|---|---|---|---|
| 1 | mass $m_K$ | $8.1 \pm 0.5\,\mu$eV | derived | Part V | F-1c/1d |
| 2 | frequency $\nu_K$ | $1.95 \pm 0.12$ GHz | derived | Part V | F-1c/1d |
| 3 | window | $[1.59, 2.13]$ GHz | discrete ladder | Part V | F-1c |
| 4 | twin split $\Delta\nu/\nu$ | $1/35 = 2.857\%$ | theorem (symm. limit) | Part VI | F-1a |
| 5 | channel ratio $P_B/P_\gamma$ | $25/4 = 6.25$ | theorem | Part VI | F-1b |
| 6 | coupling $g_{K\gamma\gamma}$ | $(1.2 \pm 0.1) \times 10^{-15}$ GeV$^{-1}$ | derived (DFSZ) | Part V | F-1c |
| 7 | decay constant $f_K$ | $(7.1 \pm 0.5) \times 10^{11}$ GeV | bench. $\Xi_{\text{cap}}$=43 | Part V | (via 1–3) |
| 8 | misalignment angle $\theta_i$ | $\simeq 1.4$ rad | closure | Part V | (via 9) |
| 9 | abundance $\Omega_{\text{DM}}$ | $1/\pi - 1/(2\pi^2) = 0.2676$ | derived (GG-A10-Cosmos, cond.) | Part III | CMB |
| 10 | lifetime $\tau_K$ | $\simeq 4 \times 10^{32}\, t_U$ | derived | Part V | (stability) |

# Overview of Contents

# List of Figures

# List of Tables

# Contents

# Preface

This Preface is addressed, more than to anyone else, 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 actually 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 magnificent instruments and airtight null results — and not one of them could ever be pointed at a *prediction*, because no theory ever staked one. Every anomaly the field has debated — an annual modulation here, a 3.5 keV line there, most recently a $5.1\sigma$ candidate at 1.036 GHz — lingered or died ambiguous for the same structural reason: one signature, no pre-registered number, nothing to check it against.

This paper hands you something the field has not had before: a target with its numbers already on paper. A frequency — 1.95 GHz, inside a closed window, not a range to sweep but a bin to camp on. A strength — exactly on the DFSZ line you already calibrate against. A shape, an annual breathing, a sidereal wobble. And two locked ratios no other candidate offers: every line a twin pair split by exactly one part in thirty-five, and a second, baryonic channel answering at exactly twenty-five over four. Parts VI–VIII derive these; Part IX simulates the discovery week they imply, statistical and systematic errors printed on every panel; Part X designs the observatory — KAXIO — that reads all four peaks at once. The simulation code ships with the paper; 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 partner line at 2.857% spacing: it costs no hardware, no beam time, and no permission, and it executes a falsifier of this theory either way (§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 is, to our knowledge, the first dark-matter proposal in fifty years whose one-year outcome on the existence question — discovery or window closure — 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 (Part XIV) and the honesty ledger (Table 13) tell you exactly which links are theorems, which are derived, and which rest on two named conjectures. And the falsifiers cut both ways by design: a missing partner, a wrong ratio, a line outside the window, or a swept and silent window each kills the Kaxion outright — there is no fallback position, and we have written down 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 disc 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 $\sim 10^{-47}\,\mathrm{cm}^2$ 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 $\Omega_b$ and $\Omega_c$ 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 $\pi$-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 *no* adjustable dial: its mass corridor is closed by the $\pi$-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 two parameter-free fingerprints that a generic axion cannot fake: every line is a doublet split by exactly 1/35, and the baryonic-to-electromagnetic channel ratio is locked at 5:2. A signal matching the sheet identifies the Kaxion, the throat depth $K = 35$, and with them the series; a clean measurement violating either lock kills the candidate outright. 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 me 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, fixed by the geometry 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 two parameter-free fingerprints, 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 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 $\pi$-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

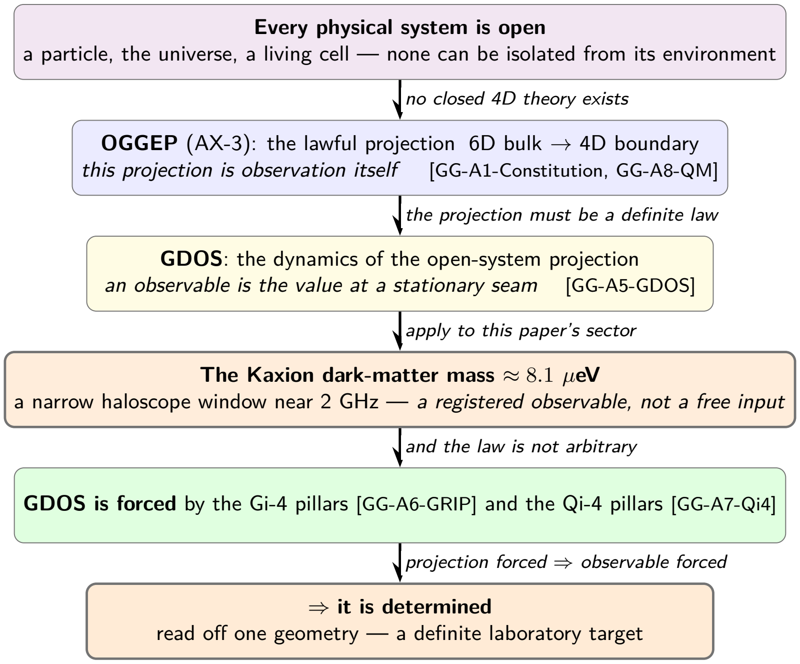


**Figure 4: The central claim, as one deductive funnel**, shared across the GG-Theory cohort. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection is observation itself; its dynamics is GDOS, forced by the 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.

(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 five primers (Appendices B.1–B.5) 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 four 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 5), and a panorama of the fifteen Parts of the present paper (§4.2, Figure 6). A reader of any later Part should be able to flip back to this Section and orient themselves in both directions: which cohort companion carries which imported result, and where the present argument stands in the 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 5 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 *GG6 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 sets the twin-peak 1/35 split. 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 $\pi$ to the phase, the angle of the compact circle that fixes the $\pi$-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 boundary shadow is the Kaxion, and fixes the 5:2 trace ratio. 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 (and forces the single-patch history, so the misalignment angle is an $\mathcal{O}(1)$ displacement later fixed by the observed abundance), and GG-A7-Qi4 (Arisaka, 2026, A7-Qi4), the quantum-information quartet whose current-locking Arisaka equation underlies the 5:2 portal.

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 $\pi$-locked $\Lambda$CDM inventory $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ that the Kaxion must saturate, together with the two named conjectures this paper develops. 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 at $\Delta\nu/\nu = 1/35$). 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 5, which closes this subsection.

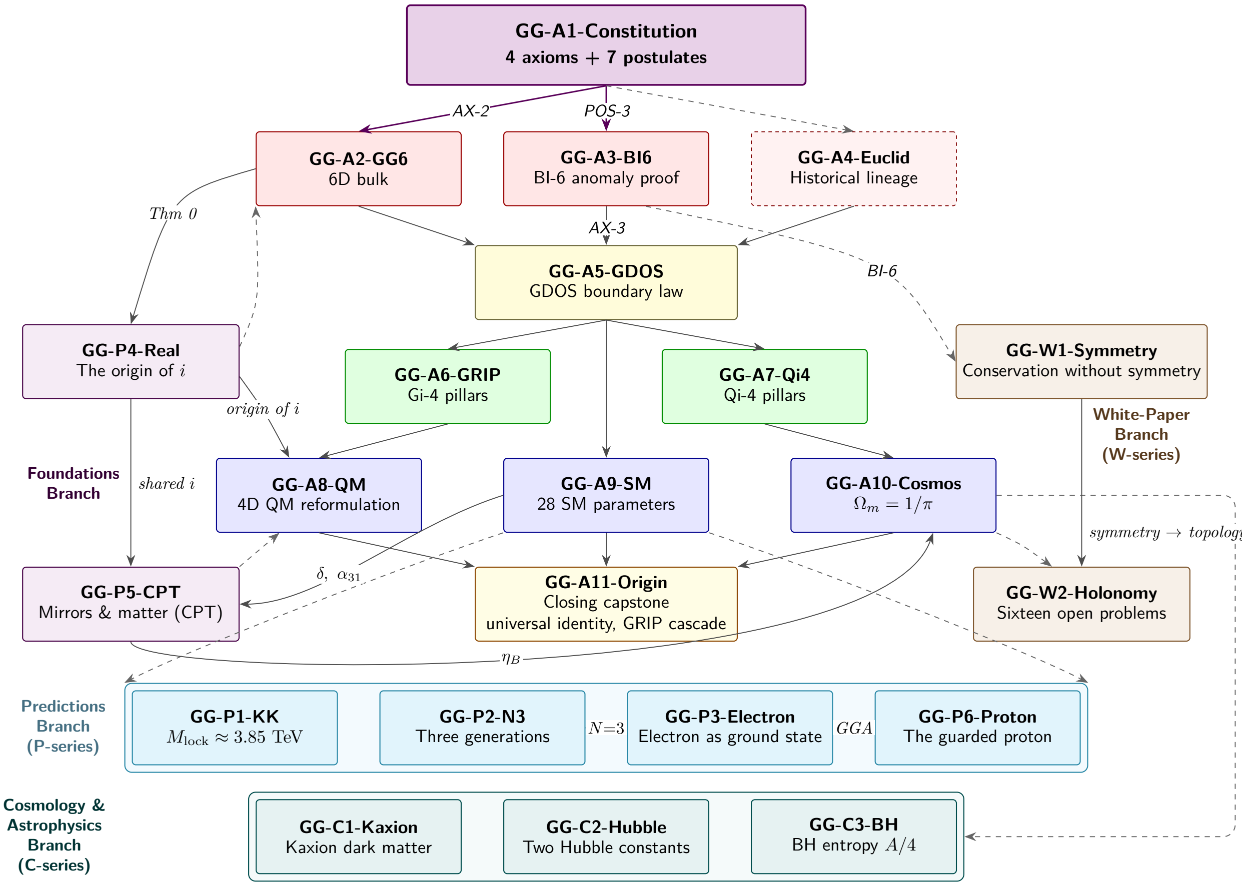


**Figure 5:** The twenty-two-paper cohort map of GG-Theory. The constitutional root GG-A1-Constitution (violet, top) supplies the four axioms and seven postulates from which the rest of the cohort descends: the GG6 Trio (GG-A2-GG6, GG-A3-BI6, GG-A4-Euclid, red) realizes the bulk geometry; GG-A5-GDOS (yellow) the boundary law; the GDOS Trio pillars (GG-A6-GRIP, GG-A7-Qi4, green) the geometric and quantum-information runtime; the 4D-Physics Trio (GG-A8-QM, GG-A9-SM, GG-A10-Cosmos, blue) the textbook readouts; and the closing capstone GG-A11-Origin (gold) the quarks-to-consciousness identity cascade. Two framed branches carry the phenomenology — the **Predictions Branch** (cyan), whose four papers hand their results along the row by the internal arrows (GG-P2-N3's $N{=}3$ 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 $\eta_B$ 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, and the walk through them is plotted in Figure 6. The present Part — the introduction — 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. Everything after this page is the argument itself, 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 that assumes no prior exposure to GG-Theory: the six-dimensional bulk with its two non-spatial coordinates, the boundary law, the five-link existence chain, the two clocks and the light mode, 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 $\pi$-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, and it is 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 $\pi$-lock, the mass $m_K \approx 8.1\ \mu$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 two *Unique fingerprints* no generic axion can imitate — the 1/35 twin peak and the 5:2 portal with its matched-channel 25/4 power lock. 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 instruments for the two parameter-free locks — the twin-peak coincidence resonator and the 5:2 harmonic booster. Part IX, *Simulating the Discovery*, computes the data of discovery week in advance — the camped spectra, the modulation fingerprints, the doublet, and the channel-ratio verdict. Part X assembles everything into *The KAXIO Proposal*: a second-generation Kaxion observatory at Fermilab — two tones, two channels, four peaks, 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 six 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).

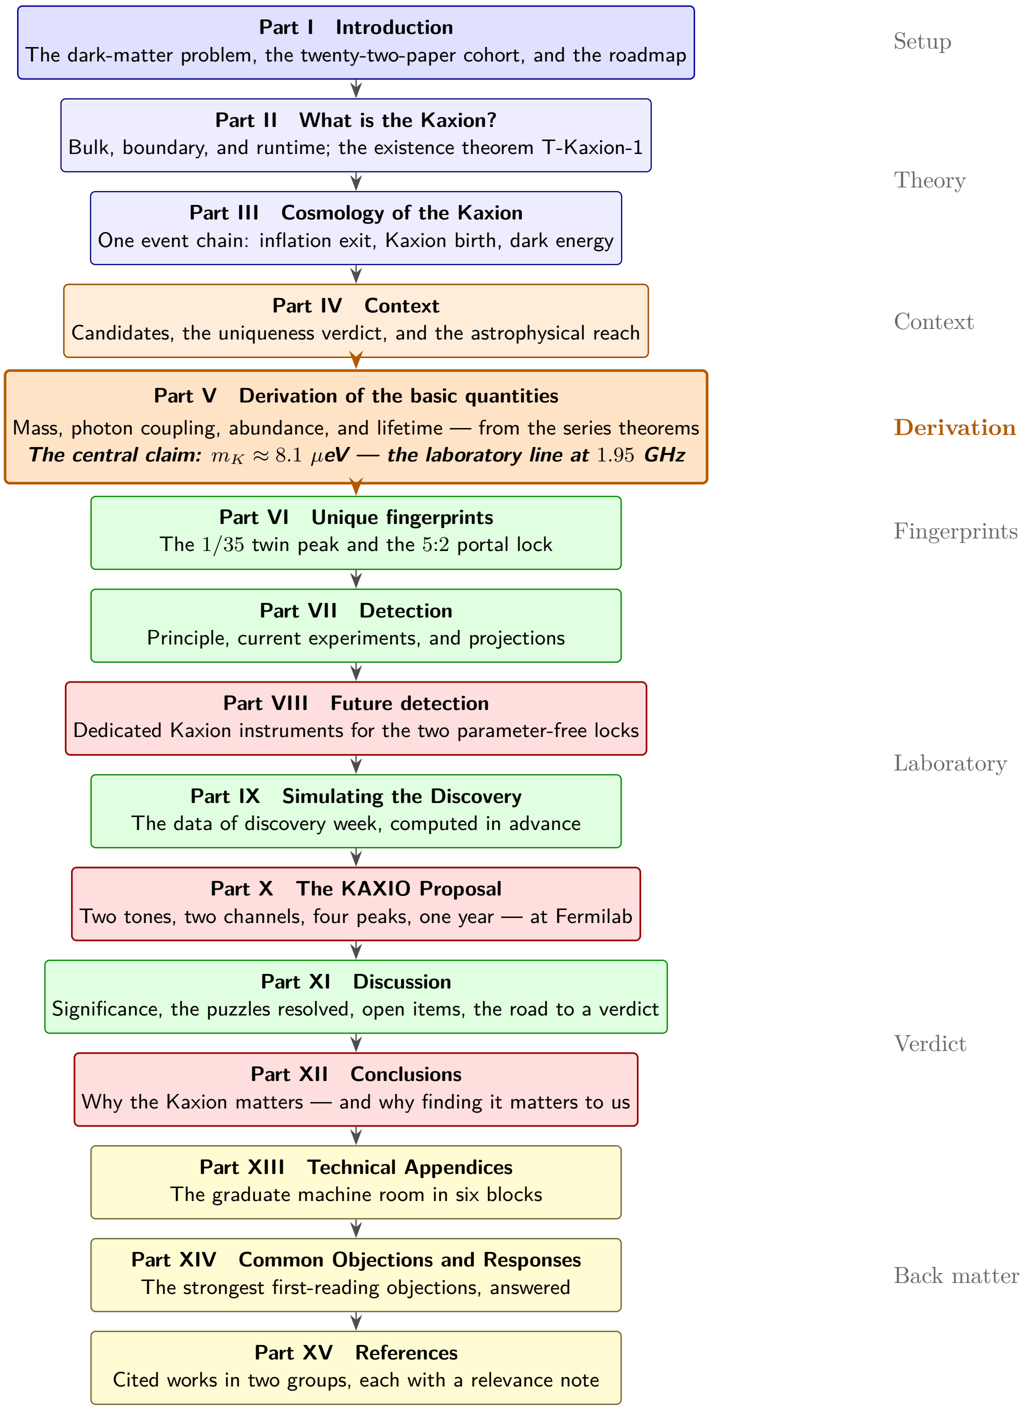


**Figure 6:** Roadmap of the fifteen Parts of GG-C1-Kaxion, grouped into eight phases — *Setup* (Part I), *Theory* (Parts II–III), *Context* (Part IV), *Derivation* (Part V), *Fingerprints* (Part VI), *Laboratory* (Parts VII–X), *Verdict* (Parts XI–XII), and the *back matter* (Parts XIII–XV). **Part V** is highlighted because it carries the paper's central claim: the parameter-free Kaxion mass $m_K \approx 8.1$ $\mu$eV — 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 6 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 — not its mass, not its coupling, not its abundance. 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 (Lorentz Closure) | 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 ancestor of the 5:2 portal. |
| 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. |
| 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

$$\mathcal{M}_6 \;=\; \mathcal{M}_4 \times I_r \times S^1_\theta, \tag{6.1}$$

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

with $G_{\theta\theta} = 1$ and *no free radius*: no adjustable $R_\theta$, no adjustable scale profile, no compactification modulus. In textbook coordinates the six directions are $(t, x, y, z;\, r, \theta) = \{$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: $r = \ln(E_{\mathrm{GUT}}/\mu)$, the logarithm of an energy ratio. Moving inward along $r$ *is* lowering the observation scale $\mu$; 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,

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

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, $E_{\mathrm{GUT}} = E_{\mathrm{Pl}}\,e^{-2\pi} \simeq 2.28\times10^{16}$ GeV. *Why is gravity* $10^{16}$ *times weaker than the weak interaction? Because thirty-five.*

**The phase coordinate.** The circle coordinate $\theta$ is a pure angle — dimensionless, periodic, carrying no unit of length. GG-A2-GG6 is explicit that it is *non-spatial*: a value of $\theta$ is an internal phase, not a location; local motion in $\theta$ 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

$$W \;=\; \mathrm{tr}\; P\exp\Big(i\oint_{S^1_\theta} A_\theta\, \mathrm{d}\theta\Big), \tag{6.4}$$

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, in essence, the slow wandering of its phase.

*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

$$G_{\mathrm{Bulk}} \;=\; \mathrm{Spin}(1,5)\times \mathrm{SU}(5)\times \mathrm{U}(1)_B \times \mathrm{U}(1)_{Qi}. \tag{6.5}$$

Each factor is forced, not chosen. $\mathrm{Spin}(1,5)$ is the Lorentz group of the six-dimensional bulk. SU(5) is the *smallest* simple group containing the Standard Model — the Georgi–Glashow spine — selected by MSMF from among all candidates. $\mathrm{U}(1)_B$ 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 SU(5), whose $X, Y$ bosons mediate $p \to e^+\pi^0$. $\mathrm{U}(1)_{Qi}$ 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 two channels whose locked ratio becomes the Kaxion's 5:2 fingerprint.

### 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,

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

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:

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

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 | $127.930 \pm 0.008$ | +0.05% |
| $\alpha^{-1}(0)$ | $2^J + N^2$ | 137.0 | 137.036 | −0.03% |
| $\sin^2\theta_W(M_Z)$ | $L^2/(4N^3)$ | $25/108 = 0.23148$ | $0.23122 \pm 0.00003$ | +0.11% |
| $\alpha_s(M_Z)$ | $1/(\pi e)$ | 0.1171 | $0.1180 \pm 0.0009$ | −0.8% |
| Higgs vev $v$ | $M_{\rm lock}(2/L)^N$ | 246.33 GeV | 246.22 GeV | +0.05% |
| Recombination $z_*$ | $e^J$ | 1096.6 | $1089.9 \pm 0.3$ | +0.6% |
| Matter share $\Omega_m$ | $1/\pi$ | 0.3183 | $0.3153 \pm 0.0073$ | +1.0% |
| Dark-matter share $\Omega_{\rm DM}$ | $1/\pi - 1/(2\pi^2)$ | 0.2676 | $0.2645 \pm 0.0070$ | +1.2% |
| Dark-energy share $\Omega_\Lambda$ | $1 - 1/\pi$ | 0.6817 | $0.6847 \pm 0.0073$ | −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 — $\hat{\alpha}^{-1} = 2^J$, $\sin^2\theta_W = L^2/4N^3$, $v = M_{\rm lock}(2/L)^N$, $z_* = e^J$ — 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 $\pi/e$ *fractions* — $\alpha_s = 1/(\pi e)$ and the three cosmological shares — are weaker, and we say so plainly: $\pi$ 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 $\alpha_s$, 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 1/35 doublet, the 5:2 lock) are rigid integer relations, not $\pi/e$ fractions.

**A note on the cosmological shares.** Three rows ($\Omega_m$, $\Omega_{\rm DM}$, $\Omega_\Lambda$) 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 $\Omega_{\rm DM}h^2 = 0.1215$, and the constant ratio of the two matter species,

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

which redshifts identically for both species and therefore holds at every epoch. We read the $\pi$-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, twin peak, portal ratio — is one more withdrawal from the same account: the integers $K = 35$, $J = 7$, $L = 5$, and the trace pair $(5, 2)$ 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 scale coordinate identifies $\mu(r) = E_{\rm GUT}e^{-r}$; the two anchor theorems are $E_{\rm GUT} = E_{\rm Pl}e^{-2\pi}$ and the lock theorem of GG-A9-SM, $M_{\rm lock} = \frac{1}{2}E_{\rm Pl}e^{-K}$. Numerically, $e^{-2\pi} = 1.867 \times 10^{-3}$ and $e^{-K} = e^{-35} = 6.305 \times 10^{-16}$, so

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

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,

$$\frac{E_{\rm GUT}}{M_{\rm lock}} = 2\,e^{\,K-2\pi} = 2\,e^{28.72} = 5.9\times 10^{12}, \tag{6.11}$$

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), $v = M_{\rm lock}(2/L)^N = 3846\times(0.4)^3 = 246.1$ 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 9.5); 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,

$$I_8 \;=\; \tfrac{1}{2}\,X_4 \wedge X_4, \tag{6.12}$$

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

$$H_3 \;=\; {\rm d}B_2 \;+\; \omega_3(\Omega) \;-\; \omega_3(A_5) \;-\; \tfrac{2}{5}\,\omega_3(A_B) \;-\; \tfrac{2}{5}\,\omega_3(A_{Qi}), \tag{6.13}$$

obeying the modified Bianchi identity that names the series:

$$\mathrm{d}H_3 \;=\; X_4 \;=\; \mathrm{tr}\,R^2 \;-\; \mathrm{Tr}\,F_5^2 \;-\; \tfrac{2}{5}\,F_B^2 \;-\; \tfrac{2}{5}\,F_{Qi}^2, \tag{6.14}$$

together with the inflow action

$$S_{\mathrm{GS}} \;=\; \int_{\mathcal{M}_6} B_2 \wedge X_4. \tag{6.15}$$

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

$$Y_0 = \mathrm{diag}\big(-\tfrac{1}{3}, -\tfrac{1}{3}, -\tfrac{1}{3}, \tfrac{1}{2}, \tfrac{1}{2}\big), \qquad B_0 = \mathrm{diag}\big(\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{1}{3}, 0, 0\big) \tag{6.16}$$

(the hypercharge and baryon-number generators in the fundamental of SU(5)),

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

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 will become the Kaxion's two-channel fingerprint.

*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 $I_8 = \mathrm{d}I_7$, and let the gauge variation of $I_7$ define $\delta_\lambda I_7 = \mathrm{d}I_6^{(1)}(\lambda)$; 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 $\delta_\lambda S_{\mathrm{GS}} = -\int X_2^{(1)} \wedge X_4 = -\delta_\lambda \Gamma$: 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 $\mathrm{tr}\,F_i^2\;\mathrm{tr}\,F_j^2$ 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 —

$$\partial_\mu J^\mu_a = 0 \;\; (a = \mathrm{em},\, B,\, Qi), \qquad \nabla_\mu T^{\mu\nu} = 0, \tag{6.18}$$

— 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, $J_E/\mathsf{Q_B} = J_e = J_B = J_{Qi}$); the dark-sector reading of that lock is the 5:2 portal ratio, derived in Part V and exploited in Parts VI–VIII.

*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 twin tones — 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,

$$\dot{\rho} = -i[H_{\text{eff}}, \rho] + \sum_k \Big( L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\}\Big); \tag{6.19}$$

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

$$\partial_t P = \nabla\cdot\Big(\frac{P\,\nabla F}{\gamma}\Big) + \frac{T_{\text{eff}}}{\gamma}\nabla^2 P; \tag{6.20}$$

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

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

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,

$$\dot{x} \;=\; -\,g^{-1}\,\nabla F_{\mathsf{GFF}}(x), \tag{6.22}$$

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: $\mathrm{diag}(\chi, M^2)$ at the origin — both eigenvalues positive, Morse index 0: this is the GA. At $(\pi, 0)$: $\mathrm{diag}(-\chi, M^2)$ — one negative eigenvalue, index 1: a GSSB saddle, the watershed of the circle. (The DGI is the $\chi \to 0$ 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,

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

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 $\chi$ — 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

$$F_{\mathsf{GFF}} \;=\; F_{\mathsf{GEP}} \;+\; F_{\mathsf{CEP}}, \tag{6.24}$$

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

$$F_{\mathsf{CEP}} \;=\; -T\,\mathcal{S} \;+\; F_{\mathrm{diff}} \tag{6.25}$$

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 $F_{\mathsf{GFF}}$ organize into a strict descent, which the series names once and uses everywhere:

$$\mathsf{DGI} \;\longrightarrow\; \mathsf{GSSB} \;\longrightarrow\; \mathsf{GA}. \tag{6.26}$$

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):

$$[\Pi_*] \;=\; \underset{[\Pi]\in\mathrm{Bdy}}{\arg\min}\, F_{\mathsf{GFF}}([\Pi]) \quad \text{exists and is unique,} \tag{6.27}$$

and is empirically identified with the Standard-Model vacuum. The same pipeline fixes the Wilson-line breaking SU(5) → 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 $F_{\mathsf{GFF}}$, which GG-A11-Origin names the *Ground Geometric Attractor*,

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

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 GRIP := DGI → (GFF) → GSSB → 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 $\tau \to \mu \to e$ 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 born in a GRIP-Out (the inflationary exit; 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 entries locked at 5:2. 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 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,

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

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 $S^1_\theta$, 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* the twin-peak split $\Delta\nu/\nu = 1/K$ of Part VI; $L = 5$ fixes the unification embedding whose anomaly ratio $E/N = 8/3$ sets the photon coupling of Part V; $J = 7$ closes the dyadic coupling $\alpha^{-1}_{\rm lock} = 2^J = 128$ of Table 2; and the pair $(L - N, L) = (2, 5)$ is the trace ratio (6.17) behind the 5:2 portal. 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 physical degree of freedom (counted in Section 9.1) plus a ledger of locked weights.

### 7.4 Link 4: the boundary provides two pseudoscalars

Projected through GDOS to the boundary, Links 1 and 3 each leave behind exactly one light, periodic, parity-odd scalar direction: the dual of the tensor field strength (the *tensor shadow* $\theta_{\mathrm{GS}}$) and the Wilson-line holonomy promoted to a field (the *holonomy echo* $\theta_W$). Their joint boundary-effective Lagrangian, derived in GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos),

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

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 detailed construction of both directions — the duality transformation, the degree-of-freedom count, the parity and periodicity arguments — 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 and a heavy orthogonal combination lifted at throat scales (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 $\mathcal{M}_4 \times I_r \times S^1_\theta$; 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 integer gear ratios: 1/35, 1/7, $2^{-7}$, 5:2 — every later observable. |
| 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 $\theta_{\mathrm{GS}}$; Wilson line $\theta_W$ | Exactly two periodic pseudoscalar directions, shift-symmetric, parity-odd. |
| 5 | GG-A6-GRIP GFF-ST7 | $[\Pi_*] = \arg\min F_{\mathsf{GFF}}$ | The light eigenstate is selected and stabilized as the Geometric Attractor. |

# 8 The Kaxion existence theorem

**Definition 8.1** (Kaxion)**.** The *Kaxion* is the light mass eigenstate of the two-direction compact-angle pseudoscalar system (7.2) — the canonically normalized combination of the tensor shadow $\theta_{\mathrm{GS}}$ and the holonomy echo $\theta_W$ that survives 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 light eigenstate of precisely two compact-angle directions: the boundary dual of the forced Green–Schwarz field strength $H_3$, and the Wilson-line holonomy of the non-spatial circle $S^1_\theta$.*
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* $(1, -1, -\frac{2}{5}, -\frac{2}{5})$*; in particular the baryonic and electromagnetic portals stand in the ratio* 5:2*.*

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 $S^1_\theta$, 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). These are the only two compact-angle directions available: GG-A2-GG6 T1 forbids further compact factors, and GG-A2-GG6 T3 forbids further gauge factors whose holonomy could supply additional angles.

Protection: $B_2$ enters the bulk action only through $H_3$ and $S_{\rm GS}$, and $A_\theta$ 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 9.4). 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 ratio of the Abelian entries is the trace ratio (6.17), giving 5:2 after the electromagnetic embedding of GG-A9-SM.

Stability: GFF-ST7 of GG-A6-GRIP guarantees a unique global minimizer of $F_{\sf GFF}$ 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 DGI $\to$ GSSB $\to$ 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* state follows because the orthogonal combination acquires its mass from the cap dynamics at throat scales, leaving exactly one eigenstate below the electroweak scale. ■

**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, $\tau_{K\to\gamma\gamma} \gg t_U$ 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 $\pi$-locked abundance closure and the susceptibility relation, and they rest on the two *named conjectures* of GG-A10-Cosmos (misalignment-completeness; compact-angle mass-uniqueness), fenced as such. The theorem-level content is the existence, the protection, the 5:2 lock, and — once the throat arithmetic of Part VI is added — the 1/35 twin-peak split (as a ratio, in the symmetric limit). The reader should hold this division firmly: the experiments of Parts VII–VIII test the geometry through the locks even if the benchmark masses shift.

*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 two clocks and the light mode

Theorem 8.2 leans on the claim that the boundary inherits *exactly two* periodic pseudoscalar directions, and that exactly one light state survives. This section constructs both directions explicitly, at the level of detail a graduate student can audit: degree-of-freedom counts, the duality transformation, parity, periodicity, the mixing matrix, and the mass-protection argument with numbers.

## 9.1 The tensor shadow: 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 $a_{\mathrm{GS}}$ enforcing its Bianchi identity:

$$\mathcal{L}_{\mathrm{1st}} = -\frac{1}{12}\,H_{\mu\nu\rho}H^{\mu\nu\rho} + \frac{1}{6\,f_{\mathrm{GS}}}\,a_{\mathrm{GS}}\,\epsilon^{\sigma\mu\nu\rho}\,\partial_\sigma H_{\mu\nu\rho} \;\longrightarrow\; \begin{cases} \delta a_{\mathrm{GS}}: & \mathrm{d}H_3 = 0 \;\Rightarrow\; H_3 = \mathrm{d}B_2 \;\text{ locally,} \\ \delta H_3: & H^{\mu\nu\rho} = f_{\mathrm{GS}}^{-1}\,\epsilon^{\mu\nu\rho\sigma}\,\partial_\sigma a_{\mathrm{GS}}. \end{cases} \tag{9.1}$$

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

$$H_{\mu\nu\rho} \;=\; f_{\mathrm{GS}}^{-1}\,\epsilon_{\mu\nu\rho\sigma}\,\partial^\sigma a_{\mathrm{GS}}. \tag{9.2}$$

**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 $a_{\mathrm{GS}} \to -a_{\mathrm{GS}}$ under spatial inversion. The dual is a *pseudo*scalar — which is why it can couple to $F\widetilde{F}$ and why its condensate violates no parity test.

**Periodicity.** The flux of $H_3$ through any closed three-cycle is quantized (the differential-cohomology refinement of GG-A3-BI6); through (9.2) this becomes the statement that $a_{\rm GS}$ is an *angle*: $a_{\rm GS} \to a_{\rm GS} + 2\pi f_{\rm GS}$ is a gauge identification. We write $\theta_{\rm GS} = a_{\rm GS}/f_{\rm GS} \in [0, 2\pi)$ and call it 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 $\theta_{\rm GS}$, 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,

$$\vartheta(x) \;=\; \oint_{S^1_\theta} A_\theta(x, \theta)\, \mathrm{d}\theta, \qquad a_\theta(x) \;=\; f_\theta\, \vartheta(x), \tag{9.3}$$

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 $\vartheta \to \vartheta + 2\pi$ 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 $\vartheta$. Its potential, too, arises only from effects that wrap the circle — holonomy-sensitive, topological effects. The same mechanism, run at the unification scale with the heavy quark content, is the Hosotani breaking (Hosotani, 1983) that selects the Standard-Model vacuum inside SU(5); run at the topological-susceptibility scale with the light content, it is the Kaxion potential of Part V. One mechanism, two florets.

**The Higgs is the Kaxion's sibling.** The two florets deserve to be named as one family, because the identification is useful in both directions. The electroweak Higgs and the Kaxion are the *same kind of object*: a continuous modulus of the compact angle, gauge-invariant, admitting no local counterterm, and therefore light for a reason no four-dimensional theory can state. They differ in three respects and in no other. They live on different strata — the Wilson–Hosotani Higgs stratum and the compact-angle stratum. They are massed by different topological effects — gauge and matter loops winding the circle at the unification scale in one case, the topological susceptibility of the light content in the other. And they are caught at different points of the same cascade: the Higgs completed its GRIP descent at the electroweak transition and has sat at its attractor ever since, while the Kaxion is still ringing, the last such descent the observable universe has to run.

Two payoffs follow, and both are recorded at the reading grade rather than claimed as new theorems. The first is expository: a reader who accepts the Higgs — and the community does, at five sigma — has already accepted an object of this type, so the Kaxion is not an exotic addition to physics but the second member of a class with one confirmed member. The second is methodological: whatever technique evaluates the curvature of the holonomy potential at its minimum for one of them should transfer to the other, which is the natural attack on the one step the codebook paper leaves open in its Higgs sector (Arisaka, 2026, A9-SM). The reading also carries an expectation worth stating plainly: light fundamental scalars ought to *be* compact-angle moduli, the bulk has one compact fiber with few strata, and the cohort therefore expects few of them. It has exactly two, one found and one staked on an instrument and a decade.

**Parity.** The echo couples through the same anomaly-descended $\vartheta F\widetilde{F}$ operators as the shadow (they share the ledger), fixing its parity assignment as odd. Both clocks are pseudoscalars.

## 9.3 The kinetic matrix and the light eigenstate

**Why the clocks mix.** The Chern–Simons improvement in (6.13) contains the gauge potential whose holonomy is $\theta_W$; dimensional reduction of the $|H_3|^2$ kinetic term therefore generates, alongside the diagonal kinetic terms, an off-diagonal $\partial\theta_{\rm GS}\,\partial\theta_W$ overlap — the kinetic matrix $K_{ij}$ of (7.2), with entries given by overlap integrals of the two zero-mode profiles over $(r,\theta)$. The two clocks are geared.

**Diagonalization.** Canonical normalization and a rotation bring the system to mass eigenstates:

$$\begin{pmatrix} a_K \\ a_H \end{pmatrix} = \begin{pmatrix} \cos\beta & \sin\beta \\ -\sin\beta & \cos\beta \end{pmatrix} \begin{pmatrix} a_{\rm GS} \\ a_\theta \end{pmatrix}. \tag{9.4}$$

Here $a_H$ is the heavy *radial cap* mode, lifted by the cap dynamics to throat scales and decoupled (Appendix C.3); $a_K$ is the canonically normalized light eigenstate — the Kaxion of Definition 8.1 — which keeps the flat potential and the locked portals. The decoupling of the cap mode is what fixes the light sector's canonical decay constant $f_K$ (the cap renormalization $\Xi_{\rm cap}$ of §26.3); nothing in the present Part depends on its value. One fact must be stated at the Lagrangian level now, because Part VI leans on it: the light compact-angle sector is *not* one field but *two* shift-protected zero-modes — the Green–Schwarz dual $a_{\rm GS}$ and the Wilson-line phase $a_\theta$ — and both stay light, because the shift symmetry that protects one protects the other. They are the two members of the twin-peak doublet (Part VI); $a_K$ is their QCD-massive light eigenstate, and the doublet is its fine structure, resolved just below. The full spectrum is therefore *three* states: two light zero-modes and one heavy cap mode.

**The two-by-two algebra, explicitly.** For the student who wants the diagonalization rather than its announcement. The reduced Lagrangian of the two angle fields is

$$\mathcal{L} = \tfrac{1}{2}\,\dot{\mathbf{a}}^{\mathsf{T}} K\,\dot{\mathbf{a}} - \tfrac{1}{2}\,\mathbf{a}^{\mathsf{T}}\mathcal{M}^2\,\mathbf{a}, \qquad K = \begin{pmatrix} 1 & \epsilon \\ \epsilon & 1 \end{pmatrix}, \qquad \mathcal{M}^2 = \begin{pmatrix} \mu^2 & 0 \\ 0 & \mu^2 + \Delta^2 \end{pmatrix}, \tag{9.5}$$

with $\epsilon = 1/K = 1/35$ the dimensionless kinetic gearing (computed geometrically in Appendix C.2) and $\Delta$ the cap contribution. The matrix (9.5) pairs *one* light direction with the heavy cap mode; with the full three-state content in view, the diagonalization acts on two distinct pairs, and keeping them apart is what reconciles this Part with Part VI.

*Light against heavy — the cap lift.* Because $\Delta \gg \mu$ (the QCD-induced light scale), diagonalizing (9.5) leaves the cap mode at $\simeq \Delta$ — decoupled, the heavy state of Section 9.5 and Appendix C.3 — and renormalizes the surviving light kinetic norm. This canonical normalization *is* the cap renormalization $\Xi_{\text{cap}}$ that fixes $f_K$ (§26.3): the lift removes the heavy mode and sets *where* the line sits, but it does not split the light sector.

*Light against light — the doublet.* The split lives in the two surviving light zero-modes $(a_{\text{GS}}, a_\theta)$. They carry a *diagonal, equal* QCD-induced curvature $\mu^2 = \chi_{\text{QCD}}/f_K^2$ — each inheriting it through its own geometric route ($a_{\text{GS}}$ via the BI-6 inflow, $a_\theta$ via the Wilson-line holonomy of the color subgroup of SU(5)) — with the diagonal, equal-curvature structure protected by the throat's reflection symmetry that exchanges the two directions. Their only off-diagonal coupling is then the kinetic gearing $\epsilon$, so the normal modes are $u_\pm = (a_{\text{GS}} \pm a_\theta)/\sqrt{2}$ with kinetic coefficients $1 \pm \epsilon$, and canonical rescaling gives

$$\omega_\pm^2 = \frac{\mu^2}{1 \pm \epsilon} \qquad \Longrightarrow \qquad \frac{\Delta\nu}{\nu} = \frac{\omega_- - \omega_+}{\omega} \simeq \epsilon = \frac{1}{K}, \tag{9.6}$$

the twin-peak formula, here previewed and in Part VI made the centerpiece.

*The degeneracy, and its one assumption.* The split $1/K$ is the *symmetric-limit* value: it is exact insofar as the two zero-modes' curvatures are equal, which the exchange symmetry enforces. This is the crucial physical point, and it is worth seeing why the doublet is genuine rather than the generic outcome. With a single anomaly and a *shared* coupling $\propto (a_{\text{GS}} + a_\theta)$ one would get one massive state plus a massless partner — no doublet. Here the two routes deliver a *diagonal* potential (each zero-mode coupled to the topological charge through its own geometric channel, the off-diagonal potential term forbidden by the exchange symmetry), so both acquire $\mu^2$ and the only splitting is kinetic. A residual curvature asymmetry $\delta \equiv |\omega_1^2 - \omega_2^2|/\mu^2$ — were the symmetry broken — would add in quadrature, $\Delta\nu/\nu = \sqrt{\epsilon^2 + \delta^2}$, so $1/K$ is the value when the routes are symmetric and a lower bound otherwise. We therefore carry the doublet's *existence* — two light compact-angle zero-modes, not one massive state and an invisible massless partner — as structural, and its *exact ratio* $1/K$ as the symmetric-limit prediction resting on the equal-curvature assumption: the one model-level input behind the first fingerprint's precise value, flagged as such in the ledger of Part V.

*Physical picture.* Two clock hands geared on one dial: the bulk Green–Schwarz normalization ($a_{\text{GS}}$) and the canonical light eigenstate ($a_K$) the apparatus couples to. The cap renormalization relates them by the integer $\Xi_{\text{cap}}$; the physical decay constant, mass, and frequency are those of $a_K$, the single propagating field — the bulk normalization is an intermediate of the diagonalization, not a second particle.

## 9.4 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 $e^{-S_{\rm inst}}$ of their Euclidean action. Such effects generate a periodic potential of the generic form $V = \chi\,[1-\cos\theta]$, where $\chi$ — 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:

$$m_K \;=\; \frac{\sqrt{\chi}}{f_K}. \tag{9.7}$$

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

$$m_K = \frac{(75.5\,\mathrm{MeV})^2}{7.1\times 10^{11}\,\mathrm{GeV}} = \frac{5.70\times 10^{15}\,\mathrm{eV}^2}{7.1\times 10^{20}\,\mathrm{eV}} \simeq 8.0\times 10^{-6}\,\mathrm{eV}, \tag{9.8}$$

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

## 9.5 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 $S^1_\theta$. Because the circle is rigid with unit metric, these quanta sit at the ultraviolet end of the scale interval, with masses of order $|n|\,E_{\rm GUT} \sim |n|\times 2.3\times 10^{16}$ GeV — thirty-one 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 $M_{\rm lock} \approx 3.85$ TeV (6.3) — nineteen orders above the Kaxion, and the collider business of P1-KK. Third, the *compact-angle zero-modes*: the $n = 0$ angle variables themselves — the two clocks — 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 (9.7). 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 ($e^{in\theta}$, $n \neq 0$) | $\sim \lvert n\rvert\, E_{\mathrm{GUT}} \sim 10^{16}$ GeV | — (inaccessible) | Quantized wiggles around the phase circle; frozen out at unification. |
| Radial lock cluster | $M_{\mathrm{lock}} \approx 3.85$ TeV | P1-KK | Resonances of the scale interval at $r = K$; the collider signature. |
| Compact-angle zero-modes | $\mu$eV corridor | **this paper** | The dial positions themselves; protected by shift symmetry, massed only by topology, Eq. (9.7). |

### 9.6 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 1/35 and 5:2 fingerprints that no Peccei–Quinn construction possesses (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 $\Sigma_\theta$ is the two-torus spanned by the clocks, $(\theta_{\mathrm{GS}}, \theta_W) \in T^2$, with the free energy $F_{\mathsf{GFF}}|_{\Sigma_\theta}$ 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 $2\pi M/L = 2\pi/5$ — a discreteness inherited from the integer quantization of the Wilson loop (6.4). 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

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

the explicit Hessian computation, including the off-diagonal clock mixing, being Appendix C.3. 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 $m_K/f_K = \chi^{1/2}/f_K^2 \sim 10^{-59}$ 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 $5 \times 10^{-38}$: 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 $F_{\mathsf{GFF}}|_{\Sigma_\theta}$,

$$\mathsf{GGA}_\theta \;=\; \underset{(\theta_{\mathrm{GS}},\theta_W)\in\Sigma_\theta}{\arg\min}\; F_{\mathsf{GFF}}, \tag{10.2}$$

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 $\Sigma_\theta$ 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 perfect coherence (a population of excitations of a single type, Part III), and the universality of the gear ratios — the 1/35 and the 5:2 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 $\mathrm{U}(1)_B$ 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 → GSSB → 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 $S^1_\theta$ 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 circle — the compact fiber's own angular coordinate, promoted to a boundary field by OGGEP, 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 holonomy coordinate itself, 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 7 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): both parents are creatures of the closed fiber — the holonomy echo $\theta_W$ *is* holonomy, the Wilson line of $S^1_\theta$ promoted to a field, and the tensor shadow $\theta_{\rm GS}$ is the circle's *bookkeeping made field*, the boundary dual of the tensor that exists because the anomaly ledger must balance (GG-A3-BI6). *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 $(1, -1, -\frac{2}{5}, -\frac{2}{5})$, the 5:2 ratio, 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 *twin-peak split* $\Delta\nu/\nu = 1/K$ of Part VI is a ratio of the circle's integers, scale-free and immune to every residual uncertainty in this paper. 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 analogue 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 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 gearbox: the one place the fibers multiply.** Exactly one structure in this paper belongs to both fibers at once, and it is exactly where the one open number lives. The kinetic matrix of the two clocks (§9.3) carries the Green–Schwarz direction with an infrared-enhanced norm $\sim \pi e^{2K}$ — the *interval's depth appearing inside the circle sector's kinetic matrix*, $e^{70} \simeq 2.5 \times 10^{30}$ — against the Wilson direction's order-one norm, and the cap that closes the throat renormalizes the canonical decay constant by the integer loading of (26.3): $f_K = f_{\text{bulk}}/\Xi_{\text{cap}}$. Note the *shape* of that formula: an interval-descended 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 theorem-shaped; the *which* integer is the compact-angle mass-uniqueness conjecture (§26.3); and the seam between the two is the only place in the entire construction where a number can still move. The one open item of this paper lives at the one point where the two fibers mesh — which is exactly where a two-fiber theory should keep its last unknown.

**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.

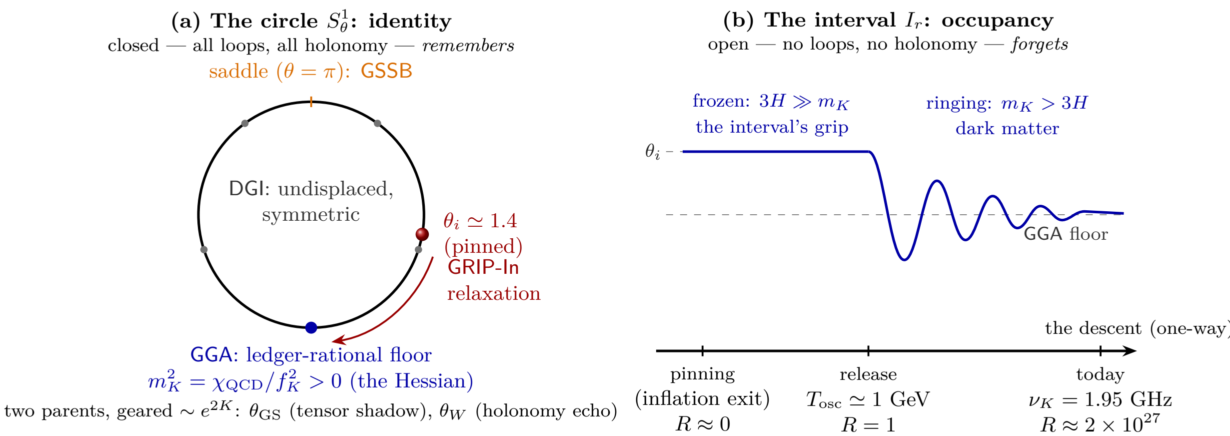


**Figure 7: The Kaxion on the two fibers.** *(a) The circle: identity.* The compact fiber $S^1_\theta$ carries the two parent angles ($\theta_{\rm GS}$, $\theta_W$, geared by the throat norm $\sim e^{2K}$); its vacua sit at discrete ledger-rational phases (dots); the GRIP cascade runs DGI (center) $\to$ GSSB (the $\theta = \pi$ saddle) $\to$ GGA (the floor), and the curvature at the floor *is* the dark-matter mass, Eq. (10.1). Everything on this panel is exact and unrunnable: portals 5:2, split $\Delta\nu/\nu = 1/K$, one frozen angle. *(b) The interval: occupancy.* The one-way descent pins the angle at inflation exit ($R \equiv m_K/3H \approx 0$), 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 $T_{\rm osc} \simeq 1$ 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 its inflationary-exit saddle, leaving the compact-angle sector displaced from its eventual attractor. The field is born displaced *by the geometry of the cascade* — a saddle crossing does not deposit a system at the bottom of the next basin — rather than by accident.

*Selection.* The cascade leaves the compact-angle sector displaced from its attractor by an amount set by the inflationary-exit kinematics — a structured $\mathcal{O}(1)$ angle, not a tuned one. The runtime does not predict its precise value; Part III fixes it from the observed abundance ($\theta_i \simeq 1.4$ 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 inflationary exit, 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 $S^1_\theta$ with quantized holonomy, and by Theorem 8.2 its two light periodic descendants on the boundary are then not optional. 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 inflation-exit 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 reflection-symmetric dressing of GG-A10-Cosmos pinning one angle coherently across the whole future observable universe. This is the one conditional step, and it is named, owned, and falsifiable: a future detection of axion-sector isocurvature kills the pinning outright (§13.5). Nothing about the step is a new hypothesis of this paper — it is a segment of the event chain already required upstream; "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 pinning (killable by isocurvature), and the quantitative address rides on the cap conjecture (§26.3; the discrete window of Part V) — 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 *on* it (labeled) | matter *guarded by* it (welded) | matter *of* it (the coordinate itself) |
| Stratum | charge stratum | baryon stratum | the compact-angle stratum $\Sigma_\theta$ |
| 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 $(1, -1, -\frac{2}{5}, -\frac{2}{5})$; split $1/K$; ledger-rational vacuum |
| GSSB event | family cascade | confinement weld | inflation-exit displacement |
| 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 locked 5:2 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 (Theorem 8.2) — the light eigenstate of two geared clocks (Section 9), exponentially below the geometric scales for a structural reason (Section 9.5), 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:

$$\mathsf{BI\text{-}6} \implies \mathsf{TGG} \implies \mathsf{CCI} \implies \mathsf{GFF} \implies \mathsf{GRIP} \implies \text{parameter block.} \tag{12.1}$$

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 inflationary-exit saddle is crossed, the compact-angle sector is left 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,

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

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

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

with $g_*(T)$ the effective number of relativistic degrees of freedom ($g_* \simeq 73$ at $T \sim 0.6$ 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

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

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 selected at the inflationary exit survives unmodified across thirty-some orders of magnitude in temperature — 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

$$3H(T_{\rm osc}) \;\simeq\; m_K(T_{\rm osc}), \tag{12.5}$$

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 combination 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 the two clocks: 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 $T_{\rm osc} \simeq 1$ 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 inflationary exit 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 not chosen; it is what the observed abundance requires.

| **Epoch** | **Clock** | **Event** |
|---|---|---|
| Inflationary exit | — | GRIP crosses the exit saddle; the compact-angle sector is born displaced at an $\mathcal{O}(1)$ angle, fixed by the abundance to $\theta_i \simeq 1.4$ rad (Section 13.3). |
| $T \simeq 1.0$ GeV | $t \simeq 0.3\,\mu$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 | $t \simeq 13\,\mu$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 $\sim 10^{19}$ oscillations old, begins to dominate structure growth. |
| Today | $t = 13.8$ Gyr | The dial still rings: $\sim 2.0 \times 10^{25}$ 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 $t \simeq 0.3\,\mu$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, $m_K = 8.1\,\mu\text{eV} = 8.1 \times 10^{-15}$ GeV (Part V). Then (12.5) with (12.3) gives

$$T_{\text{osc}}^{\text{const}} = \left( \frac{m_K\, M_{\text{Pl}}}{3 \times 1.66\,\sqrt{g_*}} \right)^{1/2} = \left( \frac{(8.1 \times 10^{-15})(1.22 \times 10^{19})}{42.3} \right)^{1/2} \text{GeV} \simeq 48\,\text{GeV}, \tag{13.1}$$

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

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

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,

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

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

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

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

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

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 $t \to 0$ is

$$\boxed{a_K(t) = a_i\,\Gamma(\tfrac{5}{4})\Big(\frac{2}{m_K t}\Big)^{1/4} J_{1/4}(m_K t)} \tag{13.6}$$

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

$$a_K(t) \;\longrightarrow\; \begin{cases} a_i\Big[1 - \dfrac{(m_K t)^2}{10} + \cdots\Big], & m_K t \ll 1 \quad \text{(frozen, with the leading slip)}, \\ 0.86\, a_i\,(m_K t)^{-3/4}\cos(m_K t - \frac{3\pi}{8}), & m_K t \gg 1 \quad \text{(oscillating)}, \end{cases} \tag{13.7}$$

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 analogue 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; our $\pm 8\%$ pipeline systematic (Part V) is the residual of such coefficients after matching to the precision literature.

*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 thirty orders of magnitude of cosmic history leave 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 $m_K = 8.1\,\mu$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, $\theta_i$. In generic axion cosmology $\theta_i$ 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 decay constant is fixed by the geometry (Section 14), it is no longer $f_K$ that is free and $\theta_i$ that is scanned — it is the reverse: $f_K$ is predicted, and the observed abundance *measures* $\theta_i$.

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 field's frozen value and the nearest vacuum, set by the kinematics of the inflationary-exit saddle. 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,

$$\theta_i \;\simeq\; 1.4 \text{ rad}, \tag{13.8}$$

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 $\theta_i$ is homogenized across the observable universe; the post-inflationary alternative would replace $\theta_i^2$ by the circle average $\langle\theta^2\rangle = \pi^2/3$ and add minicluster substructure, a contingency fenced in Part IV. Second, at $\theta_i \simeq 1.4$ rad the cosine potential is moderately anharmonic, so the harmonic estimate (15.7) carries an $\mathcal{O}(1)$ correction; this shifts the inferred angle within the 1.1–1.5 rad range and is folded into the $\pm 8\%$ abundance systematic of Part V — it does not touch the frequency, which the geometry fixes.

## 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 inflationary exit (taking, generously, even a value as high as $H \sim 10^{13}$ GeV — far above the lock-epoch scale that actually sets the angle) to onset, the accumulated classical drift is dominated by the last Hubble time before onset and is of order $(m_K/3H)^2 \to \mathcal{O}(1)$ only *at* onset — before that, the integrated slip is utterly negligible (at $T = 10$ GeV, $(m_K(T)/3H)^2 \sim 10^{-30}$). Quantum (de Sitter) fluctuations during inflation shift the angle by $\delta\theta \sim H_{\rm inf}/(2\pi f_K)$ per e-fold; for the lock-epoch scale of GG-A10-Cosmos that sets the angle this is $\delta\theta \simeq 7.7 \times 10^{-25}$ — the runtime's selected angle is delivered to the onset epoch effectively untouched. (The same estimate is what keeps the isocurvature contamination far below current CMB bounds; we record the qualitative statement here and defer the transfer-function computation, as GG-A10-Cosmos does, to future work.)

## 13.5 The isocurvature question, quantified

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

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

Planck bounds $\beta_{\rm iso} < 0.04$ (Planck Collaboration, 2020). The scale that governs the Kaxion angle is the lock epoch used throughout this paper (and in the scenario-exclusion argument of Part VII), *not* the primary GUT-adjacent inflation ($H_1 \sim 1.6 \times 10^{13}$ GeV, which does not set the compact angle): $V^{1/4} \sim M_{\rm lock}$ gives the lock-epoch curvature $H_2 = M_{\rm lock}^2/(\sqrt{3}\,\bar{M}_{\rm Pl}) \simeq 3.5 \times 10^{-12}$ GeV. The per-mode fluctuation is then $\delta\theta = H_2/(2\pi f_K) \simeq 7.7 \times 10^{-25}$, and (13.9) evaluates to

$$\beta_{\rm iso} \;=\; \Big(\frac{2 \times 7.7 \times 10^{-25}}{1.4}\Big)^2 \Big/\, 2.1 \times 10^{-9} \;\simeq\; 6 \times 10^{-40}:$$

the isocurvature channel is dead silent, thirty-eight orders below the bound. Indeed the constraint can be inverted into a consistency statement: the Kaxion with $\theta_i \simeq 1.4$ and $f_K = 7 \times 10^{11}$ GeV tolerates an angle-setting curvature scale up to $H^{\max} = \pi\theta_i f_K \sqrt{0.04\,\mathcal{P}_\zeta} \simeq 2.8 \times 10^7$ GeV before (13.9) bites. The primary GUT-adjacent inflation ($H_1 \sim 1.6 \times 10^{13}$ GeV) sits *above* this tolerance — the familiar high-scale-inflation axion isocurvature problem — but it does not set the Kaxion angle; the lock epoch does ($H_2 \simeq 3.5 \times 10^{-12}$ GeV), passing by nearly nineteen orders of magnitude. The standard danger is thus converted into a feature: the compact angle is fixed at the lock crossing, geometrically single-patch, and never scanned during primary inflation. (In the post-inflationary alternative the question is moot: there is no super-horizon angle to perturb, and the relevant inhomogeneities are the order-unity patch fluctuations that become the miniclusters of Part IV.)

## 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 $\theta_i$. The logic here is the reverse of generic axion cosmology — *the geometry fixes the decay constant, and the observed abundance then fixes the angle* — and it is worth stating once, plainly.

**One decay constant, fixed by the geometry.** The light eigenstate of the two geared clocks (Section 9.3) is a single canonically normalized field with one decay constant $f_K$, set by the scale-throat kinetic norms of Appendix C.1: the Green–Schwarz direction carries an infrared-enhanced norm $\sim \pi e^{2K}$ and the Wilson direction an order-one norm, and the cap that closes the throat renormalizes the canonical decay constant down from the bulk normalization scale by an integer cap loading (Section 26.3). The benchmark loading gives

$$f_K = (7.1 \pm 0.5) \times 10^{11}\,\text{GeV}, \tag{14.1}$$

with a residual theoretical uncertainty — the cap computation is not yet complete — spanning a discrete window that Part V converts into $\nu_K \in [1.59, 2.13]$ GHz (benchmark 1.95 GHz). The susceptibility relation

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

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

**The angle follows from the abundance.** With $f_K$ geometric, the displacement at birth $\theta_i$ is the one remaining input, and it is no longer free: requiring the misalignment condensate to saturate the derived inventory $\Omega_{\rm DM}$ (Section 15.3) fixes it at the generic value $\theta_i \simeq 1.4$ rad. This is the standard $\mathcal{O}(1)$ misalignment angle, fitted to the abundance exactly as in every axion-dark-matter calculation — the only continuous freedom in the cosmology, removed by the data rather than by hand.

**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 $\rho_{\rm loc} \simeq 0.45\,{\rm GeV/cm^3}$ through $\rho = \frac{1}{2}m_K^2 a_K^2$ and the invariant (14.2):

$$\theta_0 = \frac{\sqrt{2\rho_{\rm loc}}}{m_K f_K} = \frac{2.6\times10^{-3}\,{\rm eV}^2}{(75.5\,{\rm MeV})^2} \simeq 4.6\times10^{-19}, \tag{14.3}$$

a half-attoradian swing oscillating at $\nu_K = 1.95$ GHz — the $10^{-19}$-radian bell detection must hear (Part VII). *(ii) The twin-peak ratio* $\Delta\nu/\nu = 1/K = 1/35$ of Part VI is a ratio of frequencies and so carries no scale dependence whatsoever: $\Delta\nu \simeq 56$ MHz about 1.95 GHz, fixed by the integer $K$ irrespective of the cap calculation's residual uncertainty — which is why Part VI makes it the primary identification test.

# 15 Abundance and the $\pi$-locked inventory

## 15.1 The target is derived, not fitted

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

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

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

$$\boxed{\Omega_{\rm DM} \;=\; \frac{1}{\pi} - \frac{1}{2\pi^2} \;=\; 0.26765} \tag{15.2}$$

— the *target* that any dark-matter candidate of the program must saturate. Two cautions of principle travel with it. The fraction $\Omega_{\rm DM}$ is a present-epoch quantity; the timeless content the closure below actually uses is the comoving $\Omega_{\rm DM}h^2 = 0.1215$ and the constant species ratio $\rho_{\rm DM}/\rho_b = 2\pi - 1$ (Eq. 6.9). And the $\pi$-locked inventory, like the $\alpha_s = 1/(\pi e)$ entry of Table 2, is among the *weaker* ($\pi/e$-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 $\Omega_c$. Against observation the chart performs as follows:

**Table 7:** The $\pi$-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 | $\pi$-lock | Value | Planck 2018 | Deviation |
|---|---|---|---|---|
| $\Omega_m$ | $1/\pi$ | 0.3183 | $0.3153 \pm 0.0073$ | $+1.0\%$ $(+0.4\sigma)$ |
| $\Omega_\Lambda$ | $1 - 1/\pi$ | 0.6817 | $0.6847 \pm 0.0073$ | $-0.4\%$ |
| $\Omega_b$ | $1/(2\pi^2)$ | 0.05066 | $0.0493 \pm 0.0003$ | $+2.8\%$ |
| $\Omega_{\rm DM}$ | $1/\pi - 1/(2\pi^2)$ | 0.26765 | $0.2645 \pm 0.0070$ | $+1.2\%$ |
| $\Omega_{\rm DM}h^2$ | — | 0.1215 | $0.1200 \pm 0.0012$ | $+1.3\%$ |

### 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:

$$\rho_{\rm osc} = \chi_{\rm QCD}(T_{\rm osc})\big[1 - \cos\theta_i\big] \simeq \tfrac{1}{2}\, m_K^2(T_{\rm osc})\, f_K^2\, \theta_i^2, \tag{15.3}$$

the small-angle form carrying an $\mathcal{O}(1)$ anharmonic correction at (13.8), folded into the $\pm 8\%$ pipeline systematic (Section 15.3).

**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

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

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, $E/\omega$ 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 $\dot m/m^2 = bH/m_K(T_{\rm osc}) = b/3 \simeq 1.4$, 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 literature's transfer functions resolve numerically, and is contained in our $\pm 8\%$ pipeline systematic.

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

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

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

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

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

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

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

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 — every constant evaluated explicitly in Part V, where the pipeline is also cross-validated against the precision misalignment literature — gives the master formula

$$\Omega_K h^2 \;\simeq\; 0.090\, \theta_i^2 \left(\frac{f_K}{10^{12}\,\mathrm{GeV}}\right)^{1.165}. \tag{15.7}$$

(This is the one misalignment normalization the series carries: GG-A10-Cosmos states the identical prefactor at its misalignment equation, and the cross-check that guards the shared pipeline is Section 25.5.)

*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* ($T_{\rm osc}$ 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 $T_{\rm osc}^{-1} \propto f_K^{1/(b+2)} \cdot (\ldots)$, landing at $(b{+}3)/(b{+}2)$. A reader who can reproduce this paragraph owns the misalignment mechanism.

### 15.3 The closure: from the $\pi$-chart to the misalignment angle

Now run the logic in the direction unique to this program. Equation (15.7) has *one* unknown left: with the decay constant fixed by the geometry at $f_K = (7.1 \pm 0.5) \times 10^{11}$ GeV (14.1), the exponent fixed by lattice QCD, and the target fixed by the derived chart (15.2) — converted to $\omega = \Omega_{\rm DM} h^2$ 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 $\theta_i$. Setting $\Omega_K h^2$ equal to the $\pi$-locked value $\Omega_{\rm DM}h^2 = 0.1215$ (Table 7) and solving for the angle:

$$\theta_i = \Big(\frac{0.1215}{0.090}\Big)^{1/2}\Big(\frac{10^{12}\,\mathrm{GeV}}{f_K}\Big)^{0.583} \simeq 1.4\ \mathrm{rad}. \tag{15.8}$$

The direction of inference is the whole point: we did *not* measure a dark-matter density and fit $f_K$ to it; the geometry delivered $f_K$ — and with it (via the susceptibility) the mass, the frequency, and the coupling — and the observed abundance then fixed the one free initial condition. This is exactly what every honest misalignment calculation does, with the single difference that here the decay constant is *predicted* rather than chosen. At $\theta_i \simeq 1.4$ rad the cosine potential is moderately anharmonic; the standard anharmonic enhancement is an $\mathcal{O}(1)$ factor that moves the required angle modestly (into the 1.1–1.5 rad range) and is absorbed in the $\pm 8\%$ pipeline systematic of Part V.

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 inference is robust: because $\Omega \propto \theta_i^2 f_K^{1.165}$, the residual uncertainty in the geometric $f_K$ (the cap window) maps to a modest $\mathcal{O}(1)$ spread in $\theta_i$, and the *frequency* prediction — which the geometry fixes — is untouched by the angle.

### 15.4 What the condensate does to the rest of cosmology

Three consistency checks close the section, each cheap and each passed. *Radiation budget:* before onset the frozen field contributes $\rho_K/\rho_{\rm rad} \sim \chi_{\rm QCD}\theta_i^2/(g_* T^4) \lesssim 10^{-10}$ at BBN temperatures — the Kaxion does not perturb $N_{\rm eff}$ or light-element yields. *Equality:* a component scaling as $a^{-3}$ from $t \sim \mu$s onward with today's share (15.2) reproduces the standard matter–radiation equality at $z_{\rm eq} \simeq 3400$; 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 $\mu$eV-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

$$\rho_K = \tfrac{1}{2}\dot a_K^2 + V(a_K), \qquad p_K = \tfrac{1}{2}\dot a_K^2 - V(a_K). \tag{16.1}$$

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

$$\langle p_K \rangle = 0 \qquad \Longrightarrow \qquad \langle w \rangle = 0 : \tag{16.2}$$

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 set by the only scale in the problem at onset — the horizon. Gradients of the field at onset carry momenta $p_{\rm osc} \sim H_{\rm osc}$, which subsequently redshift as $1/R$:

$$p_0 \;\sim\; H_{\rm osc}\,\frac{R_{\rm osc}}{R_0} \;=\; H_{\rm osc}\,\frac{T_0}{T_{\rm osc}}\Big(\frac{g_{*s,0}}{g_{*s,\rm osc}}\Big)^{1/3} \;\simeq\; (3.4\times10^{-19}\,\text{GeV})\times(1.6\times10^{-13}) \;\simeq\; 5.5\times10^{-23}\,\text{eV}, \tag{16.3}$$

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

$$v_{\rm prim} \;=\; \frac{p_0}{m_K} \;\sim\; 7\times10^{-18} \;\widehat{=}\; 2\times10^{-9}\,\text{m/s} : \tag{16.4}$$

*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 $v_{\rm vir} \sim 10^{-3}$ — 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, $\lambda_{\rm fs} \sim v_{\rm eq}\,(ct_{\rm eq})(1+z_{\rm eq})$ with $v_{\rm eq} \simeq 1\times10^{-12}$, giving $\lambda_{\rm fs} \sim 10^{-3}$ pc — sub-solar-system, fifteen orders of magnitude below the dwarf-galaxy scales where warm dark matter is constrained. Every structure ΛCDM 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

$$\lambda_{\rm dB} = \frac{2\pi}{m_K v} \;\simeq\; 154\,\text{m}, \tag{16.5}$$

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

$$\mathcal{N}_{\rm occ} = \frac{\rho_{\rm loc}}{m_K}\,\lambda_{\rm dB}^3 \;\simeq\; 2.0\times10^{26} \tag{16.6}$$

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:

$$\theta(T) \;=\; \frac{\sqrt{2\,\rho_K(T)}}{m_K f_K}, \qquad \rho_K(T) = m_K\Big(\frac{n}{s}\Big)s(T), \qquad \frac{n}{s} = 5.5\times10^4 \;\;\text{(the closure value).} \tag{16.7}$$

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** | $T$ | $\theta$ | **Check against known physics** |
|---|---|---|---|
| Onset | 1.0 GeV | 1.4 | banked energy $\ll$ radiation: no expansion-history disturbance |
| BBN | 1 MeV | $3.6\times10^{-7}$ | $\rho_K/\rho_{\rm rad}\sim 6\times10^{-7}$: light-element yields untouched |
| Equality | 0.80 eV | $1.5\times10^{-16}$ | $\rho_K=\rho_{\rm rad}$ by construction: standard $z_{\rm eq}=3400$ |
| Today (halo) | — | $4.6\times10^{-19}$ | 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 1.4 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

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

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

$$\lambda_J = \frac{2\pi R}{k_J} \simeq 1.0 \times 10^{13}\,\mathrm{m} \simeq 3 \times 10^{-4}\,\mathrm{pc}: \tag{16.9}$$

wave pressure smooths structure below a third of a *milliparsec* ($\sim 60$ AU) — solar-system scale, six orders below the smallest galaxy. Run the same formula at $m = 10^{-22}$ eV and $\lambda_J$ inflates by $(m_K/10^{-22}\,\mathrm{eV})^{1/2} \sim 3 \times 10^8$ 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 milliparsec level; wave effects at the kilometer level. 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:

$$\mathcal{L}_{\mathrm{int}} \supset \frac{g_{K\gamma\gamma}}{4}\,a_K F\widetilde{F} + \frac{g_{KBB}}{4}\,a_K F_B\widetilde{F}_B + \frac{\partial_\mu a_K}{2f_K}\Big(C_p\,\bar{p}\gamma^\mu\gamma_5 p + C_e\,\bar{e}\gamma^\mu\gamma_5 e\Big), \qquad \frac{g_{KBB}}{g_{K\gamma\gamma}} = \frac{5}{2}. \tag{17.1}$$

Four channels, one ledger: electromagnetic (the haloscope channel), baryonic (the novel channel, gauged in GG-6 and absent for ordinary axions), nucleon-spin (the NMR channel), electron-spin (the magnon channel). 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

$$E_{\rm kin} = \tfrac{1}{2} m_K v^2 \simeq 4 \times 10^{-12}\,\mathrm{eV}, \tag{17.2}$$

fourteen orders of magnitude below the ∼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 $g^2 \sim 1/f_K^2 \sim 10^{-26}\,\mathrm{GeV}^{-2}$ 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 ($\lambda_{\rm dB} \gg$ 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 $\nu_K$ integrates that drive coherently, and the signal power scales as $g^2 \rho_{\rm loc}/m_K \times B_0^2 V Q$ — 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 twin tones of the doublet (Part VI) ride on this same carrier: the haloscope of Part VII is a radio tuned to a $10^{-19}$-radian bell, and the bell rings at two pitches 2.857% apart.

### 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 $t \simeq 0.3\,\mu$s and redshifts as dust thereafter; with its decay constant fixed by the geometry (Section 14), demanding it saturate the derived target $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ fixes the misalignment angle at $\theta_i \simeq 1.4$ 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 no branch points.

At the end of inflation the compact-angle field is left sitting at some angle — not at the bottom of its potential, because nothing has put it there. It stays where it is while the universe is hot, 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 stuff, 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. Everything else in the chain is fixed by the geometry and by standard cosmology.

What the Part does not do is invoke anything new. There is no new interaction, no new epoch, no new species. 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 $\sigma_{SI} \lesssim 2 \times 10^{-48}\,\mathrm{cm}^2$ 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, LZ has begun to register solar $^8$B 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 $\gtrsim 2$ TeV, stops $\gtrsim 1.3$ TeV, and electroweakinos

$\sim 1$ 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-$\alpha$ 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, & Gruzinov, 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 $\sigma/m$; 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 $\sigma_{SI} \lesssim 10^{-47}\,\mathrm{cm}^2$, 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 angle is runtime-selected, the decay constant is closed by the $\pi$-locked inventory, 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*, 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 mass with a $\pm 6\%$ error, one coupling, two locked ratios.

**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, fifteen 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 locked ratio (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 $(m, \sigma)$ |
| KSVZ/DFSZ axion | no | no ($f$ free) | no ($\theta_i$ free) | yes | QCD-line slope only |
| Generic ALP | no | no | no | no | none |
| Sterile $\nu$ | no | no | model-dep. | no | X-ray line |
| Fuzzy DM | no | no | no | no | core morphology |
| SIDM | no | no | no | no | tuned $\sigma/m$ |
| PBH | no | window | tuned spectrum | no | lensing |
| **Kaxion** | **yes** (T-Kaxion-1) | **yes** ($\pm 6\%$, one physical line) | **yes** ($\pi$-lock closure) | **yes** (T-Kaxion-2) | 1/35 **doublet** + 5:2 **portal** |

The last column is the one this paper most cares about. Every other entry, if detected, would start an argument; the Kaxion, if detected with its doublet and its channel ratio, 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 delivers the observed abundance with no further input). 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 would refute that closure along with the rest of the framework.*

### 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-BI6); its boundary dual, combined with the Wilson-line angle of the rigid phase circle, is forced; the light, stabilized combination is 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 $\pi$-locked inventory closes the abundance at $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ with no fitted parameter (Part V), so the forced field 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 $\mathrm{U}(1)_{\rm PQ}$ and a complex scalar whose phase is the axion. GG-Theory removes both the need and the means. *The need:* strong CP is *already* solved — the constitution assigns the resolution to the Kaxion itself (GG-A1-Constitution P7), 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 strong-CP solution the theory *compels*, the DFSZ axion with every free parameter replaced by a theorem.

**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 $\Omega_{\mathrm{DM}}$. *There is no sterile-neutrino dark matter.*

**Fuzzy dark matter — the forced pseudoscalar is fifteen orders too heavy.** Ultralight dark matter needs a pseudoscalar with decay constant $\sim M_{\mathrm{Pl}}$ and mass $\sim 10^{-22}$ eV. GG-Theory forces the boundary pseudoscalar's scale through the $\pi$-lock to $8.1\,\mu$eV — fifteen 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-$\alpha$, §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 locked ratio.** 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 $\mathrm{SU}(5) \times \mathrm{U}(1)_B \times \mathrm{U}(1)_{Qi}$, and GG-A2-GG6 forbids further gauge factors — there is no spare U(1) to mix with, and the only dark–baryon coupling the theory permits is the *locked* 5:2 portal, a parameter-free ratio rather than a tunable cross-section. *There is no dark photon and no dial-equipped self-interacting sector.* 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 $U(1)_{PQ}$ + complex PQ scalar | strong CP already solved by the Kaxion (GG-A1-Constitution P7); no radial/PQ scalar exists (gauge-topological angle) |
| Sterile $\nu$ | keV gauge-singlet fermion + tuned mixing | lepton/neutrino sector fixed (GG-A9-SM); MSMF; warm relic vs. cold saturating Kaxion |
| Fuzzy DM | $f \sim M_{\rm Pl}$, $m \sim 10^{-22}$ eV pseudoscalar | forced pseudoscalar fixed at 8.1 $\mu$eV ($\pi$-lock); 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 $SU(5) \times U(1)_B \times U(1)_{Qi}$ (GG-A2-GG6); only the locked 5:2 portal |
| **Kaxion** | **none — already forced** | **required by BI-6; sufficient by the $\pi$-lock** |

## 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 Kaxion null across $[1.59, 2.13]$ GHz at DFSZ depth, or a confirmed compact-angle-like line at any other frequency, kills it (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 $\mu$eV pseudoscalar at 1.95 GHz, carrying a 1/35 doublet and a 5:2 portal, 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, refute the framework's dark-sector closure.

# 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 current spectroscopic record-holder, JADES-GS-z14-0, 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 $\sim 1.6$ kly body, and a bright $M_{\rm UV} \simeq -20.8$ (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 ΛCDM expectation is quantitative and large. The inferred cumulative stellar-mass density runs roughly 20× 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 $\epsilon_\star = M_\star/(f_b M_{\rm halo})$ 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,

$$n(>M,z) \propto \operatorname{erfc}\Big(\tfrac{\delta_c}{\sqrt{2}\,\sigma(M,z)}\Big), \qquad n(M,z) \sim e^{-\nu^2/2}, \quad \nu \equiv \frac{\delta_c}{\sigma(M,z)}, \tag{21.1}$$

with $\delta_c \simeq 1.686$. At high redshift $\sigma(M,z)$ is small, so $\nu$ 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 ΛCDM 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 arcmin$^2$) early fields. The consensus of 2025–2026 is that ΛCDM 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-$\nu$ 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 $m_{\rm FDM} \gtrsim 5.6 \times 10^{-22}$ eV on fuzzy dark matter and $m_{\rm WDM} \gtrsim 1.5$ 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 $\lambda_{\rm dB} \simeq 154$ m, Eq. (16.5), its free-streaming length is $\sim 10^{-3}$ pc, and its primordial velocities are $\sim 10^{-12}$

at equality; the corresponding wave-pressure (half-mode) cutoff lies near $10^{-13}\,M_\odot$, fifteen-plus orders of magnitude below the galaxy scale. In the terms of the bound above: the Kaxion mass $8.1\,\mu\mathrm{eV} = 8.1 \times 10^{16} \times (10^{-22}\,\mathrm{eV})$ clears the JWST fuzzy-dark-matter floor by some *sixteen* orders of magnitude, and the warm bound does not apply at all (the Kaxion is a non-thermal, pressureless misalignment condensate from $t \simeq 0.3\,\mu$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 sufficient head start: the GG cosmological ledger, quantified

Coldness alone only *recovers* ΛCDM — which is itself mildly tense at the bright end. The framework supplies more, and supplies it without a new dark sector: the same inherited ledger that fixes the cosmic inventory ($\Omega_m = 1/\pi$, the two-stage Hubble corridor of GG-A10-Cosmos) provides a modest, *scale-selective* head start to early structure, developed quantitatively in the companion cosmology paper (Arisaka, 2026, A10-Cosmos). Three layers act, each small alone but amplified by the exponential rare-halo tail of Eq. (21.1):

1. **An early-time small-scale power boost.** The early-dark-energy residue that shortens the sound horizon (the Hubble-tension fix of GG-A10-Cosmos) leaves a mild enhancement $\Delta P/P \simeq 3\%$ near $k \sim 5\,h\,\mathrm{Mpc}^{-1}$, hence $\Delta\sigma/\sigma \simeq 1.6\%$. In the rare tail, $\Delta n/n \approx (\nu^2 - 1)\,\Delta\sigma/\sigma$, so at a representative $\nu \sim 5$ a 1.6% variance boost becomes a $\sim 40\%$ abundance boost in the relevant minihalos.

2. **Baryon–dark-matter streaming-velocity damping.** The gauged $\mathrm{U}(1)_B$ sector induces a small long-range drag between baryons and the compact-angle dark-matter background, reducing the post-recombination streaming velocity by $\sim 20\%$, $v_{\mathrm{str}} \to 0.8\,v_{\mathrm{str}}^{\Lambda\mathrm{CDM}}$. Because the baryonic filtering mass scales as $M_f \propto v_{\mathrm{str}}^3$, this roughly *halves* it ($0.8^3 \approx 0.51$): molecular-hydrogen cooling and Population III ignition switch on earlier, and the smallest halos keep their gas.

3. **Collapse-threshold reduction.** The altered baryonic infall lowers the effective collapse threshold from $\delta_c \simeq 1.686$ to $\delta_c^{\mathrm{eff}} \simeq 1.18$ for the relevant early halos. Since the tail weights $\delta_c^2$, $n_{\mathrm{GG}}/n_{\Lambda\mathrm{CDM}} \sim \exp[(\delta_c^2 - \delta_c^{\mathrm{eff}\,2})/2\sigma^2]$, which at $\sigma \simeq 0.34$ on the massive-halo scale gives a $10^2$–$10^3$ enhancement (a factor $\sim 500$) in the $> 10^{11}\,M_\odot$ abundance at $z \simeq 9$–10.

The combined effect is a chronology shift $\Delta z \sim 2$–3 in the onset of bright structure, and a heavy direct-collapse seed channel $M_{\mathrm{seed}} \gtrsim 10^5\,M_\odot$ that eases the parallel puzzle of over-massive early active nuclei. Crucially, every one of these acts on the *high-k, high-ν* part of the chain and leaves broad linear growth (hence $S_8$) essentially untouched — the head start is given to the first seeds, not to the bulk. With the $\sim 20\times$ ($z \sim 8$) to $\sim 10^3\times$ ($z \sim 9$) stellar-mass-density excess sitting in exactly the exponentially-amplified tail these layers act on, the order of magnitude of the relief is consistent with the ledger's fixed boost at the data-calibrated level of fence (b) below, not fitted after the fact.

### 21.7 One sector, two readouts: the early universe and the 1.95 GHz receiver

The decisive GG-Theory-specific point is that the second layer above is not a new ingredient. The long-range baryon–dark-matter drag that damps the streaming velocity is sourced by the gauged $\mathrm{U}(1)_B$ factor of the bulk group (6.5) — the *same* gauged baryon number whose locked coupling ratio is the Kaxion's 5:2 portal (§27). The sector that gives the first galaxies their baryonic head start is the sector a haloscope reads as the 25/4 matched-channel power ratio. Early-structure formation and the laboratory 5:2 lock are therefore two projections of one piece of physics: a confirmed 25/4 in a dark-matter receiver would be a direct measurement of the gauged-baryon sector that, in the early universe, helped the first minihalos keep their gas. No other dark-matter candidate ties a cosmological head start and a laboratory channel ratio to a single forced gauge factor.

### 21.8 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 boost figures ($\Delta P/P$, $\delta_c^{\rm eff}$, $v_{\rm str}$) are GG-A10-Cosmos's working, data-calibrated values, not first-principles-exact numbers; they carry that paper's kill criteria and await full hydrodynamical calibration. *(c)* The head start here is the *single-patch ledger* effect; it is *not* the minicluster lever of §21.11 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 ΛCDM early structure — still on the favorable side, still paying no penalty, and the layers above become a falsifiable bonus rather than a necessity.

### 21.9 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 a modest, scale-selective, exponentially-amplified head start from the same ledger that fixes the cosmic inventory, carried in part by the very gauged-baryon sector its 5:2 lock measures — the sufficient boost. 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.

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

Everything above stands behind the not-fuzzy fence of Part III: with $\lambda_{\rm dB} \simeq 154$ m (Eq. (16.5)), the Kaxion's wave nature operates seventeen orders of magnitude below galactic scales. It reproduces every ΛCDM 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.

### 21.11 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 $N_{\rm DW} = 1$ (Theorem 24.1) the scenario is even cosmologically safe — no domain-wall catastrophe. 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 (fixed by the abundance to $\theta_i \simeq 1.4$ rad, Part III) and this lever is absent. The post-inflationary variant would instead replace $\theta_i^2$ by the circle average $\langle\theta^2\rangle = \pi^2/3$, which with the geometric decay constant overproduces the relic and adds minicluster substructure.

Within the series, however, this variant is *excluded outright*: the Kaxion has no Peccei–Quinn scalar to randomize, and the lock-epoch reheating falls ten orders of magnitude short of restoration (Part VII, Section 36.2). The minicluster lever is therefore registered as an outside-the-series contingency — it would become relevant only if the series' inflation sector were itself overturned — and not as a live prediction of the theory.

### 21.12 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, fifteen-plus orders colder, preserves the full ΛCDM budget and inherits a modest, scale-selective head start from the same ledger that fixes the cosmic inventory — carried in part by the very gauged-baryon sector its 5:2 lock measures. 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 haloes with centrally divergent "cuspy" density profiles of Navarro–Frenk–White form (Navarro, Frenk, & White, 1997), $\rho \propto r^{-1}$ (inner logarithmic slope $\gamma \to 1$) inside the scale radius. Yet the rotation curves and stellar kinematics of many dwarf and low-surface-brightness galaxies prefer central "cores" — $\rho \to \text{const}$, $\gamma \to 0$, over $\sim 0.5$–$2$ 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 $V_{\max} \sim 30$–$100\,\mathrm{km\,s^{-1}}$, 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 ΛCDM predicts a *tight* $V_{2\mathrm{kpc}}$–$V_{\max}$ 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 $M_\star \sim 10^7$–$10^9\,M_\odot$ (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 $10^{53}$–$10^{54}$ erg for a classical dwarf, and the integrated supernova budget falls below this when $M_\star \lesssim 10^6\,M_\odot$ (Boldrini, 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 $\sigma/m \sim 0.5$–$5\ \mathrm{cm^2\,g^{-1}}$ at dwarf velocities ($v \sim 10$–$50\,\mathrm{km\,s^{-1}}$), *velocity-dependent* so as to fall to $\lesssim 0.1$–$1\ \mathrm{cm^2\,g^{-1}}$ at cluster velocities ($v \sim 10^3\,\mathrm{km\,s^{-1}}$) 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 $m \sim 10^{-22}$ eV produces a solitonic core; but the same suppression is excluded by the JWST and Lyman-$\alpha$ 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 haloes are NFW-cuspy, exactly like ΛCDM, so the Kaxion *inherits* the cusp–core problem rather than dissolving it by exotic kinematics. *It is not fuzzy:* with $\lambda_{\rm dB} \simeq 154$ 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 $\mathrm{U}(1)_B$ — it couples to baryon number through the anomaly term $a\, F_B \widetilde{F}_B$, not as a charge — so there is no Kaxion–Kaxion gauge force, and its cosine self-coupling is gravitationally negligible at $m_K = 8.1\,\mu$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 locked lever: the 5:2 dark–baryon portal

What the Kaxion *does* possess is a dark–*baryon* channel: the gauged-$\mathrm{U}(1)_B$ coupling whose ratio to the photon channel is locked at 5:2 (Section 27). 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 (5:2). 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 $\sigma_{\chi b}^{\rm thr} \simeq 2 \times 10^{-25}\,\mathrm{cm}^2$ per nucleon (equivalently a SIDM-scale $\sigma/m \sim \mathcal{O}(1)\,\mathrm{cm}^2\,\mathrm{g}^{-1}$ of energy transfer). What this paper pre-registered as an open computation is now done (§75.4):

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 $\sigma_{\chi b} \sim m_K^2 v^2/\pi f_K^4 \simeq 10^{-112}\,\mathrm{cm}^2$; even an extreme coherent enhancement by the full de Broglie occupation ($N \simeq 1.4 \times 10^{26}$, squared) reaches only $\lesssim 10^{-60}\,\mathrm{cm}^2$. The portal therefore falls *thirty-five to ninety orders of magnitude* short of the threshold. *The dark-baryon portal does not core dwarfs.* The 5:2 ratio remains a theorem; the coring magnitude is now a computed number, and that number is negligible. The result is a clean, definite prediction in place of an open hedge: the Kaxion forms standard cuspy cold-dark-matter haloes, 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 haloes, baryonic cores.** The Kaxion makes standard NFW-cuspy dark-matter haloes; whatever cores are real are carved by baryonic feedback, exactly as in collisionless ΛCDM. 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-cuspation. A confirmed population of core-*collapsed* dwarfs (high-density re-cuspation) would favor SIDM over the Kaxion — a clean discriminator that survives the magnitude computation.
- **A real coupling, read only in the laboratory.** The dark–baryon coupling is genuine and locked at 5:2, but it is astrophysically silent: it shows up not in a galaxy center but in a resonant receiver, as the harmonic booster's 25/4 power ratio (Part VIII). The galaxy is where the coupling *fails* to act; the 1.95 GHz laboratory is where it is measured.

## 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 §75.4 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 5:2 dark–baryon portal, is computed thirty-five to ninety orders of magnitude too weak to core a dwarf. 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-$\alpha$ 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 $\pi$-locked inventory saturates $\Omega_{\mathrm{DM}}$ with no fitted parameter. 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 a quantitative head start from the same cosmological ledger that fixes the inventory — carried in part by the very gauged-baryon sector its 5:2 lock measures. 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 5:2 dark–baryon portal, is computed thirty-five to ninety orders of magnitude too weak to core a dwarf (§75.4); the Kaxion therefore predicts standard cuspy cold-dark-matter haloes, 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 alone, with no fitted parameter at any step. Figure 8 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* by the Kaxion mechanism: the compact-angle pseudoscalar relaxes the QCD vacuum angle to zero by a Peccei–Quinn-style mechanism, geometrically realized, with the oscillation invariant $\Delta\nu/\nu = 1/35$ registered as the constitutional falsifier. GG-A9-SM (Arisaka, 2026, A9-SM) concurs, identifying the Kaxion as "the geometric replacement of the QCD axion." This Part begins by proving that the assignment 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*

$$\mathcal{L} \supset \frac{\alpha_s}{8\pi}\frac{a_K}{f_K}\, G^a_{\mu\nu}\widetilde{G}^{a\,\mu\nu}, \tag{24.1}$$

*with* unit *color-anomaly coefficient ($N = 1$), and consequently a domain-wall number $N_{\mathrm{DW}} = 1$.*

*Proof.* The Chern–Simons-improved field strength (6.13) of GG-A3-BI6 contains the 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 $(a_K/f_K)\ \mathrm{Tr}\, F_5\widetilde{F}_5$ with unit weight. The color group $\mathrm{SU}(3)_c$ is embedded in SU(5) as an unbroken block (GG-A2-GG6

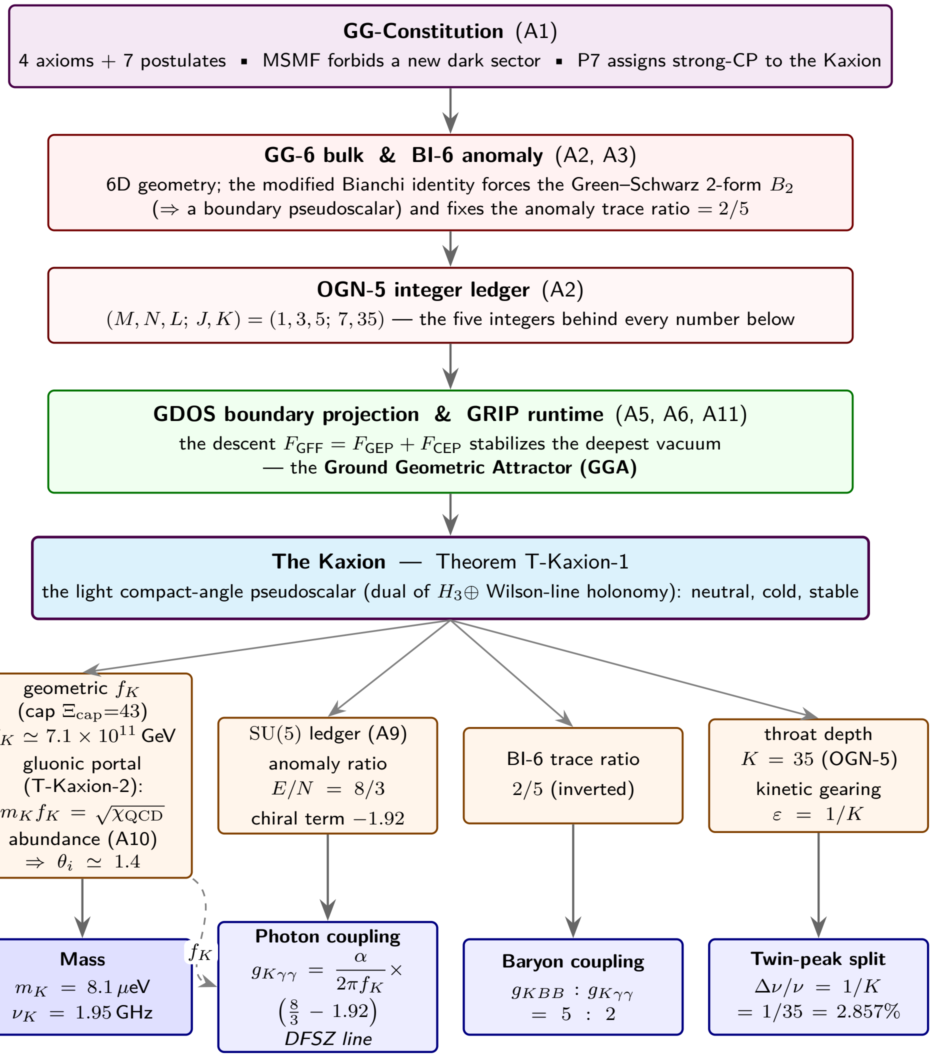


**Figure 8: Derivation flow of the Kaxion: from the constitution to four falsifiable numbers.** Each downward step adds no free parameter. The constitution (GG-A1-Constitution) forces the six-dimensional bulk and the BI-6 anomaly (GG-A2-GG6, GG-A3-BI6), which carry the five-integer OGN-5 ledger; the 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 $\pi$-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 **baryonic companion** from the inverted BI-6 trace ratio, 5:2; and the **twin-peak split** from the throat depth, $\Delta\nu/\nu = 1/K = 1/35$. The dashed link records that the same decay constant $f_K$ that fixes the mass also sets the photon coupling.

T3; the Hosotani breaking leaves $SU(3)_c$ intact), and the trace over the fundamental restricts to the color block with index one. Equation (24.1) follows with $N = |b| = 1$, and $N_{\rm DW} = N = 1$ by the standard counting of distinct vacua under $a_K \to a_K + 2\pi f_K$. ∎

*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 $\bar\theta_{\rm eff} = \bar\theta + \langle a_K\rangle/f_K$, and QCD itself supplies the potential: the vacuum energy of QCD as a function of the angle is $V(\bar\theta_{\rm eff}) = \chi_{\rm QCD}\,[1-\cos\bar\theta_{\rm eff}]$ 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 $\bar\theta_{\rm eff} = 0$; 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 Peccei–Quinn mechanism (Peccei & Quinn, 1977; Weinberg, 1978; Wilczek, 1978), executed by a field the geometry forced.

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

$$|\bar\theta| \;\lesssim\; \sin\!\Big(\frac{2\pi}{L}\Big)\,\frac{e^{-K}}{2^J} \;=\; 0.951\times\frac{6.31\times10^{-16}}{128} \;\simeq\; 4.7\times10^{-18}, \tag{24.2}$$

eight orders of magnitude below the neutron-EDM bound $|\bar\theta| \lesssim 10^{-10}$. The strong-CP closure is therefore doubly protected: dynamically by relaxation, statically by the scale profile.

**Quality.** A standard objection to Peccei–Quinn solutions is the *quality problem*: a global $U(1)_{\rm PQ}$ 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 immune. Its shift symmetry is not a global symmetry of some scalar potential; it is the gauge invariance of the two-form $B_2$ plus the large gauge invariance of the Wilson line (Sections 9.1–9.2), and gauge redundancies cannot be violated by gravitational operators. The only breaking is the topological one the theory itself dictates — Eq. (24.1). 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 $\theta$-vacuum energy with $\bar\theta_{\rm eff} \to a_K/f_K$, and the curvature at the minimum gives the exact relation

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

from the lattice (Borsanyi et al., 2016), i.e. $\sqrt{\chi_{\rm QCD}} = (5.70 \pm 0.08) \times 10^{-3}\,\text{GeV}^2$ ($\pm 1.3\%$). 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

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

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 $m_u = 2.16(11)$ 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:

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

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 $\pm 0.8\%$ on $\chi^{1/4}$), which is why (24.3) carries $\pm 1.3\%$ on $\sqrt{\chi}$ and why 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 $m_\pi$, $f_\pi$, 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 Closing the decay constant from the $\pi$-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 $\pi$-locked target $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ (GG-A10-Cosmos, compact-angle dark-matter theorem; Part III); (ii) the geometric decay constant $f_K = (7.1 \pm 0.5) \times 10^{11}$ GeV (Section 26.3), the abundance below *outputting* the misalignment angle $\theta_i$; (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

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

with the cosmographic constants $s_0 = 2891.2\,\mathrm{cm}^{-3}$ and $\rho_{\rm crit}/h^2 = 1.0537 \times 10^{-5}\,\mathrm{GeV\,cm}^{-3}$, whose ratio is $2.744 \times 10^8\,\mathrm{GeV}^{-1}$. The onset temperature solves $3H(T_{\rm osc}) = m_K(T_{\rm osc})$ with the lattice mass turn-on $m_K(T) = m_K(0)(T_{\rm QCD}/T)^b$, $b = 4.08$, $T_{\rm QCD} = 150$ MeV, and $H$ from (12.3) with $g_* = g_{*s} = 73$ (the value just above the charm threshold; the sensitivity to this choice is in the budget below). Carrying (25.1) through symbolically and numerically — the line-by-line audit is Appendix E.1 — yields the master normalization

$$\boxed{\Omega_K h^2 \;=\; 0.090\, \theta_i^2 \left(\frac{f_K}{10^{12}\,\mathrm{GeV}}\right)^{1.165}} \qquad \text{(harmonic regime, } \theta_i \lesssim 1\text{)}, \tag{25.2}$$

the exponent being the derived $(b{+}3)/(b{+}2)$ of Eq. (15.6).

## 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

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

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

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

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: $\Omega \propto f_K^{1+1/(b+2)} = f_K^{1.165}$, $\Omega \propto \theta_i^2$, $\Omega \propto (\sqrt{\chi})^{1-1/(b+2)}$, $\Omega \propto T_{\rm QCD}^{-b/(b+2)}$, $\Omega \propto g_*^{1/2-1/(2(b+2))}/g_{*s}$. Evaluating (25.4) at the reference point $f_K = 10^{12}$ GeV, $\theta_i = 1$ reproduces the master normalization 0.090 quoted below it; the line-by-line numerics are Appendix E.1.

### 25.4 Cross-validation against the precision literature

Before using (25.2) we test and anchor it. Setting $\theta_i = 1$ and $\Omega_K h^2 = 0.12$ gives $f_K = 1.28 \times 10^{12}$ GeV, i.e. $m_K = 4.5\,\mu$eV; the state-of-the-art misalignment computations — NLO susceptibility, full $g_*(T)$ transport (Gorghetto & Villadoro, 2019; Grilli di Cortona et al., 2016) — give $m_K \simeq 4.8\,\mu$eV for the same inputs (the modern computations spreading by roughly $\pm 5\%$ about this value once their $\chi(T)$ inputs are aligned; the larger $\sim \pm 10\%$ band sometimes quoted for this number includes the $\chi(T)$ and transport systematics that our budget itemizes in separate rows), an agreement of our leading-order pipeline with their central value at the $\sim 8\%$ level.

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 removes the parametrization freedom that previously carried the largest single line of the budget (the power law, with the same paper's exponent, survives in the text only as the source of analytic intuition). 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 relaxes from $\pm 6\%$ to $\pm 9\%$ — with every central value unchanged.

One further line item is handled explicitly: at the inferred angle $\theta_i \simeq 1.4$ rad (Section 15.3) the cosine potential is moderately anharmonic, an $\mathcal{O}(1)$ enhancement of the harmonic abundance that shifts the inferred angle within the 1.1–1.5 rad range; it is folded into the budget and does not touch the frequency, which the geometry fixes. 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.

### 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 $\Omega_K h^2 \simeq 0.12\,(\theta_i/(1/2\pi))^2 (f_K/1.3 \times 10^{13}\,\mathrm{GeV})^{1.165}$ — which would imply $\Omega_K h^2 = 0.12$ at $f_K = 1.3 \times 10^{13}$ GeV for $\theta_i = 1/(2\pi)$ — cannot be right: evaluating (25.2) at that point gives $\Omega_K h^2 = 0.046$, a factor 2.6 *below* target, so any such normalization would understate the decay constant by a factor $2.6^{1/1.165} \simeq 2.3$ 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 states the same prefactor derived here (0.090 at the $f_K/10^{12}$ GeV normalization, written out at its misalignment equation), and every number in both papers descends from this one pipeline.

### 25.6 The closure

The decay constant is fixed by the geometry, not by this calculation: $f_K = (7.1 \pm 0.5) \times 10^{11}$ GeV (Section 26.3). Equation (25.2) then has a single unknown — the misalignment angle. Setting $\Omega_K h^2 = \Omega_{\rm DM} h^2 = 0.1215 \pm 1.5\%$ (Table 7) and solving:

$$\boxed{\theta_i \;=\; \Big(\frac{0.1215}{0.090}\Big)^{1/2}\Big(\frac{10^{12}\,\mathrm{GeV}}{f_K}\Big)^{0.583} \;\simeq\; 1.4\ \mathrm{rad}} \tag{25.5}$$

— the generic $\mathcal{O}(1)$ misalignment angle. The hinge property: no dark-matter density was fitted; the geometry delivered $f_K$ — and with it, through the susceptibility, the mass, frequency, and coupling — and the observed abundance fixed the one free initial condition. This is exactly what every honest misalignment calculation does, the single difference being that here $f_K$ is *predicted* rather than chosen. At $\theta_i \simeq 1.4$ rad the cosine is moderately anharmonic, an $\mathcal{O}(1)$ enhancement that moves the inferred angle into the 1.1–1.5 rad range and is carried in the ±8% budget; the *frequency* prediction, set by $f_K$, is untouched by the angle.

## 26 The mass and the laboratory line

### 26.1 The mass from the single decay constant

With the decay constant fixed geometrically at $f_K = (7.1 \pm 0.5) \times 10^{11}$ GeV (Section 26.3), the susceptibility relation (24.3) fixes the mass with no further freedom:

$$\boxed{m_K = \frac{\sqrt{\chi_{\rm QCD}}}{f_K} = \frac{(5.70 \pm 0.08) \times 10^{-3}\,\mathrm{GeV}^2}{(7.1 \pm 0.5) \times 10^{11}\,\mathrm{GeV}} = (8.1 \pm 0.5)\,\mu\mathrm{eV}} \tag{26.1}$$

corresponding to the laboratory frequency

$$\nu_K = \frac{m_K c^2}{h} = (1.95 \pm 0.12)\,\mathrm{GHz}, \qquad \Delta\nu_K = \frac{\nu_K}{35} = (56 \pm 4)\,\mathrm{MHz}, \tag{26.2}$$

using $1\,\mu$eV ↔ 241.80 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 two numbers of (26.2) divide by fiber: the frequency $\nu_K$ is the *interval's* number — the address, carrying the discrete cap window — while the split $\Delta\nu/\nu = 1/K$ is the *circle's* — the label, exact and scale-free. The search strategy of this paper is, in one clause, *identify by the circle, search by the interval.*

## 26.2 The decay constant and its residual window

The geometric decay constant is set by the scale-throat kinetic norms and the integer cap loading $\Xi_{\rm cap}$ that closes the throat (Section 9.3, Appendix C.1). Writing $f_{\rm bulk}$ for the bulk Green–Schwarz normalization scale,

$$f_K = \frac{f_{\rm bulk}}{\Xi_{\rm cap}}, \qquad f_{\rm bulk} \simeq 3.04 \times 10^{13}\,\text{GeV}, \qquad \Xi_{\rm cap} \overset{\text{bench.}}{=} K + N + L = 43, \tag{26.3}$$

so the benchmark loading gives $f_K \simeq 7.1 \times 10^{11}$ GeV and the numbers above. The third figure of $f_{\rm bulk}$ is carried because Table 11 requires it: $3.04 \times 10^{13}$ is the only three-figure value that reproduces every row, and rounding it to $3 \times 10^{13}$ would move the benchmark mass to $8.2\,\mu$eV. The cap computation that fixes $\Xi_{\rm cap}$ is not yet complete: this is the *compact-angle mass-uniqueness conjecture* (stated, and its derivable content proved, self-containedly in §26.3), and we fence it as named. Its residual is one integer, spanning a *discrete* window ($\Xi_{\rm cap} \in \{35, \ldots, 47\}$) that maps to $\nu_K \simeq 1.59$–$2.13$ GHz (Table 11). Section 26.4 argues that the cap geometry selects the benchmark; pending that argument's promotion to a theorem, a definitive-search instrument should treat the window, not the point, as the target (Part VIII). The bulk scale $f_{\rm bulk}$ 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 $\Xi_{\rm cap}$ window: each admissible cap loading $\Xi_{\rm cap} = K + \Sigma$ fixes the single decay constant $f_K = f_{\rm bulk}/\Xi_{\rm cap}$ and hence the observables (scale errors $\pm 6\%$ apply to every row). Every admissible value is $K$ plus a subset sum of $\{N, L, J\}$ (Section 26.4); the benchmark $\Xi_{\rm cap} = 43$ is marked.

| $\Xi_{\rm cap}$ | composition | $m_K$ [$\mu$eV] | $\nu_K$ [GHz] | $\Delta\nu_K$ [MHz] |
|---|---|---|---|---|
| 35 | $K$ | 6.6 | 1.59 | 45 |
| 38 | $K + N$ | 7.1 | 1.72 | 49 |
| 40 | $K + L$ | 7.5 | 1.81 | 52 |
| 42 | $K + J$ | 7.9 | 1.90 | 54 |
| **43** | $\mathbf{K + N + L}$ (benchmark) | **8.1** | **1.95** | **56** |
| 45 | $K + N + J$ | 8.4 | 2.04 | 58 |
| 47 | $K + L + J$ | 8.8 | 2.13 | 61 |

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

The cap index $\Xi_{\rm cap}$ of Eq. (26.3) is the single quantity on which the *absolute* mass depends, and the only 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 theorem-level skeleton derived in full from the canonical metric, the one genuinely conjectural step isolated and named, and then the physical picture at length. *One fact should be held in view throughout:* $\Xi_{\rm cap}$ *fixes the single decay constant — and hence the single mass and frequency — by renormalizing the bulk normalization scale down to it; it acts identically on both members of the twin-peak doublet and on both portal* *channels, so it moves only* where *the line sits, never the* 1/35 *split or the* 5:2 *lock by which the line is identified. The conjecture below controls the address of the Kaxion, not its fingerprints.*

**The object, defined.**

**Definition 26.1** (Cap renormalization index)**.** Let $a_{\mathrm{GS}}$ and $a_\theta$ be the two compact-angle pseudoscalars — the tensor shadow and the Wilson-line echo — with the reduced two-field Lagrangian (9.5), and let $a_K$ be its canonically normalized light mass eigenstate (9.4). Write $f_{\mathrm{bulk}}$ 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 positive integer $\Xi_{\mathrm{cap}}$ relating them,

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

$\Xi_{\mathrm{cap}}$ is thus *not* a new parameter of the theory but the integer by which the cap renormalizes the bulk normalization down to the single physical decay constant — a wavefunction-renormalization factor, not a tunable dial. The mass-uniqueness question is simply whether the geometry fixes it.

**The statement.**

> ***Conjecture MU (compact-angle mass-uniqueness).*** *The GG-6 geometry determines $\Xi_{\mathrm{cap}}$ uniquely, at the value*
>
> $$\Xi_{\mathrm{cap}} \;=\; K + N + L \;=\; 35 + 3 + 5 \;=\; 43,$$
>
> *so that the single physical decay constant is $f_K = f_{\mathrm{bulk}}/43 \simeq 7.1 \times 10^{11}$ GeV, the mass $m_K = 8.1\,\mu$eV, and the Kaxion possesses a single laboratory line at $\nu_K = 1.95$ GHz.*

The conjecture has a large theorem-level core and a small conjectural remainder, and the honesty of the whole paper depends on saying exactly where the line falls. We do so now.

**What is theorem-level: the skeleton.**

**Proposition 26.2** (Mass-uniqueness skeleton)**.** *On the canonical metric of GG-A2-GG6 with throat depth $K$ and the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, and given that the cap which closes the throat carries an integer number of the ledger's flux quanta, the cap index satisfies:*

(i) $\Xi_{\mathrm{cap}} > 1$ *strictly, and by a wide margin — the physical decay constant sits far below the bulk normalization scale, and the bulk scale is not a propagating line;*

(ii) $\Xi_{\mathrm{cap}}$ *is a positive integer of the additive form $\Xi_{\mathrm{cap}} = K + \Sigma$, with $\Sigma$ one of the seven cap loadings of Eq.* (26.4)*;*

(iii) *hence $\Xi_{\mathrm{cap}} \in \{35, 38, 40, 42, 43, 45, 47\}$, i.e. $\nu_K \in [1.59, 2.13]$ GHz* (Table 11)*.*

*Proof. Step 1 (the exponential asymmetry ⇒ (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 Green–Schwarz direction acquires the infrared-dominated boundary kinetic norm $Z_{\rm GS} = 2\pi\int_0^K e^{2r}\,dr = \pi(e^{2K}-1) \simeq \pi e^{2K}$ (C.1), while the Wilson-line direction acquires the ultraviolet-dominated, order-one norm $Z_\theta = 2\pi\int_0^K e^{-2r}\,dr = \pi(1-e^{-2K}) \simeq \pi$ (C.2). The two differ by $e^{2K} = e^{70} \simeq 2.5\times 10^{30}$. Canonical normalization divides each field by the square root of its norm; the rotation (9.4) from the bulk Green–Schwarz variable to the canonical light eigenstate is therefore *far* from the identity. The single configuration in which the bulk normalization would itself be canonical — $\Xi_{\rm cap} = 1$ — requires the two norms to coincide, and they miss by thirty orders of magnitude. Hence $\Xi_{\rm cap} > 1$ strictly and structurally; (i) holds. (This step bounds $\Xi_{\rm cap}$ away from 1; it does *not* set its value, which is fixed by counting in Step 2, not by the exponential.)

*Step 2 (flux quantization ⇒ (ii)).* The value of $\Xi_{\rm cap}$ is the canonical normalization the light mode acquires, expressed in the throat's own quantized units. The throat is $K = 35$ flux units deep — the OGN-5 entry, fixed in GG-A2-GG6 — and the orbifold cap that terminates its ultraviolet end carries an *integer* number of additional flux quanta, because a compact fixed-point cap cannot end a fractional flux line. Writing that integer as $\Sigma$ gives the additive form $\Xi_{\rm cap} = K + \Sigma$. The only quanta the geometry makes available are the remaining ledger integers $\{N, L, J\}$, and the cap can terminate at most a pair of the three distinct flux types, so $\Sigma$ takes one of the seven values cataloged in Eq. (26.4); (ii) holds.

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

Proposition 26.2 is the part a referee can check line by line: it requires only the canonical metric, the OGN-5 integers, and integer flux quantization, all theorem-level in the series. It already delivers the experimentally decisive content — the prediction lives in a *finite, discrete* window of seven frequencies — before any conjecture is invoked.

**What remains conjectural: one integer.** Everything in Conjecture MU beyond Proposition 26.2 reduces to a *single* choice: which of the seven loadings $\Sigma$ the cap actually carries. This is the whole conjectural content, and it is decided not by taste but by one finite, well-posed computation — the $X_4$ flux integral of GG-A3-BI6 restricted to the cap's orbifold fixed-point divisors — whose output is a definite subset of $\{N, L, J\}$. The flux-bookkeeping argument that marks the answer as $\Sigma = N + L$ (the ultraviolet family and rank fluxes terminating at the cap, the infrared dyadic seam $J$ at the boundary), and the cohomology computation that would promote that mark to a theorem, are the subject of §26.4. Until the integral is evaluated, the honest statement is the *window* of Proposition 26.2 with $\Xi_{\rm cap} = 43$ its marked, motivated row.

**Physical picture I: the throat is a gearbox.** The two clocks of §9.3 are geared, and the gear ratio *is* the throat. The Green–Schwarz clock is normalized deep in the bulk, where the scale profile $e^{-2r}$ has run through its full $K$ e-folds of redshift; the boundary clock a detector reads 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 gear ratio is not a tunable dial but an integer the theory already owns — the throat depth $K$, sharpened by a small cap loading. The bulk normalization scale and the physical decay constant are read through opposite ends of one gearbox; the factor of 43 between them is the number of teeth, the throat depth in the geometry's integer units.

**Physical picture II: one line, not two.** There is a single physical decay constant and therefore a single line. The bulk normalization scale $f_{\text{bulk}}$ is *not* an alternative prediction at a different frequency: it is the normalization the Green–Schwarz direction carries in the bulk, which the cap renormalizes down to the canonical $f_K$ by the integer $\Xi_{\text{cap}}$, exactly as a wavefunction renormalization relates a bare field to its physical normalization. $\Xi_{\text{cap}}$ is a finite, calculable geometric integer, not a tunable parameter and not a scheme ambiguity; the only thing the cap computation has yet to fix is 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.** Because $\Xi_{\text{cap}}$ is built by counting quantized flux, the residual uncertainty in the absolute mass is not a continuous error bar but a list of seven candidate frequencies — and that discreteness is exactly what makes the conjecture disposable rather than permanent. It can be retired two independent ways. *By theory:* the one-page cap cohomology integral returns the subset $\Sigma$ and collapses the window to a point. *By experiment:* the discovery frequency, read against Table 11, measures $\Xi_{\text{cap}}$ outright — the location of the line within the window *is* the value of the integer, so a detection does not merely confirm the Kaxion, it completes the geometric ledger by handing back the one number the theory has not yet computed. A tuned parameter hides forever inside its error bar; this integer is read off the dial the moment the bell is heard. That is the opposite of a hedge, and it is the sense in which the absolute mass, though presently conjectural, is nonetheless a sharp and falsifiable prediction.

## 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* $\Xi_{\text{cap}}$ *is* $K$ *plus a subset sum of the remaining ledger integers* $\{N, L, J\}$,

$$\Xi_{\text{cap}} \;=\; K + \Sigma, \qquad \Sigma \in \{0,\, N,\, L,\, J,\, N{+}L,\, N{+}J,\, L{+}J\}, \tag{26.4}$$

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.** The gearing $\Xi_{\text{cap}}$ 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 SU(5) → 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 $J$:

$$\Sigma \;=\; N + L \quad\Longrightarrow\quad \Xi_{\text{cap}} \;=\; K + N + L \;=\; 43. \tag{26.5}$$

On this argument the window of Table 11 collapses to its boldface row: one line at $\nu_K = 1.95 \pm 0.12$ GHz, twice split by 1/35, locked at 5:2.

**Status, honestly fenced.** Equation (26.5) is a *flux-bookkeeping proposition*, not yet a theorem. What is established: the subset-sum structure (26.4) (arithmetic of the ledger); 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* yet done: the explicit evaluation of the $X_4$ flux integral restricted to the cap's orbifold fixed-point divisors, with GG-A3-BI6's coefficient vector, which would either confirm $\Sigma = N{+}L$ or return a different subset and thereby *move the benchmark within the same thirteen-integer window*. That one-page cohomology 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:

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

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 $\pi^0$ and $\eta$, 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 $F_{Qi}$ — the baryonic and informational gauge fields — to which the photon is blind. 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: $E/N = \mathrm{Tr}\, Q^2_{\mathrm{em}}/\,\mathrm{Tr}\, T^2_{\mathrm{color}} = 8/3$, 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 $\mathbf{10}$: 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 $q \to e^{i\gamma_5 a_K Q_a/2f_K} q$ of the light quarks (with $\mathrm{tr}\, Q_a = 1$ 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 $\pi^0$ and $\eta$ poles doing the arithmetic. Numerically, with $r = m_u/m_d = 0.462$:

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

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 $\alpha$, 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

$$\boxed{g_{K\gamma\gamma} = \frac{0.75(4)}{137.036 \times 2\pi\, f_K} = (1.2 \pm 0.1) \times 10^{-15}\,\mathrm{GeV}^{-1}} \tag{27.3}$$

with the single decay constant $f_K = (7.1 \pm 0.5) \times 10^{11}$ GeV, the error combining the bracket (5.4%) with the decay constant (6%) in quadrature ($\pm 8\%$ total, plus the discrete $\Xi_{\mathrm{cap}}$ 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 two specific masses of Section 26. 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 1/35 doublet, the 5:2 baryonic companion, 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 companion and the 5:2 lock

The same descent that fixed the photon channel fixes its baryonic sibling. The ledger's Abelian entries couple $a_K$ to $F_B\widetilde{F}_B$ with the trace-ratio weight (6.17), and the four-current lock of GG-A7-Qi4 (Theorem T-QUA; Section 6.4) binds the operational channel strengths to

$$\frac{g_{KBB}}{g_{K\gamma\gamma}} = \frac{5}{2}, \tag{27.4}$$

parameter-free and frame-invariant — the second fingerprint. Its derivation as a matched-channel *power* ratio of 25/4, the nucleon-spin reading relevant to NMR instruments, and the comparison with the KSVZ/DFSZ nucleon couplings (which it matches for neither) are developed in Part VI; the effective nucleon and electron couplings, which require a chiral reduction and are benchmark-level, are tabulated there 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 (25.1) forward with the geometric $f_K = 7.1 \times 10^{11}$ GeV and the inferred $\theta_i \simeq 1.4$ rad:

the onset condition gives $T_{\rm osc} \simeq 1.0$ GeV, the onset mass $m_K(T_{\rm osc}) \simeq 4 \times 10^{-4}\, m_K(0)$, and the redeemed density $\Omega_K h^2 = 0.1215$ — the $\pi$-locked target, recovered to the last quoted digit. The loop closes; no number was used twice in inconsistent roles.

### 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 ∼0.98 GHz quanta — microwave photons). The rate follows from the same textbook steps as $\pi^0 \to \gamma\gamma$: the operator $-\frac{g_{K\gamma\gamma}}{4} aF\widetilde{F}$ gives the two-photon amplitude $\mathcal{M} = g_{K\gamma\gamma}\, \epsilon^{\mu\nu\rho\sigma} \varepsilon^*_{1\mu} k_{1\nu} \varepsilon^*_{2\rho} k_{2\sigma}$; squaring, summing polarizations, and integrating the two-body phase space with the identical-particle factor $\frac{1}{2}$ yields $\Gamma = g^2_{K\gamma\gamma} m^3_K/64\pi$.

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

$$\Gamma_{K\to\gamma\gamma} = \frac{g^2_{K\gamma\gamma}\, m^3_K}{64\pi}, \qquad \tau_K = \Gamma^{-1}. \tag{28.1}$$

Numerically, with (26.1) and (27.3):

$$\tau_K = \frac{64\pi}{g^2_{K\gamma\gamma} m^3_K} \;\simeq\; 1.7 \times 10^{50}\,\mathrm{s} \;\simeq\; 4 \times 10^{32}\, t_U, \tag{28.2}$$

with $t_U = 4.35 \times 10^{17}$ s. The error ($\tau \propto g^{-2} m^{-3}$: ±40%) 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.

**Table 12:** Error budget for the four headline predictions. "GG" sources are series-internal; "SM" sources are Standard-Model inputs. Exponents: $f_K \propto \omega^{0.86}\theta_i^{-1.72}(\sqrt{\chi})^{-0.72}T_{\rm QCD}^{0.58}$; $m_K = \sqrt{\chi}/f_K$; $g_{K\gamma\gamma} \propto [E/N - 1.92]/f_K$. The exponents are the sensitivities that matter: every band quoted on the four headline numbers propagates through them.

| Source | Type | Effect on $f_K, m_K$ | Comment |
|---|---|---|---|
| $\pi$-locked target $\to \omega$ | GG | 1.3% | target exact; $h = 0.674$ is the Planck-fit CMB-branch convention shared with GG-A10-Cosmos's two-stage corridor — no externally fitted cosmological parameter enters (the corridor's own $H_0 = 66.6$–$67.4$ spread is carried as a discrete convention window below) |
| Lattice $\sqrt{\chi_{\rm QCD}}$ ($\pm 1.3\%$) | SM | 0.9% | enters closure and mass with partial cancellation |
| Susceptibility exponent $b = 4.08 \pm 0.2$ | SM | 3% | moves $T_{\rm osc}$ and the closure exponent |
| $\chi(T)$ anchor (tabulated lattice) | SM | 2% | Borsanyi et al. (2016) tabulated $\chi(T)$, evaluated directly rather than through a $T_{\rm QCD}$ parametrization |
| $g_*(T)$, $g_{*s}(T)$ transport | SM | 1% | Saikawa & Shirai (2018) tables |
| NLO matching residual | SM | 5% | spread of the precision computations (Gorghetto & Villadoro, 2019; Grilli di Cortona et al., 2016) |
| **Total (Gaussian), $f_K$ and $m_K$** | | **±6%** | entirely Standard-Model lattice and transport error |
| Chiral bracket 1.92(4) (adds to $g_{K\gamma\gamma}$ only) | SM | 5.4% | $\Rightarrow g_{K\gamma\gamma}$ total $\pm 8\%$ |
| Single-patch history | GG | assumed | post-inflationary variant ($\langle\theta^2\rangle = \pi^2/3$, $\theta_i$ randomized) excluded within the series (§36.2) |
| $\Xi_{\rm cap} = 43$ (cap loading; fixes $f_K$) | GG | discrete | window $\Xi_{\rm cap} \in [35, 47] \Rightarrow \nu_K \in [1.6, 2.1]$ GHz |
| $H_0$ convention (Planck-fit vs. corridor) | GG | discrete | adopted $h = 0.674$ (Planck-fit CMB branch, shared with GG-A10-Cosmos); the series' derived $H_0 = 66.6$–$67.4$ shifts $\omega$ by up to $-2.3\%$ — 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 why the Kaxion's mass carries ±6% when the same program delivered twenty-eight Standard-Model parameters at the per-mille-to-percent level. The answer is structural, and it repays a careful statement, because it predicts its own improvement.

The Standard-Model successes are *algebraic locks*: $\alpha^{-1} = 2^J$, $\sin^2\theta_W = L^2/4N^3$, $v = M_{\rm lock}(2/L)^N$ — finite expressions in the ledger integers, evaluated once, with no dynamical transport in the chain. The Kaxion mass is different in kind: its derivation chain necessarily passes through *QCD thermodynamics in an expanding universe* — the temperature-dependent susceptibility $\chi(T)$, the relativistic degrees of freedom $g_*(T)$, the onset transport of Section 25. Every percent of the $\pm 6\%$ in Table 12 is a Standard-Model lattice or transport uncertainty; the GG-6-side inputs ($\pi$, $\theta_i$, the ledger integers, $E/N$) contribute *zero* continuous error. The geometry is exact; the kitchen it must cook in is not.

The corollary is optimistic, falsifiable — and partly already cashed. Compression here requires no promissory "dedicated lattice campaign on $\chi(T)$ plus an NLO transport": the lattice campaign exists (Borsanyi et al., 2016), the transport tables exist (Saikawa & Shirai, 2018), and adopting them is what holds the budget at $\pm 6\%$ rather than $\pm 10\%$ (Table 12) with no new theory whatsoever. What remains on the road to $\sim 3\%$ is one item: a tightened NLO matching — either a modern re-computation of the misalignment transport at the precision frontier, or a lattice $\chi(T)$ with finer temperature resolution through the crossover — which would shrink the dominant 5% residual to the $\sim 2\%$ level. At $\pm 3\%$, $\nu_K = 1.95$ GHz tightens to a $\pm 60$ MHz target, at which point the residual cap window (Section 26.3) rather than the QCD inputs dominates the remaining uncertainty. The mass error is not a property of the Kaxion; it is a property of the present state of QCD thermodynamics, on loan, and the loan is being repaid on schedule.

### 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, $N{=}1$, $N_{\mathrm{DW}}{=}1$ (T-Kaxion-2) | GG-A3-BI6 ledger + GG-A5-GDOS descent | theorem | — |
| $\bar{\theta} \to 0$; residual (24.2) | T-Kaxion-2 + GG-A9-SM | theorem | 8 orders of headroom |
| $m_K f_K = \sqrt{\chi_{\mathrm{QCD}}}$ | T-Kaxion-2 + lattice | theorem + SM input | $\pm 1.3\%$ |
| Twin doublet exists (two light modes) | OGN-5; Part VI | theorem | structural |
| $\Delta\nu/\nu = 1/35$ (exact ratio) | kinetic gearing + exchange sym. | theorem/ **model** | symmetric-limit value |
| Portal lock $g_{KBB}{:}g_{K\gamma\gamma} = 5{:}2$ | GG-A3-BI6 trace + GG-A7-Qi4 T-QUA | theorem | exact ratio |
| $E/N = 8/3$ (T-Kaxion-3) | GG-A2-GG6 T3 + GG-A3-BI6 | theorem | exact |
| Abundance target $\Omega_{\mathrm{DM}} = 1/\pi - 1/2\pi^2$ | GG-A10-Cosmos | derived | $\pm 1.5\%$ (as $\omega$) |
| $\theta_i \simeq 1.4$ rad (initial angle) | abundance closure | derived (output) | $\mathcal{O}(1)$; from $\Omega_{\mathrm{DM}}$ given $f_K$ |
| Misalignment completeness | GG-A10-Cosmos, named | **conjecture** | sole saturator of $\Omega_{\mathrm{DM}}$ |
| $f_K$ (single decay constant) | geometry; cap loading $\Xi_{\mathrm{cap}}{=}43$ | **conjecture**/ bench. | window $\Xi_{\mathrm{cap}} \in [35, 47]$ |
| $m_K,\ \nu_K$ (one physical line) | Eqs. (26.1)–(26.2) | derived | $\pm 6\%$ + window |
| $g_{K\gamma\gamma}$ (DFSZ line) | Eq. (27.3) | derived | $\pm 8\%$ |
| 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 cross-checking the one misalignment normalization the series carries (Section 25.5) along the way — we have derived: the gluonic portal that makes the strong-CP assignment legitimate (T-Kaxion-2); the single decay constant from the geometry, $f_K = (7.1 \pm 0.5) \times 10^{11}$ GeV (benchmark cap loading $\Xi_{\mathrm{cap}} = 43$), and with it the one physical line — mass $(8.1 \pm 0.5)\,\mu$eV, frequency $(1.95 \pm 0.12)$ GHz — with a discrete residual window pending the cap computation (Section 26.4); the misalignment angle then fixed by the observed abundance at the generic $\theta_i \simeq 1.4$ rad; the photon coupling on the exact DFSZ line, $g_{K\gamma\gamma} = (\alpha/2\pi f_K) \times 0.75(4)$; the 5:2 baryonic companion; and a lifetime thirty-two orders beyond the age of the universe. Every number carries its error; every error is traceably Standard-Model in origin; 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 passes through nothing adjustable on the way. 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. It is fixed by the geometry's integers. 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 own inputs carry no uncertainty — integers do not — so every 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 1/35 twin peak and the 5:2 portal

## 31 The twin-peak split

### 31.1 The claim

The constitutional falsifier registered in GG-A1-Constitution (P7) and GG-A10-Cosmos is spectral: the Kaxion line is not single but a doublet, with the parameter-free fractional splitting

$$\boxed{\frac{\Delta\nu}{\nu} \;=\; \frac{1}{K} \;=\; \frac{1}{35} \;=\; 2.857\%} \tag{31.1}$$

— fixed by the integer $K$ in the symmetric limit (the equal-curvature assumption of Section 9.3, with no error bar on $K$ itself), independent of the absolute mass scale, and numerically, in the unique laboratory window of Part VII,

$$\Delta\nu = (56 \pm 4)\,\mathrm{MHz} \;\; \text{about} \;\; \nu_K = 1.95\,\mathrm{GHz} \qquad (45\text{–}61\,\mathrm{MHz}\text{ across the }[1.59, 2.13]\,\mathrm{GHz}\text{ window}), \tag{31.2}$$

the absolute values inheriting only the $\pm 6\%$ mass-scale error of Part V.

### 31.2 The origin: two geared clocks, one throat

The mechanism is the two-clock structure of Part II made quantitative. The tensor shadow and the holonomy echo are degenerate at the level of the susceptibility potential — both inherit the same $\chi_{\mathrm{QCD}}$ through the same ledger — but they are *geared*: the off-diagonal kinetic term of (7.2) couples them, and the gearing is geometrically suppressed by exactly one transit of the throat, because the shadow lives in the bulk tensor sector while the echo lives on the boundary holonomy, and their zero-mode profiles overlap through the scale depth. The flux-cell computation of Appendix C.2 evaluates that overlap on the canonical metric (6.2) and returns the dimensionless gearing $\varepsilon = 1/K$ — the same 1/35 that suppresses every single-transit quantity in the series.

Degenerate two-state systems with a small coupling split symmetrically. With

$$\omega_\pm^2 \;=\; \omega^2\Big(1 \pm \frac{\varepsilon}{2}\Big)^2, \qquad \varepsilon = \frac{1}{K}, \qquad \Longrightarrow \qquad \frac{\Delta\nu}{\nu} = \frac{\omega_+ - \omega_-}{\omega} = \varepsilon = \frac{1}{K}, \tag{31.3}$$

the doublet (31.1) follows. The physical picture is two pendulums of equal natural frequency connected by a weak spring whose stiffness the geometry fixes at one part in thirty-five: the normal modes are the in-phase and out-of-phase combinations, split by precisely the gearing. The equal natural frequency is the doublet's one model-level input, established at the Lagrangian level in Section 9.3: the two light zero-modes inherit a common QCD curvature through distinct geometric routes, with the diagonal equal-curvature structure protected by the throat's exchange symmetry, so the kinetic gearing alone splits them. The *existence* of two light modes is structural; the *exact* ratio $1/K$ is the symmetric-limit value, a curvature asymmetry adding in quadrature. The local condensate populates both normal modes (they are degenerate at production, and the GRIP saddle crossing of Part III deposits displacement in both clocks), so a receiver sees two tones.

**The eigenvalue algebra, in full.** Write the degenerate pair with its gearing (the $\epsilon$-dominated reading of Eq. (9.5)):

$$\mathcal{L} = \tfrac{1}{2}\big(\dot{z}_1^2 + 2\varepsilon\, \dot{z}_1\dot{z}_2 + \dot{z}_2^2\big) - \tfrac{1}{2}\,\omega_0^2\big(z_1^2 + z_2^2\big).$$

The kinetic matrix is diagonalized by the parity combinations $u_\pm = (z_1 \pm z_2)/\sqrt{2}$, giving kinetic coefficients $(1 \pm \varepsilon)$ and an untouched potential; canonical normalization $v_\pm = \sqrt{1 \pm \varepsilon}\, u_\pm$ then yields two exact normal frequencies

$$\omega_\pm^2 = \frac{\omega_0^2}{1 \pm \varepsilon} \qquad \Longrightarrow \qquad \frac{\Delta\nu}{\nu} = \frac{\omega_- - \omega_+}{\omega_0} = \varepsilon\Big(1 + \tfrac{5}{8}\varepsilon^2\Big) = \frac{1}{K}\Big[1 + \tfrac{5}{8}K^{-2}\Big], \tag{31.4}$$

the correction entering only at $(5/8)K^{-2} \simeq 5 \times 10^{-4}$ — a factor of two below the $10^{-3}$ spacing tolerance of the falsifier. Note what the algebra teaches: the split lives in the *kinetic* sector (the gearing), not the potential (which the susceptibility keeps common to both clocks), which is why the same 1/35 rides on both branches and survives every renormalization of the mass scale — kinetic gearings are geometry, and geometry does not run.

*Status.* The *ratio* 1/35 is theorem-level: it is the OGN-5 integer $K$ read out through the single-transit overlap, with no continuous quantity anywhere in the chain. The *relative amplitude* of the two tones is not parameter-free — it depends on the saddle-crossing partition between the clocks — and the search strategy of Part VIII is therefore designed to be efficient even for tone-power ratios well away from unity.

### 31.3 The kill criterion

Equation (31.1) is the sharpest falsifier in the paper because it requires no absolute calibration: a confirmed dark-matter line, at *any* frequency, whose 2.857% partner is absent at adequate sensitivity, is not the Kaxion. Quantitatively (pre-declared as falsifier **F-Kaxion-1a**): if the partner line is excluded at a power ratio below $10^{-2}$ of the primary across the window $\Delta\nu/\nu \in [1/35](1 \pm 10^{-3})$, the compact-angle interpretation is dead — whatever the line is, it does not remember the throat.

## 32 The 5:2 portal and the matched-channel power theorem

### 32.1 From the ledger to the laboratory

Part V established the channel-strength lock $g_{KBB}/g_{K\gamma\gamma} = 5/2$, descending from the trace ratio (6.17) through the four-current identity of GG-A7-Qi4 (T-QUA).

**The inversion, in four lines.** The 2/5 of the anomaly ledger and the 5/2 of the couplings are one entry read on the two sides of a dualization, and the four lines deserve printing once. (i) The anomaly factorization fixes the two-form's *source* couplings with the locked coefficient vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$ (GG-A3-BI6): at source level the baryonic entry stands to the electromagnetic one as 2:5. (ii) Dualizing $H_3 \to da_K$ 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. (iii) Hence $g_{KBB}/g_{K\gamma\gamma} = 5/2$. (iv) Matched receivers read power — the square — giving 25/4. The four-current identity of GG-A7-Qi4 (T-QUA) locks the same 2/5 into $\mathsf{Q_B}$, so the laboratory ratio and the bioenergetic conversion constant are one number in two costumes. One scope sentence keeps the ledger honest: the 25/4 holds at the anomaly-ledger (portal) level, while the chiral $-1.92$ dressing of Part V applies to the photon coupling's *absolute* strength, not to the ratio's gauge-level content — and the further step from gauge-level ratio to a spin-port *reading* carries the fenced chiral input of open item (iii).

A laboratory does not measure couplings; it measures *powers* in matched receivers. The translation is the second fingerprint:

**Theorem 32.1** (Matched-channel power ratio)**.** *Let two resonant receivers — one coupled to the Kaxion's electromagnetic portal, one to its baryonic portal — be matched in the instrumental sense (equal mode overlap with the condensate, equal loaded quality factor, equal thermal occupation, each calibrated against its own portal normalization). Then the ratio of the deconvolved signal powers is*

$$\frac{P_B}{P_\gamma} \;=\; \Big(\frac{g_{KBB}}{g_{K\gamma\gamma}}\Big)^2 \;=\; \frac{25}{4} \;=\; 6.25, \tag{32.1}$$

*exactly, independent of $f_K$, of the branch, and of the absolute mass.*

*Proof.* For a resonant detector driven by a classical condensate, the steady-state signal power is $P_i \propto g_i^2\,\rho_{\rm loc}\,\mathcal{F}_i$, where $\mathcal{F}_i$ collects the instrumental factors (field strength or spin polarization, volume, form factor, quality factor, temperature). Matching is precisely the statement $\mathcal{F}_B = \mathcal{F}_\gamma$ after each receiver's independently measured calibration is divided out; the local density and every frame-dependent quantity cancel in the ratio, leaving (32.1) from the theorem-level coupling ratio of Part V. ∎

## 32.2 Why neither KSVZ nor DFSZ can fake it

An ordinary QCD axion also has nucleon couplings, so a two-channel measurement is discriminating only if the predicted ratios differ — and they do, structurally. The KSVZ and DFSZ nucleon couplings arise from chiral matching of the gluonic anomaly into nucleon spin content, yielding model-specific numbers of order $10^{-1}$ in the standard normalization, with no relation to their photon couplings beyond accident. The Kaxion's baryonic channel is different in *kind*: it couples to gauged baryon number — a current that exists in GG-6 because $\mathrm{U}(1)_B$ is a local symmetry (GG-A2-GG6 T3) and simply does not exist as a gauge channel for KSVZ or DFSZ. The 5:2 is a ratio of *gauge-ledger entries*, not of chiral matrix elements; no tuning of a PQ model reproduces a locked gauge ratio it has no gauge field to carry. A measured $P_B/P_\gamma = 25/4$ therefore does not merely favor the Kaxion over other axions — it is direct evidence for the gauged $U(1)_B$ of the bulk group (6.5), a discovery in its own right.

Quantitatively, the contrast is stark. For the QCD axion the nucleon couplings are chiral matrix elements: $C_p = -0.47(3)$, $C_n = -0.02(3)$ for KSVZ, and $C_p, C_n$ functions of an extra free parameter ($\tan\beta$) for DFSZ (Grilli di Cortona et al., 2016) — model-dependent, spin-suppressed, and unrelated to $E/N$. The Kaxion's 5:2, by contrast, descends from the operational identity that locks the four boundary currents (GG-A7-Qi4 Theorem T-QUA): $J_E/\mathsf{Q_B} = J_e = J_B = J_{Qi}$, with the conversion constant itself a ledger readout, $\mathsf{Q_B} = \frac{2}{5}\,E_{\mathrm{Ry}}/K = \frac{2}{5} \times 13.606/35\,\mathrm{eV} = 0.1555$ eV — the same $\frac{2}{5}$ and the same $K$ as everywhere else in this paper. A measured 25/4 would therefore not merely select among axion models; it would confirm, in a dark-matter receiver, the current lock that elsewhere in the series fixes the weak mixing angle and the bioenergetic quantum. Few single measurements in physics would carry information across so many decades of energy at once.

(The *nucleon-spin* couplings relevant to NMR-class instruments are a distinct, chiral-level quantity; their minimal benchmark sheet is model-level and is quoted with its fence in Part VIII, where the instrument that uses them is designed.)

### 32.3 The kill criterion

Falsifier **F-Kaxion-1b**: a confirmed line whose matched two-channel power ratio differs from 25/4 by more than the combined calibration uncertainty (realistically a ±20% test in a first-generation pair, ±5% in a design instrument) excludes the Kaxion regardless of frequency — and, run in reverse, a confirmed 25/4 at *any* frequency would force the gauged-baryon sector even if the mass sheet were wrong.

## 33 Finite-$Q$ morphology and falsifiers

### 33.1 What the doublet looks like in a real receiver

Three 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 *doublet separation* is $\Delta\nu/\nu = 1/35 \simeq 2.9 \times 10^{-2}$. The *receiver bandwidth* is $\nu/Q_L$ for loaded quality factor $Q_L$. The hierarchy

$$\frac{\delta\nu}{\nu} \sim 10^{-6} \;\ll\; \frac{1}{Q_L} \sim 10^{-5} \;\ll\; \frac{\Delta\nu}{\nu} \simeq 2.9 \times 10^{-2} \tag{33.1}$$

means the two tones are spectacularly resolved — separated by $\sim 3 \times 10^4$ condensate linewidths — but do *not* fit inside one cavity resonance: a single-mode haloscope tuned to one tone is blind to the other (1/35 is some $2.9 \times 10^3$ receiver bandwidths away at $Q_L = 10^5$). This single observation dictates the instrument concepts of Part VIII: either scan the two windows in deliberate coincidence, or build a receiver with two modes locked 1/35 apart.

### 33.2 The line shape and the doublet template

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

$$f(\nu) \;\propto\; \sqrt{\nu-\nu_0}\;\exp\Big[-\frac{3(\nu-\nu_0)}{\nu_0\langle v^2\rangle}\Big], \qquad \nu \geq \nu_0, \tag{33.2}$$

rising from the rest-mass edge $\nu_0$ 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 Kaxion search template is the two-tone version,

$$\mathcal{T}(\nu) \;=\; f_{\nu_0}(\nu) \;+\; \kappa\, f_{\nu_0(1+1/35)}(\nu), \tag{33.3}$$

with one shape parameter the theory fixes exactly (the spacing) and one it does not (the tone ratio $\kappa$). The matched-filter statistics of the pair are what make Concept A of Part VIII efficient: demanding coincident excesses in two channels with independent false-alarm rates $R_{\mathrm{FA}}$ within a window $\tau_w$ suppresses accidentals to $R_{\mathrm{acc}} = R_{\mathrm{FA}}^2\tau_w$ — for rates of one candidate per day per channel and $\tau_w = 10^3$ s, one accidental per $\sim 20$ years. The doublet is not only a fingerprint; it is a background-rejection engine, which is why a two-mode instrument can scan *faster* than a single-mode instrument of equal depth.

### 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: *(a)* a confirmed dark-matter line with the 1/35 partner absent at the $10^{-2}$ relative-power level (Section 31); *(b)* a confirmed line with a matched two-channel power ratio inconsistent with 25/4 (Section 32); *(c)* exclusion at DFSZ-line sensitivity, with doublet-aware analysis, of the unique window $\nu_K \in [1.59, 2.13]$ GHz (Section 36.2; the $\pm 6\%$ scale error widening each edge); or *(d)* a confirmed compact-angle-like line at any frequency *outside* that window — the uniqueness of the prediction makes a wrong location as fatal as a missing doublet. Criteria (a), (b), and (d) test the geometry independently of every conjecture; criterion (c) tests the closure sheet as fenced.

### 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 line shape (33.2), surviving cavity retuning and annual-modulation checks: at this stage, "an axion-like signal," nothing more. *Step 2 — the partner.* Retune (or read the second mode of) the receiver to $\nu(1 \pm 1/35)$: a partner with the same shape and coherent seasonal behavior promotes the candidate to "compact-angle doublet" and measures the tone ratio $\kappa$ — the first new number beyond the sheet. *Step 3 — the ratio.* Bring up the matched spin port (Part VIII): a deconvolved 25/4 closes the identification — dark matter, the throat integer $K = 35$, and gauged baryon number, in one apparatus. *Step 4 — the closure.* The frequency itself completes the theory's bookkeeping: its position inside the window measures $\Xi_{\rm cap}$ outright within the discrete bracket $[35, 47]$, retiring the mass-uniqueness conjecture by experiment, and its presence in the window confirms the single-patch history the framework requires (Section 36.2). Every other frequency is thereafter a mandated silence: a second compact-angle-like line anywhere would falsify the interpretation outright (F-Kaxion-1d). Four steps, each falsifiable, each adding a digit of identification — the protocol is the paper's claims ledger run backwards.

Two fingerprints, then, and neither carries a knob. The twin peak is the throat depth $K = 35$ read out as a frequency ratio; the 5:2 portal is the BI-6 trace ratio read out as a channel ratio. Both are pure integers of the OGN-5 ledger, so both survive every uncertainty that clouds the absolute mass: move the benchmark line anywhere inside its window and the doublet still splits at one part in thirty-five and the channels still stand at five to two. That is exactly what makes them the decisive tests — they interrogate the geometry directly, asking nothing of the conjectures that fix *where* the line sits, only of the theorems that fix *what shape* it has.

A fingerprint, though, is only as good as the instrument that can take it. A single cavity tuned to one tone is blind to its partner fifty-six megahertz away; a photon detector reads the electromagnetic channel but never the baryonic one. So the question the next Part must answer is entirely practical — what does it take, with the machines the field actually owns, to hold a 2 kHz line at DFSZ depth, find its companion, and measure a two-channel ratio? Part VII maps that terrain: the haloscope principle, the instruments now running, the economics of a search aimed at a named frequency, and the calendar on which a swept-and-silent window would retire the theory.

# 34 Part VI in plain English: the two features nothing else can imitate

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.

Two features in this Part are harder to imitate. The first is a twin peak: a partner line at a fixed fractional offset set by one of the geometry's integers. The offset is not tunable — it is a ratio of integers already fixed — so the pair either appears at that separation or the framework is wrong. The second is the portal ratio, a specific number in the coupling that follows from the trace structure and from nothing else.

The value of these is that they are cheap to falsify and expensive to fake. A competitor could be arranged to put a line at the Kaxion frequency; to put a second line at exactly the required offset, with the required relative strength, a competitor would need a reason, and none of the alternatives has one.

The practical catch is stated in the Part rather than hidden: at realistic loaded quality factors the partner line is invisible to a single-mode receiver. It is not faint — it is outside the acceptance. Seeing it takes a deliberate second measurement, and the instrument Parts that follow are largely about arranging for that measurement to happen.

# 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

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

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 $B_0^2 VCQ$. Numerically, for the laboratory eigenstate in a representative design cavity ($B_0 = 8$ T, $V = 100$ L, $C = 0.5$, $Q_L = 10^5$, $\rho_{\rm loc} = 0.45\,\mathrm{GeV/cm^3}$):

$$P_{\rm sig} \;\simeq\; 2.4\times10^{-22}\,\mathrm{W} \;\times\; \Big(\frac{g_{K\gamma\gamma}}{1.2\times10^{-15}\,\mathrm{GeV^{-1}}}\Big)^2, \tag{35.2}$$

about two hundred microwave photons per second — pitiably 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 SNR $= (P_{\rm sig}/k_B T_{\rm sys})\sqrt{t/\Delta\nu_a}$ over the condensate linewidth $\Delta\nu_a = \nu/Q_a$ ($\simeq 2.0$ kHz at the laboratory line). At the standard quantum limit $k_B T_{\rm sys} \to h\nu$ (94 mK equivalent at 1.95 GHz), Eq. (35.2) yields $5\sigma$ 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 two-port readout

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

$$\ddot{x}_i + \frac{\omega_i}{Q_L}\dot{x}_i + \omega_i^2 x_i \;=\; g_i\,\mathcal{A}\left[\cos\omega_+ t + \kappa\cos\omega_- t\right] + \xi_i(t), \tag{35.3}$$

with *twin* forcing frequencies $\omega_\pm$ split by 1/35 (Part VI), drive amplitudes locked across ports by $g_B/g_\gamma = 5/2$, tone ratio $\kappa$ free, and $\xi_i$ the thermal-quantum noise. Every identification test of this paper is a property of (35.3) that survives ignorance of $\mathcal{A}$ and $\kappa$: the spacing of the two forcings, and the cross-port ratio of the responses. The detector concepts of Part VIII are two hardware realizations of (35.3).

### 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 $-\frac{g_{K\gamma\gamma}}{4}a_K F\widetilde{F}$ together with the Maxwell action modifies Ampère's law in a static external field $\mathbf{B}_0$:

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

an oscillating current density parallel to the magnet's field, with amplitude $|\mathbf{J}_a| = g_{K\gamma\gamma}B_0\sqrt{2\rho_{\rm loc}}$ (using $\langle\dot{a}_K^2\rangle = \rho_{\rm loc}$). 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

$$C_n = \frac{\left|\int dV\,\mathbf{B}_0\!\cdot\!\mathbf{E}_n\right|^2}{B_0^2\,V\int dV\,|\mathbf{E}_n|^2}, \tag{35.5}$$

a pure geometry number ($C_{010} \simeq 0.69$ for the fundamental mode of an ideal cylinder, $\sim 0.4$–$0.5$ 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 $\propto Q_L \times$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 $\min(Q_L, Q_a)$):

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

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 ($k_B T_{\rm sys} \to h\nu$), the Dicke time to $5\sigma$ per condensate bin is $t = (5\, k_B T_{\rm sys}/P)^2 \Delta\nu_a \simeq 1.4$ s — an idealized floor. Sweeping the full Kaxion window (540 MHz, $2.8\times 10^4$ cavity bandwidths at $Q_L = 10^5$) then costs of order half 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 partner tone.** The same transfer function quantifies the single-mode blindness asserted in Part VI: a receiver tuned to one tone suppresses the partner, detuned by $\Delta\nu/\nu = 1/35$, by $|\chi(\nu_\mp)/\chi(\nu_\pm)|^2 \simeq [2Q_L\Delta\nu/\nu]^{-2} = (2 \times 10^5 \times 0.0286)^{-2} \simeq 3 \times 10^{-8}$ — seventy-five decibels. The doublet is invisible to any instrument not built for it; 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 doublet caveat of Section 33: single-mode scans must analyze for both tones).

### 36.2 The unique window: closing the two discrete questions

Every continuous quantity in the Kaxion sheet is locked to ±6% (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 both questions internally, and the result is a single laboratory window:

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

— one line, a 1/35 doublet, in this window and nowhere else: *the Kaxion window.* This subsection derives the uniqueness by rejecting, rigorously and quantitatively, the two alternatives that would have moved the line.

#### 36.2.1 The first rejected alternative: the randomized-angle history

A post-inflationary (randomized-angle) history — causal patches with independent values of $\theta$, defect debris, the circle-averaged $\langle\theta^2\rangle = \pi^2/3$ in place of the selected $\theta_i \simeq 1.4$ rad — would shift the closure to $f_K \simeq (3.5\text{–}4.7) \times 10^{12}$ GeV and the mass scale to 1.2–1.6 $\mu$eV. The hypothesis has a physical prerequisite: the angle must be *re-randomized after inflation*, 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 both, structurally and numerically.

**The structural exclusion: no Peccei–Quinn scalar, no randomization.** 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 $\theta$, with cosmic strings (windings of $\Phi$) as the inevitable debris. *Every step of that story requires the radial mode.* The Kaxion has none: by Theorem 8.2 its two directions are the dual of a forced two-form and the holonomy of a rigid circle — gauge-topological angles whose "decay constant" is a geometric normalization (Appendix C.2), not the vacuum value of any order parameter. There is no field whose restoration could erase the angle at any temperature below the bulk scale; no winding order parameter to seed strings (a holonomy defect would cost bulk-scale action, $\sim E_{\text{GUT}}$); and therefore no epoch, in any history that preserves the bulk, in which separate patches could acquire separate dials. This is the defining structural property of model-independent (string-type) axions — they belong to the single-patch class automatically — and for the Kaxion it is a corollary of the existence theorem: no scalar was ever added.

*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. You can turn the hands (the misalignment of Part III); you can never melt the face. One universe, one face, one initial angle.

**The quantitative exclusion: nine to twenty-four orders of margin.** Even granting, against the structure, some restoration mechanism at $T \sim f_K$, the series' cosmology could not activate it.

The scale that sets the Kaxion angle is not the primary, GUT-adjacent inflation of GG-A10-Cosmos — whose Hubble rate $H_1 \simeq 1.6 \times 10^{13}$ GeV would, for an ordinary scanned axion, badly violate the isocurvature bound — but the *lock epoch*, the middle crossing of the two-basin descent, at $V^{1/4} \sim M_{\mathrm{lock}}$. Writing $H_2$ for its curvature scale,

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

so the de Sitter kick per e-fold at the angle-setting epoch is

$$\frac{H_2}{2\pi f_K} \simeq \frac{5.6 \times 10^{-13}}{7 \times 10^{11}} \simeq 8 \times 10^{-25} \qquad \text{(twenty-four orders below randomization)}, \tag{36.3}$$

and even instantaneous reheating bounds the maximum temperature at

$$\begin{aligned} T_{\mathrm{reh}} &\lesssim \Big(\frac{90}{\pi^2 g_*}\Big)^{1/4} \sqrt{H_2 \bar{M}_{\mathrm{Pl}}} \simeq 1.6\,\mathrm{TeV} \\ \Longrightarrow \qquad \frac{T_{\mathrm{reh}}}{f_K} &\simeq 2.3 \times 10^{-9} \quad \text{(nearly nine orders below restoration)}. \end{aligned} \tag{36.4}$$

Two margins, independent, each fatal to a randomized-angle history — and both are properties of the lock epoch, not of the high-scale primary inflation that precedes it. (The primary inflation is a spectator here: it does not set the compact angle, so its large $H_1$ never enters. This is the high-scale-inflation axion isocurvature danger turned into a feature — the Kaxion evades it precisely because its angle is fixed at the lock crossing, not scanned during primary inflation.) The single-patch history is forced; the misalignment angle enters the abundance exactly as Part V used it; and the randomized-angle history is not a second prediction but a rejected hypothesis. (One corollary for Part IV: the minicluster phenomenology of the randomized history is *off* within the theory, and is fenced there as an outside-the-series contingency.)

### 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, $f_K = f_{\mathrm{bulk}}/\Xi_{\mathrm{cap}}$; the only way the line could sit near the bulk scale is the no-renormalization limit $\Xi_{\mathrm{cap}} = 1$ ($\beta = 0$), and the series closes that loophole twice over:

*(i) The gearing cannot be zero.* The off-diagonal coupling between the two clocks is sourced by the same Chern–Simons improvement (6.13) whose coefficients are the locked vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$; the flux-cell computation of Appendix C.2 evaluates the single-transit gearing as $\varepsilon = 1/K = 1/35 \neq 0$, a ratio of OGN-5 integers. Setting $\varepsilon = 0$ is not a parameter choice; it is the deletion of BI-6 itself.

*(ii) The norms are exponentially asymmetric.* The component reduction of Appendix C.1 gives the two directions' kinetic norms as $\pi e^{2K}$ (tensor, infrared-dominated) against $\pi$ (Wilson line, ultraviolet-dominated) — a mismatch of $e^{2K} = e^{70} \simeq 2.5 \times 10^{30}$. The canonical eigenbasis of so asymmetric a system is necessarily far from the bulk basis: the cap renormalization is structurally

*large*, $\Xi_{\rm cap} \gg 1$, with only its integer value — benchmark $K+N+L = 43$ inside the discrete bracket $[35, 47]$ — awaiting the throat calculation. "$\Xi_{\rm cap} = 1$" 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.

*Physical picture.* Two pendulums coupled by a spring can be described in any basis — but a microphone in the room hears the normal-mode frequencies, never the basis frequencies. The physical line is the eigenfrequency of the canonically normalized system, which the large cap renormalization places in the gigahertz window, not at the bulk normalization scale.

### 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 genuinely unfinished, finite, geometric computation of the program (the mass-uniqueness integral), bracketed by a discrete OGN-5 window rather than by a continuum. The uniqueness is itself falsifiable twice over: exclusion of the window kills the Kaxion (F-Kaxion-1c), and so does a confirmed compact-angle-like line *anywhere else* (F-Kaxion-1d) — a wrong location is as fatal as a missing doublet. This is the feature no adjustable competitor can reproduce:

**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 unique frequency. 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 point can. The falsifiers of this section are stated against the point.

| **Model** | **Continuous dials** | **Discrete forks** | **Unique frequency?** |
|---|---|---|---|
| WIMP | $m_\chi,\ \sigma$ (decades each) | — | no: a plane, not a point |
| PQ axion | $f_a$ (10 orders), $\theta_i$ | pre/post; $E/N$ | no: forks are free choices |
| Generic ALP | $m,\ g$ unconstrained | — | no |
| **Kaxion** | **none** (±6% SM input) | two, **both closed** | **yes:** 1.59–2.13 **GHz (bench.** 1.95**)** |

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. The Kaxion's two questions had calculable answers, and the calculations are done: the haloscope program of Part VIII is not scanning a parameter space; it is reading a single pre-registered dial, at 1.95 GHz, twice split by 1/35, with every other frequency a mandated silence.

## 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 |
|---|---|---|---|
| ADMX (ADMX Collaboration, 2025) | 2.7–4.2 $\mu$eV | DFSZ | DFSZ demonstrated; below the window — the technology benchmark. |
| ADMX 1.1–1.3 GHz (ADMX Collaboration, 2025b) | 4.5–5.4 $\mu$eV | near-DFSZ | adjacent to the window's lower edge. |
| CAPP-12TB (CAPP Collaboration, 2024) | $\simeq 4.55\,\mu$eV | DFSZ | DFSZ at a second site; window-relevant technology (1–2 GHz class). |
| HAYSTAC (HAYSTAC Collaboration, 2025) | 17.0–19.5 $\mu$eV | 2.9×KSVZ | above the window in mass; squeezed-state readout pioneers sub-SQL noise. |
| QUAX, ORGAN | $\sim$ 35–45; $>$ 60 $\mu$eV | ~KSVZ | above the window; high-frequency technology pathfinders. |
| ADMX SLIC (Crisosto et al., 2020) | 0.176–0.203 $\mu$eV | $\sim 10^5\times$DFSZ | historic sub-$\mu$eV pathfinder, far below the window. |
| DMRadio-m$^3$ (DMRadio Collaboration, 2022) | $\lesssim$ 0.83 $\mu$eV (lumped) | DFSZ goal | sub-$\mu$eV program, below the window. |
| ADMX-EFR (ADMX-EFR Collaboration, 2024) | 8.3–16.5 $\mu$eV (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; $\sim$ 20–100 $\mu$eV | KSVZ–DFSZ | above the window; dielectric/plasma concepts. |
| CAST (CAST Collaboration, 2024) | solar, all masses | $g < 6 \times 10^{-11}$ | contextual ceiling, six 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

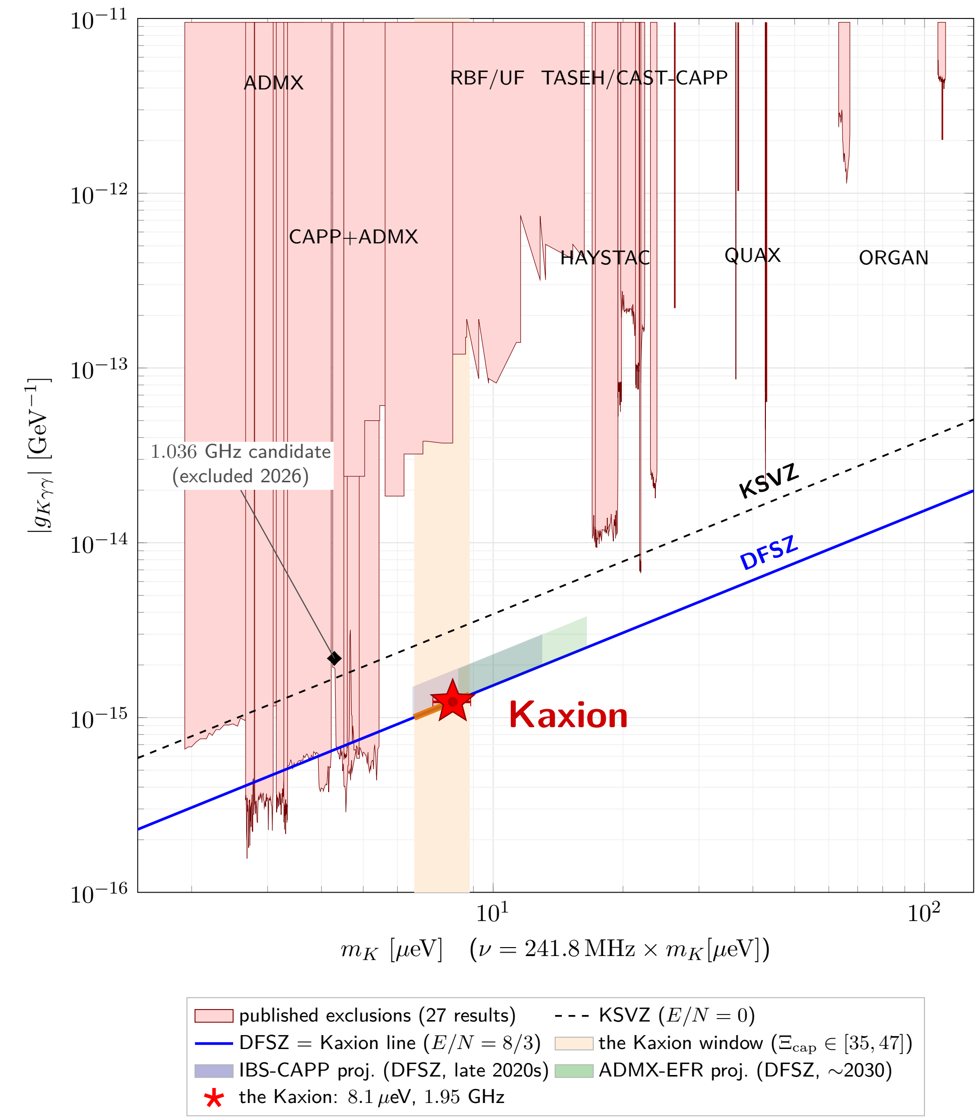


**Figure 9: The Kaxion against the haloscope landscape** in the $(m_K, |g_{K\gamma\gamma}|)$ plane. Red regions: the *published* exclusion envelopes of twenty-seven haloscope results (ADMX 2010–2025, CAPP runs including CAPP-MAX, HAYSTAC Phases I–II, TASEH, CAST-CAPP, GrAHal, QUAX, ORGAN, and the historical RBF/UF scans), digitized from the community AxionLimits compilation (O'Hare, 2020) 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 $\Xi_{\rm cap}$ bracket drawn on the line and the $\pm 6\%$ $(m)$, $\pm 8\%$ $(g)$ benchmark errors of Part V. The black star marks the 1.036 GHz candidate faced and excluded in June 2026 (§37.4) — outside the window and brighter than the DFSZ line, its confirmation would have falsified the Kaxion (F-Kaxion-1d). *Projected DFSZ-depth coverage* (on the collaborations' stated roadmaps; §36.4–36.5): IBS-CAPP approaching the window from below in the late 2020s (blue band), ADMX-EFR's 2–4 GHz campaign catching the upper edge around 2030 (green band). Published DFSZ-depth coverage ends near $5.4\,\mu$eV; inside the window the only published limits (RBF/UF) sit an order of magnitude *above* the Kaxion line: 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 runs two solenoids: one tuning 0.5–1.3 GHz, and a second spanning 1.8–10 **GHz** — a magnet that already brackets the upper $\sim 60\%$ of the Kaxion window. 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 June 2026 the group completed a textbook candidate-validation campaign — following up, and excluding, a $5.1\sigma$-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.** The 1.8–10 GHz magnet covers 1.8–2.13 GHz directly; the lower part of the window, 1.59–1.8 GHz, is the seam between CAPP's two magnets and is reachable with a window-tuned cavity. CAPP has therefore already demonstrated *both* the magnet that spans the window *and* the DFSZ depth that the Kaxion — sitting on the DFSZ line — requires; it is the only group of which both statements are presently true.

**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 1/35 doublet and the 5:2 portal — it needs the two add-ons of Part VIII: a second cavity mode at the 56 MHz doublet spacing (a modest retrofit), and a co-located spin port (the harmonic booster), for which CAPP already has the requisite cryogenics and quantum-limited readout.

**Concrete proposal (increasing cost).** (i) A *doublet-aware re-binning* of any CAPP scan that crosses $[1.59, 2.13]$ GHz — search for the 2.857% partner of every candidate, at *zero* hardware cost. (ii) Prioritize the window center 1.95 GHz on the 1.8–10 GHz magnet, the highest-value 540 MHz in the Kaxion program. (iii) Retrofit a two-mode (doublet) cavity at the center. (iv) Add the 70 mT electron-magnon port alongside, then the internal-field nucleon port (Concept B′, §40.2) as the baryonic R&D track. Steps (i)–(ii) are additive to CAPP's funded program at negligible marginal cost; (iii)–(iv) are 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 sweep of the 540 MHz window is a $\sim$ 1–2-year campaign; on CAPP's stated 1–8 GHz roadmap the window is reached around **2027–2030**. 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 (doublet + 5:2) follows with the retrofits, late-2020s to early-2030s.

## 36.5 ADMX-EFR: the funded 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 (a magnetic-field fault was under repair in spring 2026), and is developing a four-cavity coherently-combined system to move upward in frequency from its published 1.1–1.3 GHz coverage (ADMX Collaboration, 2025b) — toward the Kaxion window from below. 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** (8.3–16.5 $\mu$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 (8.3–8.8 $\mu$eV); the benchmark 1.95 GHz/8.1 $\mu$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 funded *backstop* that catches 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 $\sim$ 50–100 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 the doublet mode and a spin port; 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 *doublet-aware analysis*: the 56 MHz partner of a 2.0 GHz line falls inside the EFR band, so the 1/35 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. (iv) Add a spin port for the 5:2 test. Steps (i)–(ii) cost only analysis and run-plan priority.

**Timescale to $> 5\sigma$.** With operations beginning in the late 2020s, the 2.0–2.13 GHz overlap is covered at DFSZ in the first $\sim$ 1–2 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 2–3×. Its 2025 Phase-II results scanned 16.96–19.46 $\mu$eV ($\sim$ 4.1–4.7 GHz) to $\sim$ 2.75–2.96×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 $\mu$eV, below HAYSTAC's band — so HAYSTAC's null at 17–19.5 $\mu$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 booster 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* $[1.59, 2.13]$ *GHz discovers it or kills the window.* The plan follows the instruments' own roadmaps and asks of them, in increasing cost:

1. **Doublet-aware analysis (zero hardware).** Re-bin every existing and planned scan crossing the window for the 2.857% partner. This converts ordinary scan data into a Kaxion identification test at no instrumental cost, at both CAPP and ADMX-EFR.

2. **Window priority.** Direct CAPP's 1.8–10 GHz magnet at the window center (1.95 GHz) and ADMX-EFR's early operation at its band bottom (2.0–2.13 GHz). Between them they cover the whole window this decade.

3. **Identification retrofit.** A two-mode (doublet) cavity and a co-located spin port (the harmonic booster of Part VIII) — the upgrade that turns a discovery into an identification of gauged baryon number.

The projected DFSZ-depth coverage is the CAPP and ADMX-EFR bands of Figure 9: on the collaborations' stated (and somewhat optimistically extrapolated) roadmaps, both cross the Kaxion line inside the window between $\sim$ 2030 and $\sim$ 2035. The primary approach is IBS-CAPP (the magnet and the depth already exist at the window); the funded US backstop is ADMX-EFR (the upper edge); HAYSTAC is a readout-technology partner. In every case the dedicated contribution — the doublet-aware analysis and the harmonic-booster retrofit — is *additive* to a funded program rather than a competing build.

## 36.8 The bottom line

Three statements summarize Figure 9, each with its receipt. *First*, the window is not excluded: published at-or-near-DFSZ coverage ends near 5.4 $\mu$eV (ADMX's 1.1–1.3 GHz run; CAPP-MAX), below the window, and inside the window the only published limits — the historical RBF/UF scans — sit an order of magnitude above the Kaxion line: 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 two-mode retrofit turns any cavity in the band into a doublet instrument, and the 70 mT electron-spin port (Part VIII) adds the channel-ratio test without exotic magnets. The strategic conclusion for a dedicated program (Part VIII): strike the window with a doublet-aware instrument riding the 2030 DFSZ program, with the spin port alongside.

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

Kill criteria without dates are rhetoric. Falsifier F-Kaxion-1c — a DFSZ-depth, doublet-aware exclusion of the full window $[1.59, 2.13]$ GHz — becomes a calendar entry once each instrument's route into the window is laid beside its stated roadmap:

**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.

| Instrument | Route into $[1.59, 2.13]$ GHz | DFSZ depth | Window entry |
|---|---|---|---|
| CAPP-MAX (JPA upgrades) | up from 1.5 GHz | demonstrated | ∼2027–2029 |
| ADMX four-cavity array | up from 1.3 GHz | goal (2026 run) | ∼2028–2030 |
| ADMX-EFR | 2–4 GHz (upper edge) | design | ∼2029–2032 |
| 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, ADMX from 1.3 GHz), the EFR campaign closes the upper edge around 2030–2032, and a full DFSZ-depth, doublet-aware sweep of the window is therefore plausibly *complete by ∼2030–2032* on funded and stated programs — with the camping option standing open at every moment in between. That is the kill date this paper accepts in advance: if, by the early 2030s, the window has been swept at DFSZ depth with a doublet-aware analysis and no 1/35 pair has appeared, F-Kaxion-1c has fired and the Kaxion — and with it the dark-matter claim of the series — is retired. Part VIII closes by drawing this same calendar as the cohort's kill-test timeline (Figure 11), 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 one

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 $\nu$-dependent linewidth $\nu/Q_a$ gives $N \approx 2$–$5 \times 10^6$; the requirement below is 7.3–7.4$\sigma$ 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 5$\sigma$ excess in such a scan carries a global false-alarm probability of order unity, and discovery-grade *global* 5$\sigma$ requires local 7.4$\sigma$ — a factor 2.2 more integration time everywhere, paid on every bin of the scan. Pre-registering the Kaxion window alone ($2.8 \times 10^5$ bins) still demands local 7.0$\sigma$. Pre-registering the *frequency* — one bin, named in this paper, in advance of the data — collapses the trials factor to one: local significance *is* global significance, and 5$\sigma$ means 5$\sigma$.

The field has just supplied the textbook illustration. In June 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 5.1$\sigma$, 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 3.5$\sigma$ 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 frequency down first.

### 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 $g_{K\gamma\gamma} = 1.2 \times 10^{-15}$ GeV$^{-1}$, for three instrument classes. The volume entries are honest: a single CAPP-MAX-class cavity retuned from its published $\sim 1.1$ GHz operating point to 1.95 GHz shrinks as $V \propto \nu^{-3}$, from 37 to $\sim 6$ liters — the high-frequency volume penalty that multi-cavity arrays exist to repair, and repair well: the eighteen-cavity EFR-class array restores $\sim 200$ liters at the window.

**Table 17: Camping on the benchmark.** The registered Sikivie power (35.1) and Dicke time to $5\sigma$ over one condensate bin ($\Delta\nu_a \simeq 2.0$ kHz), at 1.95 GHz, DFSZ coupling, $\rho_{\rm loc} = 0.45$ GeV/cm$^3$, $C$ and $Q_L = 10^5$ as listed, critical antenna coupling absorbed in $Q_L$ (App. D.1; signal photon rates $3 \times 10^1$–$4 \times 10^2$ per second across the rows). Idealized live time; the 10–100× 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 $5\sigma$ within a day.

| **Instrument class** | $B_0$ | $V$ | $C$ | $T_{\rm sys}$ | $P_{\rm sig}$ | $t_{5\sigma}$/bin |
|---|---|---|---|---|---|---|
| Representative design cavity (§35) | 8 T | 100 L | 0.5 | 94 mK (SQL) | $2.4 \times 10^{-22}$ W | 1.5 s |
| CAPP-MAX-class single cavity at 1.95 GHz | 12 T | 6.4 L | 0.6 | 200 mK | $4.1 \times 10^{-23}$ W | 3.7 min |
| EFR-class 18-cavity array | 9.4 T | 200 L | 0.4 | 150 mK | $5.3 \times 10^{-22}$ W | 0.76 s |
| CAPP-MAX-class cavity, 4 K HEMT receiver | 12 T | 6.4 L | 0.6 | 4 K | $4.1 \times 10^{-23}$ W | 24 h |

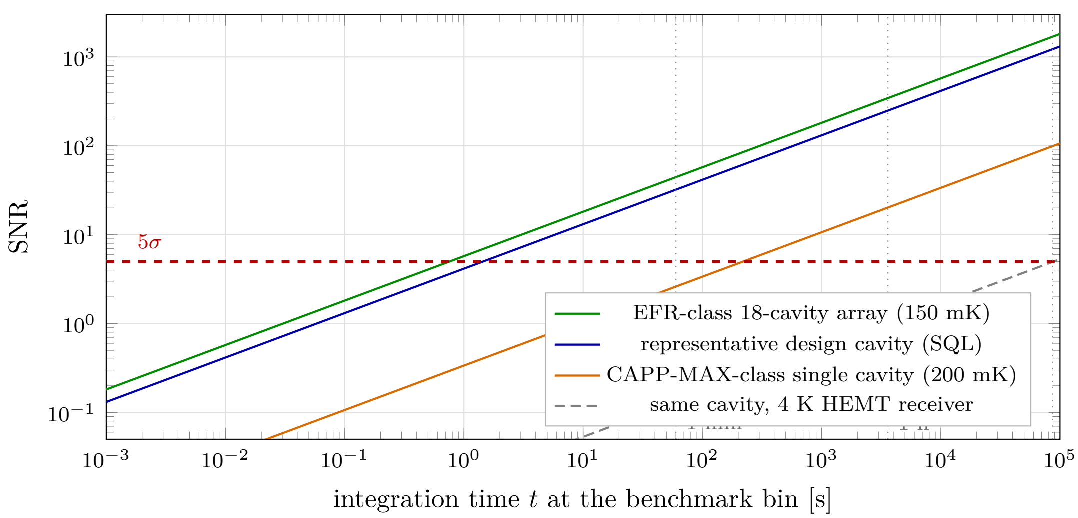


**Figure 10: SNR versus integration time, camped on the benchmark.** Dicke radiometer growth SNR $\propto \sqrt{t}$ at 1.95 GHz and DFSZ coupling for the instrument classes of Table 17 (idealized live time). Every curve crosses discovery-grade $5\sigma$ within seconds to minutes — even the deliberately 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 $5\sigma$ 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 Standard-Model-side error, not the geometry: the ±6% on the benchmark (Part V) spreads the target over a $\sim 240$ MHz band, which a CAPP-MAX-class cavity covers at DFSZ depth, at one $5\sigma$-local dwell per cavity bandwidth, in $\sim 1$ idealized month —

against $\sim 70$ idealized days for the full 540 MHz window, which the scanning overheads of 10–100 turn into the 1–2-calendar-year campaign already quoted in §36.4; an EFR-class array covers the entire window in under six idealized hours, days to weeks of calendar time. The path from $\pm 6\%$ to $\pm 3\%$ (Part V) halves the band, and every lattice-QCD improvement to $\chi_{\mathrm{QCD}}$ narrows the camp site further. (These band-coverage times are $5\sigma$-*local* dwell; claiming a discovery anywhere in the band at *global* $5\sigma$ costs a factor $\times 2$ in dwell, §37.1.) *Third*, the doublet is free at these depths: the 56 MHz partner of a benchmark line sits well inside every listed instrument's tuning range, so the identification test of Part VI rides along with no additional exposure.

### 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 1/35 doublet template, and the 25/4 channel ratio are fixed in advance by this paper and its falsifier registry (Appendix E.2); the trials factor of §37.1 is thereby one for the benchmark and $\leq 2.8 \times 10^5$ for the window — both declared before the data.

2. **Benchmark first.** Camp on 1.95 GHz and its $\pm 6\%$ band — the highest-value 240 MHz in the field — then step the discrete $\Xi_{\mathrm{cap}}$ ladder of Part V outward from the center. The ladder is finite; the campaign is enumerable.

3. **Doublet-aware from day one.** Analyze every spectrum against the two-tone template. A candidate anywhere in the band is confirmed or killed within days: the partner line at $\Delta\nu/\nu = 1/35$ and persistence tests (field-off, cross-instrument, re-scan) settle it, precisely the validation discipline of Ahn et al. (2026).

4. **The second channel on standby.** The 70 mT electron-magnon port (Part VIII) is a sub-tesla add-on; commissioning it alongside the camp turns a confirmed doublet into a measured channel ratio — identification, not detection.

**The Kaxion Search Card — everything a haloscope collaboration needs, on one page**

**Target.** $\nu_K = 1.95 \pm 0.12$ GHz ($m_K = 8.1 \pm 0.5\,\mu$eV); closed window $\nu_K \in [1.59, 2.13]$ GHz; discrete ladder $\Xi_{\rm cap} \in \{35, 38, 40, 42, \mathbf{43}, 45, 47\}$ (benchmark bold; per-rung frequencies in Table 33). A confirmed compact-angle line at any frequency *outside* the window falsifies the theory (F-Kaxion-1d).

**Strength.** Exactly on the DFSZ line: $g_{K\gamma\gamma} = (\alpha/2\pi f_K)\,[8/3 - 1.92] = (1.2 \pm 0.1) \times 10^{-15}$ GeV$^{-1}$; expected cavity power $P_{\rm sig} = g^2_{K\gamma\gamma}(\rho_{\rm loc}/m_K)B_0^2 VC \min(Q_L, Q_a)$ — worked budgets at three instrument classes in Table 17.

**Line shape.** Standard-halo form $\sqrt{u}\,e^{-u}$: sharp low edge riding $\simeq 0.6$ kHz above $\nu_K$ (the solar-motion offset), thermal tail, width $\simeq 2.0$ kHz ($Q_a \simeq 10^6$); breathing $\pm 74$ Hz annually and $\sim\pm 2$ Hz with the *sidereal* day (Eqs. (48.1)–(48.2)).

**The two locks.** (i) Every line is a *pair*: partner at $\Delta\nu/\nu = 1/35 = 2.857\%$ (55.7 MHz at the benchmark), tone ratio unconstrained, partner threshold $10^{-2}$ relative power — a single-mode receiver suppresses it by 75 dB, so it must be asked for. (ii) Matched two-channel ratio $P_B/P_\gamma = 25/4 = 6.25$ (electron-magnon port at $B_e = 70$ mT; staging $\pm 20\% \to \pm 5\%$, App. D.2).

**Monday morning, in increasing cost.**

- re-bin any archived scan crossing the window for the 2.857% partner (zero hardware);
- camp the benchmark: at DFSZ depth the matched line reaches $5\sigma$ in seconds-to-minutes of live time (Table 17); the $\pm 6\%$ band is hours (array-class) to a month (single cavity) of idealized dwell;
- retune to $\nu(1 \pm 1/35)$ on any candidate — F-Kaxion-1a settles it within a day;
- add the 70 mT spin port for the 25/4 verdict.

**The four falsifiers** (App. E.2): ☐ F-1a: partner absent at 95% CL at $10^{-2}$ relative power ⇒ dead. ☐ F-1b: converged ratio $\neq 25/4$ at $> 5\sigma$ (incl. calibration syst.) ⇒ dead. ☐ F-1c: window swept at DFSZ, doublet-aware, silent ⇒ dead. ☐ F-1d: confirmed line outside the window ⇒ dead.

**Reproduce everything.** The discovery simulation (spectra, modulations, doublet response, ratio convergence; fixed seeds) is a single public standard-library Python script (§53); 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 Kaxion outright, wherever it appears and however faint. 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× 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 funded instruments — IBS-CAPP and ADMX-EFR — are being 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 doublet-aware sweep of $[1.59, 2.13]$ GHz completed at DFSZ depth with no 1/35 pair 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* tone in a *single* channel. None can see the 1/35 partner and the 25/4 ratio in one run — the very fingerprints Part VI made the identification test. *Confirming* the Kaxion, as opposed to merely detecting an axion-like line, therefore calls for apparatus built to the prediction's own shape: two cavities locked a doublet-spacing apart, and a second port that hears the baryonic channel. Part VIII designs those dedicated instruments, and shows that neither asks for 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

Dedicated Kaxion instruments for the two parameter-free locks

## 39 The twin-peak coincidence resonator

**Principle.** A single-mode haloscope answers "is there power at $\nu$?"; identification needs "is there power at $\nu$ *and* $\nu(1-1/35)$, coherently, always?" Concept A is a receiver with two simultaneously instrumented modes locked at the geometric spacing: a dual-mode cavity (or two cavities phase-locked to one synthesizer) with resonances at $\nu_\pm = \nu_0(1 \pm 1/70)$, scanned together so that the pair spacing is $\Delta\nu/\nu = 1/35$ at every step. The Kaxion drives both modes (Eq. (35.3)); thermal and technical artifacts do not conspire across both at fixed fractional spacing. Coincidence is the discriminator: demanding simultaneous excesses suppresses the accidental background quadratically, so each mode may run faster and shallower than a discovery-grade single-mode scan — the pair is more sensitive *and* more diagnostic than the sum of its parts.

**Design quantities.** Center 1.95 GHz, mode spacing $56 \pm 4$ MHz, condensate linewidth 2.0 kHz, required relative-spacing accuracy $10^{-3}$ of the splitting (trivial for modern synthesis); single-photon-counting readout (transmon-qubit class) preferred over linear amplification — at 1.95 GHz and $\sim$10 mK, dark counts rather than the standard quantum limit set the floor, buying the factor that DFSZ-depth scanning otherwise pays in years.

## 40 The 5:2 harmonic booster

**Principle.** Concept B reads the second lock: operate the electromagnetic port and a *spin* port — nuclear or electron — in the same condensate, and form the deconvolved power ratio of Theorem 32.1, target 25/4, with every instrumental factor independently calibrated and divided out. The spin port couples to the baryonic/nucleon channel (NMR-class, CASPEr-style (CASPEr Collaboration, 2021)) or the electron channel (ferromagnetic-magnon class); resonance requires matching the Larmor frequency to $\nu_K$.

**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,

$$\nu_K = 1.95 \text{ GHz}: \qquad B_e = (70 \pm 7)\,\text{mT}, \qquad B_p = 45.8\,\text{T (impractical)}, \tag{40.1}$$

so the practical spin port is the *electron* channel: a $(70 \pm 7)$ 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; the full baryonic 25/4 lock therefore proceeds in two stages: the matched electron-port ratio first (its calibration sheet is the model-level chiral input fenced in Part V), with non-Larmor nucleon-readout concepts registered as the design study's R&D track. A confirmed 25/4 in the matched pair is direct laboratory evidence for gauged baryon number (Section 32).

## 40.1 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 $\nu_K$ 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-$\mathrm{U}(1)_B$ portal reduces onto ordinary matter.

*The axial (spin) reduction.* Chirally reduced, the $a\,F_B\widetilde{F}_B$ 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 $\nu_K$ 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.* $\mathrm{U}(1)_B$ 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 $\nu_K$, read mechanically with no magnet and no spins (Concept B‴). This route reads the very quantity $\mathrm{U}(1)_B$ gauges, and is the conceptual complement of the spin ports.

We develop all three. Each is a candidate baryonic port for the matched-ratio test of Theorem 32.1; each is fenced by the same model-level input — the chiral reduction that converts the theorem-level gauge ratio 5:2 into an instrument response (open item (iii), Section 75).

## 40.2 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* $B_{\mathrm{hf}}$ 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,

$$\nu_n = \frac{\gamma_N}{2\pi}\,(B_{\mathrm{hf}} + B_0), \tag{40.2}$$

and, decisively, the electronic system *enhances* both the rf drive on the nucleus and its radiated signal by the hyperfine enhancement factor $\eta \sim B_{\mathrm{hf}}/B_{\mathrm{a}} \sim 10^2$–$10^4$ (with $B_{\mathrm{a}}$ the anisotropy field) (Turov & Petrov, 1972); the electron-spin analogue of an enhanced magnon port is the quantum-limited QUAX ferromagnetic haloscope (Crescini et al., 2020). A survey of ordered nuclei brackets the window:

$$^{55}\text{Mn}\,(\text{MnF}_2): \;\; \nu_n \simeq 0.66\,\text{GHz}; \qquad ^{159}\text{Tb}: \sim 3\,\text{GHz}; \qquad ^{165}\text{Ho}\,(B_{\text{hf}} \sim 7\text{–}8{\times}10^2\,\text{T}): \sim 6\text{–}7\,\text{GHz},$$

so material engineering plus a sub-tesla trim field $B_0$ places an enhanced transition on $\nu_K = 1.95$ 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 $h\nu_n/k_B = 93.6$ mK at 1.95 GHz; in a dilution refrigerator at $T = 10$ mK the equilibrium polarization is

$$p = \tanh\frac{h\nu_n}{2k_BT} = \tanh(4.68) = 0.9998, \tag{40.3}$$

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 $1\,\text{cm}^3$ 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 $\mathcal{F}_B$ in Theorem 32.1 is calibrated against the *same* condensate the photon port sees — exactly the matched configuration the power ratio 25/4 requires. The one cost is the model-level chiral input that converts $g_{KBB}$ into the enhanced nucleon response; this is the calibration sheet fenced in Part V, and is the single theory deliverable the port waits on.

## 40.3 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 $\omega_p$ that dresses the spins; the Kaxion at $\omega_K$ 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 analogue lets a low-field nuclear sample respond to $\nu_K$ 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.

## 40.4 Concept B‴: the baryon-number force readout

The conceptually cleanest port, reading the vector charge directly. The Kaxion's gauged-$\text{U}(1)_B$ 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 $\nu_K$, 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 $\mathrm{U}(1)_B$ 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 18 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 18:** 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 of the four-current lock (GG-A7-Qi4 T-QUA) the port accesses; the *absolute* normalization of every nucleon row carries the model-level chiral fence of Part V (open item (iii)), while the 5:2 *ratio* it tests is theorem-level.

| 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 ($g_{ae}^{\mathrm{eff}}$) | mature, QUAX-class (Crescini et al., 2020); reads the $J_e$ lock, not $J_B$ |
| **B′** internal-field nuclear | $\lesssim 1$ T trim | hyperfine $B_{\mathrm{hf}} \sim 10$–$10^2$ 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 ($\mathrm{U}(1)_B$ 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 baryonic magnitude: the 5:2 *ratio* is theorem-level, but the integration time any nucleon port needs is set by the chiral reduction of the gauged-$\mathrm{U}(1)_B$ portal (open item (iii)), a theory deliverable that should precede metal. The virtue is that the harmonic-booster observable is a *ratio* in a shared condensate: $\rho_{\mathrm{loc}}$, 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 recommended sequence is therefore unchanged: the electron port demonstrates the matched-ratio method now; B′ is the baryonic-port R&D target once its chiral calibration is in hand; B″ and B‴ are the pathfinder and the magnet-free complement.

### 40.5 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

$$g_{aNN} = c_N \frac{m_N}{f_K} = c_N \times 1.3 \times 10^{-12}\,\mathrm{GeV}^{-1}, \qquad c_N = \mathcal{O}(1)\ \text{(uncomputed; open item (iii))}, \tag{40.4}$$

with $m_N = 0.94$ 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; only the dimensionless coefficient $c_N$ — the chiral matrix element of the gauged-U$(1)_B$ current in the nucleon — is uncomputed. Its plausible range is set by analogy with the QCD axion, whose nucleon couplings are $\mathcal{O}(10^{-2}$–$10^{-1})$ (Grilli di Cortona et al., 2016); we carry $|c_N| \in [0.05, 1]$ as the honest prior.

A resonant spin haloscope accumulates significance as a radiometer, SNR $\propto g_{aNN}^2\sqrt{t_{\rm int}}$, so the integration time to a fixed significance scales as

$$t_{\rm int} \;\propto\; g_{aNN}^{-4} \;\propto\; c_N^{-4}. \tag{40.5}$$

At the Concept-B′ benchmark — $N \simeq 2 \times 10^{22}$ polarized nuclei ($p \simeq 0.9998$, Eq. 40.3), hyperfine enhancement $\eta \sim 10^3$, and a quantum-limited (ideally photon-counting) readout co-located with an *already detected* photon line — the matched-ratio test of Theorem 32.1 is reached within a campaign comparable to the photon discovery run ($\mathcal{O}$(months–year)) for $c_N \sim 1$, and lengthens sharply as $c_N$ falls: by (40.5), $c_N \simeq 0.3$ costs a factor $\sim 10^2$ in time and $c_N \simeq 0.1$ a factor $\sim 10^4$ — the difference between a one-year confirmation and "not with this method." Below $c_N \sim 0.05$ the co-located ratio test is impractical with present spin readouts, and identification would lean on the electron port or a dedicated sensitivity upgrade. This is the booster's feasibility band, and it is one chiral computation wide.

Two facts keep the band from being a fatal uncertainty. First, what the booster delivers is the *deconvolved ratio* $P_B/P_\gamma = 25/4$, in which $\rho_{\rm loc}$, the halo velocity distribution, and the coherence time cancel: the absolute $c_N$ sets only *how long* the test takes, never *whether* the ratio is 25/4 when it is measured. Second, the test rides on the photon detection — the booster is a retrofit to an instrument that has already found the line (Part VIII), so its integration time is spent confirming the second channel of a known signal, not searching blind. The chiral computation that fixes $c_N$ (open item (iii)) therefore converts this band into a schedule, and is the recommended theory precursor to any baryonic-port construction.

# 41 Search strategy and roadmap

Table 19 condenses the strategy of Parts VII–VIII into three stages — what runs now at zero marginal cost (the doublet-aware 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 19:** 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–) | Doublet-aware reanalysis of existing 1–2 GHz data; coincidence pathfinder (two-mode retrofit) | the window at KSVZ depth; doublet logic proven | method validation; first 1/35-aware limits |
| II (2030 era) | Concept-A dual-mode in the window, photon counting; ride the announced DFSZ program | full $[1.59, 2.13]$ GHz at DFSZ depth | *discovery, or closure of the window* |
| III (2030s) | Electron-port 25/4 booster at the window (70 mT magnon + matched cavity); non-Larmor nucleon-port R&D | matched two-channel ratio in the window | identification: doublet + portal lock 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 frequency (±6%, improvable to ±3% by lattice work alone, Section 29.2), its coupling (DFSZ-exact), its doublet (1/35, exact), and its channel ratio (25/4, exact) all written down in advance. A Stage-III confirmation of the 25/4 would simultaneously identify the dark matter, validate the gauged-baryon sector of GG-6, and measure the throat depth of a six-dimensional geometry with a sub-tesla spin port. Conversely, the full program executed null retires the Kaxion under F-Kaxion-1 — and with it the dark-matter claim of the series, exactly as a falsifiable framework should die.

## 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 11 draws it in the cohort's kill-test-timeline convention: the test already faced and survived (June 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 1/35 pair 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 twice: a dual-mode receiver for the 1/35 pair, and a spin port for the 25/4 ratio, both riding hardware the field already knows how to build, both 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.

**Two-sided and dated.** One pre-registered window; a confirmed line anywhere else is fatal (F-Kaxion-1d — already faced once, June 2026), and a completed doublet-aware DFSZ sweep of the window with no 1/35 pair is fatal (F-Kaxion-1c, ∼2030–2032). Between the standing camping option and the funded scans, the Kaxion is found or dead within the decade.

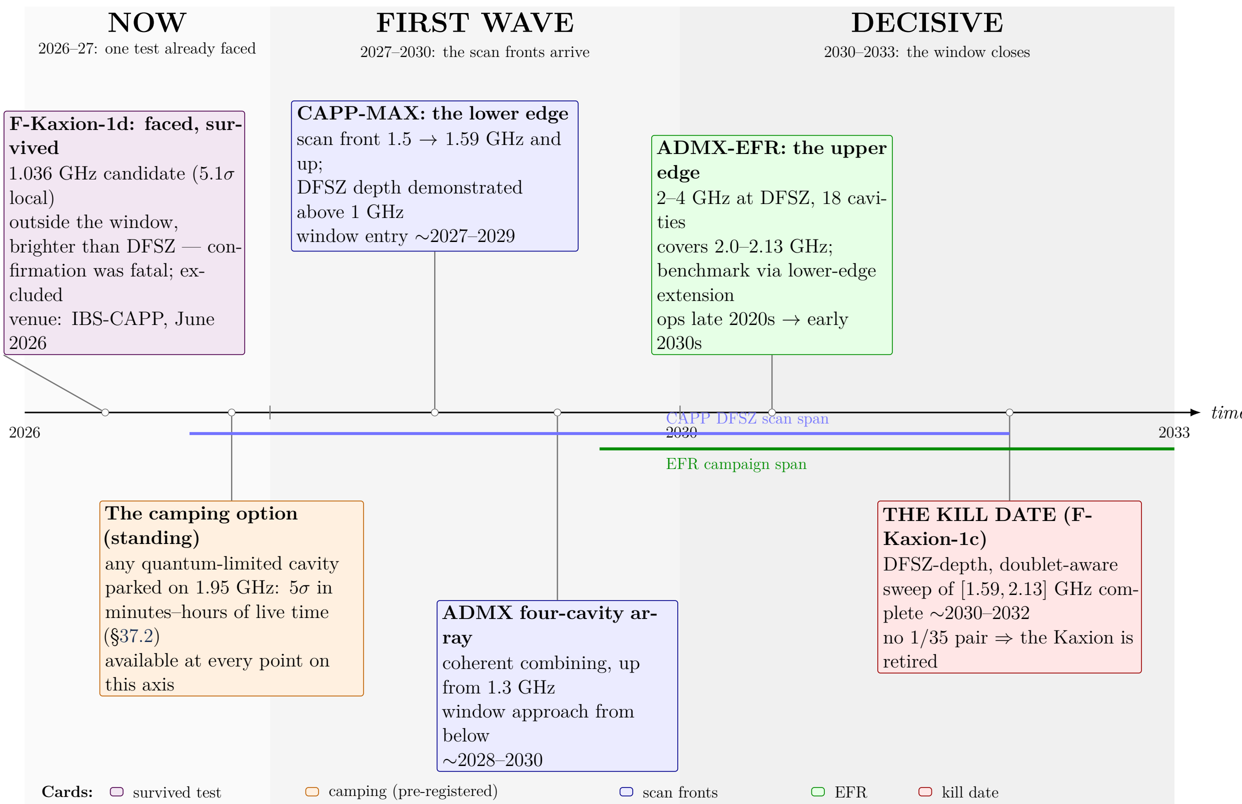


**Figure 11: Falsification / kill-test timeline for the Kaxion.** The discovery calendar of Table 16 drawn on a forward time axis (2026–2033, broad data-window eras rather than precise launch dates), in the cohort's kill-test-timeline convention. Violet: the F-Kaxion-1d test already faced and survived (the 1.036 GHz candidate, June 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, doublet-aware sweep of $[1.59, 2.13]$ GHz with no 1/35 pair retires the Kaxion, and with it the dark-matter claim of the series, by ∼2030–2032 on stated roadmaps.

# 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 second lock. A machine that scans looks for one line and moves on. A machine built for this prediction has to see two lines at a fixed separation, which means it needs two modes live at once rather than one mode swept past both. That single requirement drives most of the design: the mode structure, the receiver chain, the coincidence 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 two-lock geometry in mind, and the argument of this Part is that building one is a modest step rather than a new program.

# Part IX
# Simulating the Discovery

The data of discovery week, computed in advance: inputs, equations, error budgets, and the four 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 free parameters, and that changes what this Part can do: every property of the signal — its frequency, its strength, its spectral shape, its seasonal breathing, its partner tone, its second channel — 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 (§51). The generator is a single deterministic script, released with the paper (§53), 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; the response of a single tuned mode to the 1/35 partner; and the convergence of the two-channel power ratio. *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 27); 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, doublet spacing, channel ratio — it cannot mimic.

## 44 The simulation pipeline and its inputs

Figure 12 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. The three blocks are deliberately separated because their epistemic status differs: the *theory sheet*

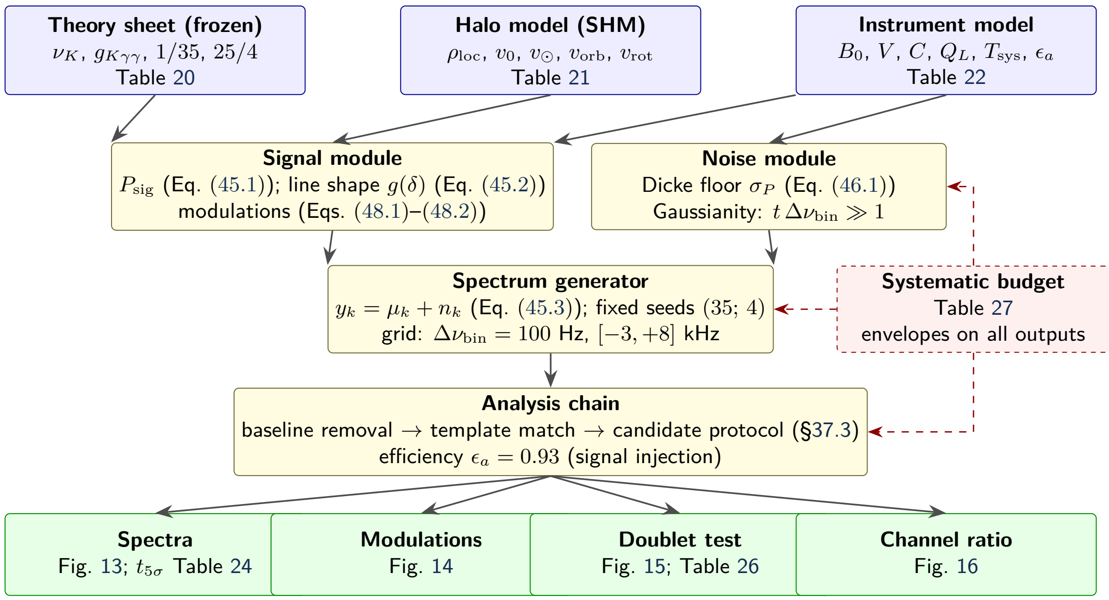


**Figure 12: The simulation pipeline.** Three input blocks (theory sheet, frozen by Parts II–V; halo model; instrument model — Tables 20–22) feed the signal and noise modules, whose equations are stated in §45–46; the spectrum generator draws fixed-seed Gaussian noise on a 100 Hz grid; the analysis chain applies baseline removal, template matching, and the candidate protocol of §37.3 with the signal-injection efficiency $\epsilon_a$. The systematic budget (Table 27, red dashed) propagates onto every output as the envelopes drawn in Figures 13–16. The generator is a single self-contained script (§53).

**Table 20: 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 $\nu_K$ | 1.95 GHz | ±6% (SM: $\chi_{\rm QCD}$, window) | Part V, Eq. (26.1) |
| mass $m_K$ | $8.1\,\mu$eV | ±6% | Part V |
| photon coupling $g_{K\gamma\gamma}$ | $1.2 \times 10^{-15}$ GeV$^{-1}$ | ±8% | Part V (DFSZ, $E/N = 8/3$) |
| doublet spacing $\Delta\nu/\nu$ | 1/35 | exact (corr. $\frac{5}{8}K^{-2} \simeq 5 \times 10^{-4}$) | Part VI |
| channel ratio $P_B/P_\gamma$ | 25/4 | exact | Theorem 32.1 |
| window | $[1.59, 2.13]$ GHz | discrete $\Xi_{\rm cap}$ ladder | Part V |

(Table 20) is frozen by Parts II–V and carries only the Standard-Model-side errors of Part V; the *halo model* (Table 21) 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 22) 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 27.

**Table 21: 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 cancel entirely in the channel-ratio observable.

| Quantity | Value | Assigned syst. | Justification |
|---|---|---|---|
| local density $\rho_{\rm loc}$ | 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 $v_0$ | 270 km/s | $\pm 8\%$ | Maxwellian $f(v) \propto v^2 e^{-v^2/v_0^2}$ in the rms-speed convention ($v_0 = \sqrt{3/2}\, v_c$; the collaborations quote circular speeds $v_c \simeq 220$–226 km/s); sets line width via Eq. (45.2) |
| solar speed $v_\odot$ | 232.8 km/s | $\pm 3$ km/s | galactocentric solar motion |
| orbital speed $v_{\rm orb}$ | 29.79 km/s | — | ephemeris (exact) |
| orbit projection $\cos\gamma$ | 0.49 | $\pm 5\%$ | angle between orbital plane and solar apex |
| rotation speed $v_{\rm rot}\cos\lambda$ | 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,

$$P_{\rm sig} = g_{K\gamma\gamma}^2 \frac{\rho_{\rm loc}}{m_K} B_0^2\, V\, C\, \min(Q_L, Q_a), \tag{45.1}$$

identical to Eq. (35.1); nothing is re-derived here. What this Part adds is the worked budget: Table 23 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, $f(v)\, dv \propto v^2 e^{-v^2/v_0^2} dv$. 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 the lab-frame line shape

$$g(\delta)\, d\delta = \frac{2}{\sqrt{\pi}} \frac{\sqrt{\delta}}{\Delta^{3/2}}\, e^{-\delta/\Delta}\, d\delta, \qquad \Delta \equiv \frac{\nu_K v_0^2}{2c^2} = 791 \text{ Hz}, \tag{45.2}$$

normalized to unit area for $\delta > 0$: a sharp low-frequency edge at $\nu_K$ (the condensate at rest), a maximum at $\delta = \Delta/2$, and a thermal tail upward. Its full width at half maximum is $1.79\,\Delta = 1420$ Hz, and the registered working bandwidth $\Delta\nu_a = 1.95$ kHz of Part VI contains $\approx 86\%$ of the line power — the two statements describe the same shape. (The simulation carries the solar motion as a rigid $\simeq 0.6$ kHz offset of this isotropic form; the exact lab-frame shape is skewed at the $\sim 10\%$ level, an approximation absorbed into the $v_0$ systematic of Table 21 that moves no conclusion of this Part.)

**Table 22: 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 $\sim 1.1$ 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 | ±1% (±2% on $B_0^2$) | NMR probe map of the solenoid bore |
| cavity volume $V$ | 6.4 L | ±2% | machining + cryogenic contraction; $37\,\mathrm{L} \times (1.09/1.95)^3$ |
| form factor $C_{010}$ | 0.60 | ±5% | bead-pull measurement cross-checked against FEM eigenmode computation |
| loaded quality factor $Q_L$ | $10^5$ | ±5% | vector-network-analyzer $S$-parameters at each tuning |
| system noise $T_{\mathrm{sys}}$ | 200 mK | ±6.4% | Y-factor + shot-noise calibration of the JPA chain; dominant instrument systematic |
| analysis efficiency $\epsilon_a$ | 0.93 | ±1% | software signal injection through the full analysis chain |
| spectral bin $\Delta\nu_{\mathrm{bin}}$ | 100 Hz | — | analysis choice; $\sim 1/14$ of the line width |
| condensate bandwidth $\Delta\nu_a$ | 1.95 kHz | from $v_0$ | registered Part-VI value ($\approx$ 86% of line power; FWHM 1420 Hz) |
| B-port sensitivity $\mathrm{SNR}_B$ | $5\sqrt{t/5\,\mathrm{d}}$ | design req. | Concept-B electron-magnon port requirement (Part VIII); fenced |

**Table 23: The worked power budget.** Equation (45.1) evaluated term by term for the instrument of Table 22. 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 |
|---|---|---|---|
| $g^2_{K\gamma\gamma}$ | $(1.2 \times 10^{-15}\,\mathrm{GeV}^{-1})^2$ | ±17% | theory (SM-side) |
| $\rho_{\mathrm{loc}}/m_K$ | $0.45\,\mathrm{GeV/cm^3}\,/\,8.1\,\mu\mathrm{eV}$ | ±22% | halo |
| $B_0^2$ | 144 $\mathrm{T}^2$ | ±2% | instrument |
| $V$ | 6.4 L | ±2% | instrument |
| $C$ | 0.60 | ±5% | instrument |
| $\min(Q_L, Q_a)$ | $10^5$ | ±5% | instrument |
| $P_{\mathrm{sig}}$ | $4.1 \times 10^{-23}$ W | ±23% (halo+instr.) | $\Rightarrow 3.2 \times 10^1$ photons/s |

**Per-bin expectation.** On the analysis grid of bin width $\Delta\nu_{\mathrm{bin}}$, the expected signal-to-noise in bin $k$ after live time $t$ is

$$\mu_k = \epsilon_a \frac{P_{\mathrm{sig}}\, w_k}{k_B T_{\mathrm{sys}}} \sqrt{\frac{t}{\Delta\nu_{\mathrm{bin}}}}, \qquad w_k = \int_{\mathrm{bin}\ k} g(\delta)\, d\delta, \tag{45.3}$$

**Table 24: Time to $5\sigma$, 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 ±24% SNR systematic of Table 27, dominated by the halo density. Idealized live time; the duty and calibration conventions of Table 17 apply (scanning-campaign overheads 10–100×; camped runs far less).

| **Instrument class** | $t_{5\sigma}$ (central) | **syst. range** |
|---|---|---|
| Representative design cavity (SQL 94 mK) | 1.7 s | 1.1 s–2.5 s |
| CAPP-MAX-class single cavity (200 mK) | 4.2 min | 2.9 min–6.3 min |
| EFR-class 18-cavity array (150 mK) | 0.88 s | 0.59 s–1.3 s |
| CAPP-MAX-class cavity, 4 K HEMT | 28 h | 19 h–42 h |

with $w_k$ the fraction of line power in the bin (peak bin: $w = 0.061$). This is the Dicke radiometer statement of Part VII applied bin-wise, with the analysis efficiency $\epsilon_a$ applied explicitly.

# 46 The noise model and the time to detection

**Noise.** Over an integration $t\,\Delta\nu_{\rm bin} \gg 1$ (here $\geq 10^3$ even in the shortest panel), the averaged power in each bin is $\chi^2$-distributed with $2t\,\Delta\nu_{\rm bin}$ degrees of freedom, hence Gaussian to excellent approximation, with standard deviation

$$\sigma_P = \frac{k_B T_{\rm sys}\,\Delta\nu_{\rm bin}}{\sqrt{t\,\Delta\nu_{\rm bin}}} \qquad \Longrightarrow \qquad y_k = \mu_k + n_k, \quad n_k \sim \mathcal{N}(0,1) \text{ i.i.d.} \tag{46.1}$$

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 27.

**Time to detection.** Matched-filtering the full line (bandwidth $\Delta\nu_a$) gives the time to a $5\sigma$ *statistical* detection, $t_{5\sigma} = \Delta\nu_a\,[\,5\,k_B T_{\rm sys}/(\epsilon_a P_{\rm sig})\,]^2$, evaluated per instrument class in Table 24 together with its systematic range: since $t \propto \mathrm{SNR}^{-2}$, the ±24% envelope on the expected SNR (Table 27) maps to a factor ≈ 1.5–1.7 either way on the time, which changes nothing qualitative: seconds stay seconds, minutes stay minutes, hours stay hours.

# 47 Results I: the camped spectra

Figure 13 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 $\mu_k$ of Eq. (45.3): everything about it is computed, nothing fitted — amplitude from Table 23, shape from Eq. (45.2). The dashed orange envelope is the systematic

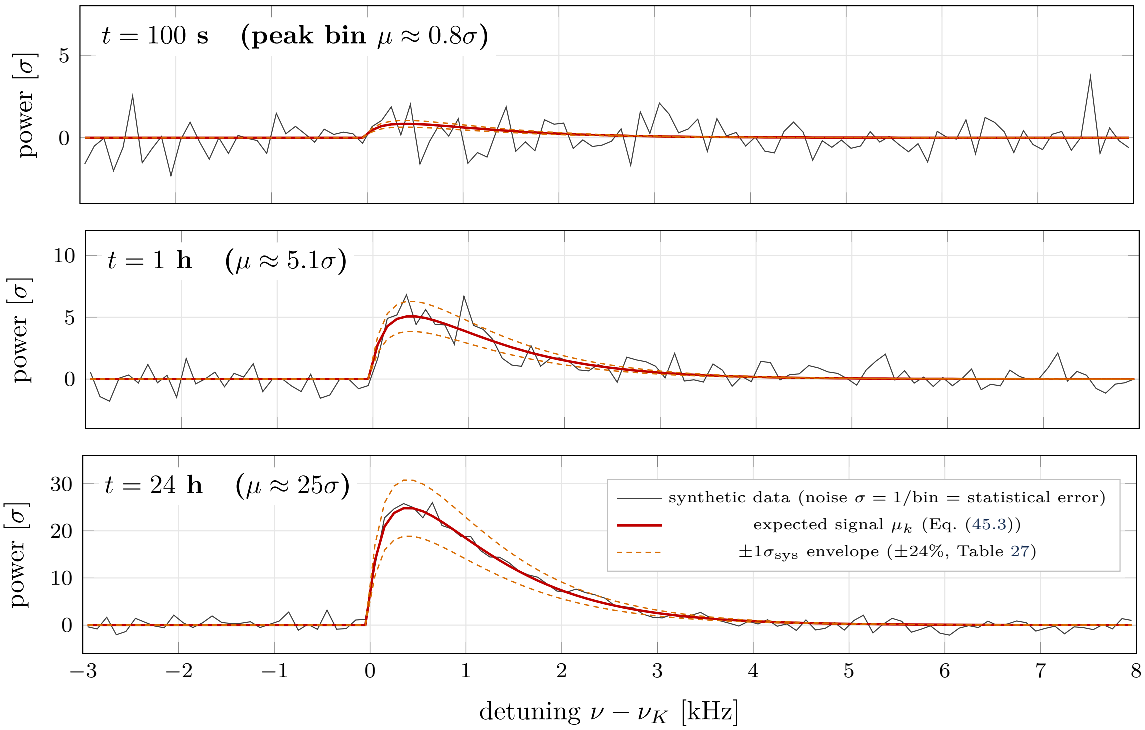


**Figure 13: Results I — the camped spectra, with statistical and systematic errors shown.** Synthetic power spectra ($\Delta\nu_{\mathrm{bin}} = 100$ Hz, units of the per-bin thermal noise $\sigma$; identical fixed-seed noise across panels, §53) for the CAPP-MAX-class camp of Table 22 at $t = 100$ s, 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 $\mu_k$ of Eq. (45.3) — amplitude from the worked power budget of Table 23, shape from Eq. (45.2). Dashed orange: the $\pm 1\sigma_{\mathrm{sys}}$ envelope, $\pm 24\%$ on $\mu_k$ from the total systematic budget of Table 27 (halo density dominated). The peak-bin expectations are $\mu \approx 0.8\sigma$, $5.1\sigma$, and $25\sigma$; the frequency *position* carries the separate theory-side $\pm 6\%$ of Part V, which moves the camp band, not the line within these panels.

band: $\pm 24\%$ on $\mu_k$, the quadrature total of Table 27, dominated by the halo density — visibly a second-order caveat on an unmistakable signal by the third panel, where the peak-bin expectation reaches $25\sigma$ 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 $\pm 6\%$ on $\nu_K$ moves the whole camp band, not the line within a panel — it is the $\pm 120$ MHz band-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 1.5 of Table 24: 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-edge offset above $\nu_K$ oscillates annually as the orbital velocity adds to and subtracts from the Sun's galactic motion,

$$\delta\nu_{\mathrm{edge}}(t) = \frac{\nu_K}{c^2}\Big[\frac{v_\odot^2}{2} + v_\odot v_{\mathrm{orb}} \cos\gamma\, \cos\Big(\frac{2\pi\,(t - t_{\mathrm{Jun}})}{1\,\mathrm{yr}}\Big)\Big] = 588\ \mathrm{Hz} + 74\ \mathrm{Hz} \times \cos(\cdots), \tag{48.1}$$

**Table 25: 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 ±23 Hz, sinusoid-fit precision $\sigma_A = \sigma\sqrt{2/N}$; systematics from Table 21.

| Observable | Period | Amplitude | Stat. (1 yr) | Syst. |
|---|---|---|---|---|
| mean edge offset | — | 588 Hz | ±4.4 Hz | ±5% ($v_\odot^2$) |
| annual modulation | 365.25 d | 74 Hz | ±6.2 Hz ($\sim 12\sigma$) | ±10% ($v_\odot v_{\rm orb}\cos\gamma$) |
| daily modulation | 23.934 h | 1.9 Hz | ±6.2 Hz ($\sim 0.3\sigma$/yr) | site ($\cos\lambda_{\rm lab}$) |

and daily as the rotation velocity sweeps the laboratory through the halo wind,

$$\delta\nu_{\rm day}(t) = \frac{\nu_K\, v_\odot\, v_{\rm rot}\cos\lambda_{\rm lab}}{c^2}\,\cos\Big(\frac{2\pi\, t}{23.934\,\rm h}\Big) \approx 1.9\ \text{Hz}\ \ \text{at Daejeon's latitude.} \tag{48.2}$$

Both are small against the 1.95 kHz 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:

$$\sigma_{\rm cent} \simeq \frac{\rm FWHM}{{\rm SNR}_{\rm line}(t)}, \qquad {\rm SNR}_{\rm line}(t) = \epsilon_a \frac{P_{\rm sig}}{k_B T_{\rm sys}}\sqrt{\frac{t_{\rm live}}{\Delta\nu_a}} \quad\Longrightarrow\quad \sigma_{\rm cent} \approx 23\ \text{Hz per fortnight} \tag{48.3}$$

at a deliberately conservative 1/30 duty cycle ($t_{\rm live} \simeq 11$ h per fortnight; at full duty the same fortnight gives ≈ 4 Hz precision). These are the statistical error bars of Figure 14; the daily panel folds one year of high-duty camping into three-sidereal-hour bins (±2.3 Hz per bin). The systematic envelope on the curve itself (dashed) collects the halo rows of Table 21: ±10% on the annual amplitude, ±5% on the mean offset. Table 25 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: the doublet and single-mode blindness

The Kaxion drives two tones at $\nu_\pm = \nu_0(1 \pm 1/70)$ (Part VI). A receiver reads them through its mode's Lorentzian acceptance,

$$\Big|\frac{\chi(\nu)}{\chi(\nu_+)}\Big|^2 = \Big[1 + \Big(\frac{2Q_L(\nu-\nu_+)}{\nu_+}\Big)^2\Big]^{-1}, \tag{49.1}$$

which at the partner detuning $\Delta\nu/\nu = 1/35$ evaluates to −75 dB for $Q_L = 10^5$ — the single-mode blindness asserted in Part VI, now drawn with its calibration systematic in Figure 15 and tabulated against $Q_L$ in Table 26. The engineering meaning of the ±5% on $Q_L$ is a ∓0.4 dB shift on a 75 dB suppression: the conclusion — *no single-mode instrument sees the doublet by accident* — is systematics-proof by four orders of magnitude. The test is therefore a deliberate act (retune, or

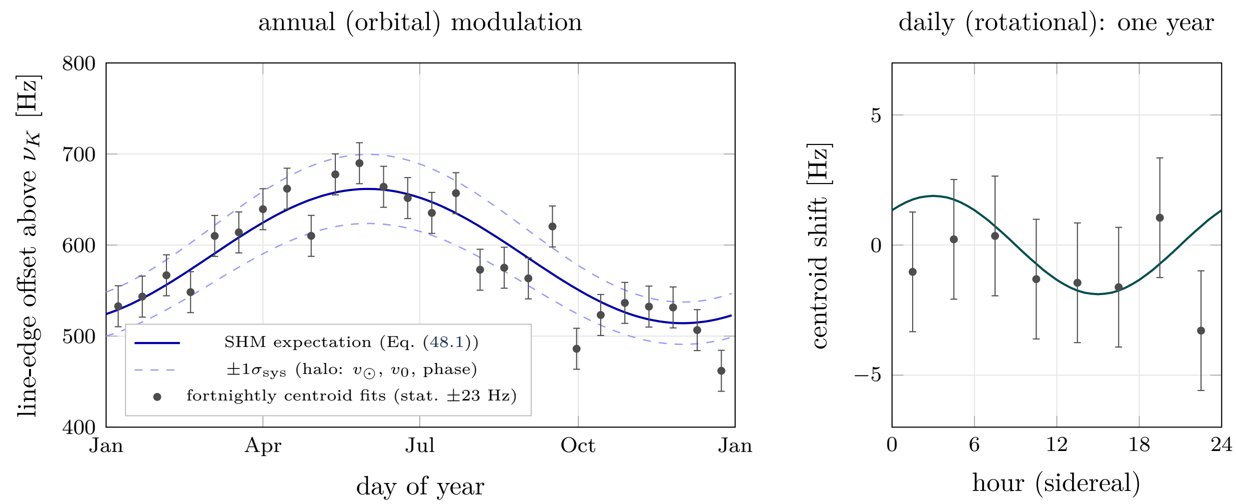


**Figure 14: Results II — the modulation fingerprints, with errors.** Left: the line edge rides at $\nu_K v_\odot^2/2c^2 \simeq 588$ Hz above $\nu_K$ and breathes annually with amplitude $A_{\rm ann} \simeq 74$ Hz (Eq. (48.1); solid blue). Dashed: the halo-side systematic envelope — $\pm 10\%$ on $A_{\rm ann}$ from $v_\odot = 232.8 \pm 3.0$ km/s and the orbit-projection geometry, $\pm 5\%$ on the mean offset ($\propto v_\odot^2$) (Table 21). Points: simulated fortnightly centroid fits with their *statistical* error of $\pm 23$ Hz (Eq. (48.3), at a conservative 1/30 duty cycle). Right: the daily (sidereal) signature at the latitude of Daejeon, amplitude $\simeq 1.9$ Hz (Eq. (48.2)); site-dependent through $\cos\lambda_{\rm lab}$. Points: one *year* of high-duty camping, folded into three-sidereal-hour bins, statistical error $\pm 2.3$ Hz per bin — even a high-duty year leaves the daily signature at only $\sim 2\sigma$, and the conservative 1/30-duty fortnightly program does not resolve it at all (Table 25). At single-cavity aperture the sidereal fingerprint is a hint, not a verdict — 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.

**Table 26: Single-mode suppression of the 1/35 partner versus $Q_L$** (Eq. (49.1) at $\Delta\nu/\nu = 1/35$). At every realistic $Q_L$ the partner is invisible without a deliberate second measurement; conversely a dual-mode receiver regains it in full. The row to read is the one at the realistic loaded $Q$, not the best case.

| $Q_L$ | **partner suppression (dB)** |
|---|---|
| 50000 | −69.1 |
| 100000 | −75.1 |
| 200000 | −81.2 |
| $10^6$ | −95.1 |

read the Part-VIII second mode), and at the camped SNRs of Table 24 its verdict arrives within a day, carrying the same statistical and systematic structure as Figure 13.

# 50 Results IV: the channel-ratio verdict

The second lock is read by operating the electromagnetic port and the electron-magnon port (Part VIII) in the same condensate and forming the deconvolved ratio $\widehat{R}$ of Appendix D.2, whose

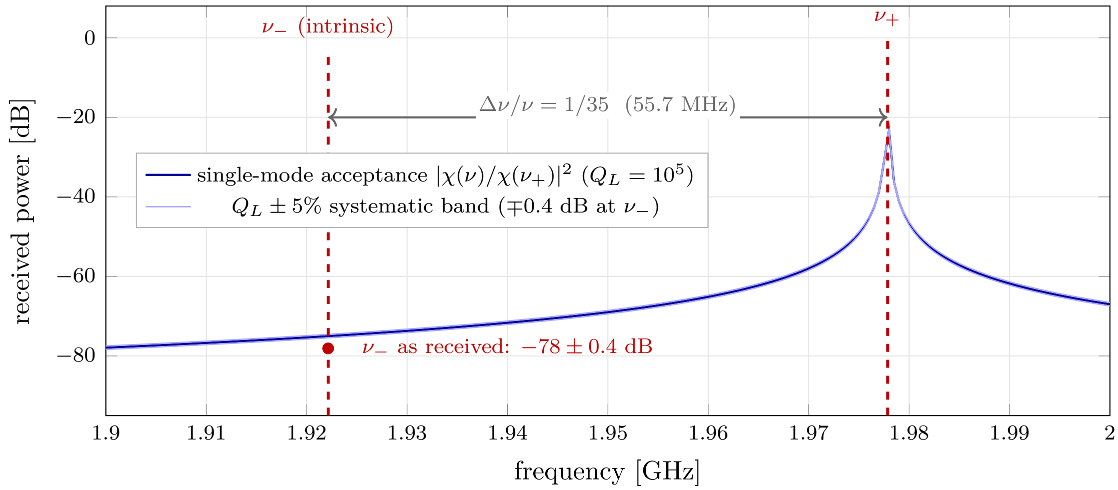


**Figure 15: Results III — the doublet, and why it must be asked for.** The Kaxion drives two tones at $\nu_\pm = \nu_0(1 \pm 1/70)$ (dashed red; tone ratio $\kappa$ not locked, partner drawn at $-3$ dB for illustration). A single-mode receiver tuned to $\nu_+$ applies the Lorentzian acceptance of Eq. (49.1) (dark blue) and receives the partner suppressed by 75 dB — at $-3$ dB intrinsic, $-78$ dB, invisible in any realistic exposure. Light blue: the $Q_L \pm 5\%$ calibration systematic (Table 22), $\mp 0.4$ dB on the suppression — immaterial to the conclusion (Table 26). The 1/35 test (F-Kaxion-1a) is therefore a deliberate act: retune, or instrument a second mode at the geometric spacing (Part VIII), and the partner appears at camped SNR within a day. Statistical errors at the retuned mode are those of Figure 13.

expectation is the locked $P_B/P_\gamma = 25/4$. Its two error classes separate cleanly:

$$\left.\frac{\sigma_{\widehat{R}}}{R}\right|_{\text{stat}} = \sqrt{\frac{1}{\mathrm{SNR}_B^2} + \frac{1}{\mathrm{SNR}_\gamma^2}} \xrightarrow{\mathrm{SNR}_\gamma \gg \mathrm{SNR}_B} \frac{1}{\mathrm{SNR}_B(t)}, \qquad \left.\frac{\sigma_{\widehat{R}}}{R}\right|_{\text{sys}} = \sigma_{\text{cal}}, \tag{50.1}$$

with $\mathrm{SNR}_B(t) = 5\sqrt{t/5\,\mathrm{d}}$ the B-port design requirement of Table 22 (a fenced engineering input, not a theorem); in the error bars and milestones we conservatively carry $\sqrt{2}\times$ this statistical limit per ratio, absorbing sub-dominant $\gamma$-side and spectral-window losses. $\sigma_{\text{cal}}$ is the two-port calibration: $\pm 20\%$ for a first-generation pair, $\pm 5\%$ for the design instrument — the F-Kaxion-1b staging of Part VI, drawn as the stepped band of Figure 16. Everything else cancels in the ratio by construction (Theorem 32.1): the halo density, the magnet, the local amplitude, every frame-dependent factor — which is why the observable that carries the deepest structural content also carries the shortest systematic table. A converged $\widehat{R}$ off 25/4 kills the Kaxion; on it, the gauge vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$ — and with it $K = 35$, a second, independent time — is read in a laboratory.

# 51 The systematic-error budget

Table 27 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 $\rho_{\text{loc}}$ 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

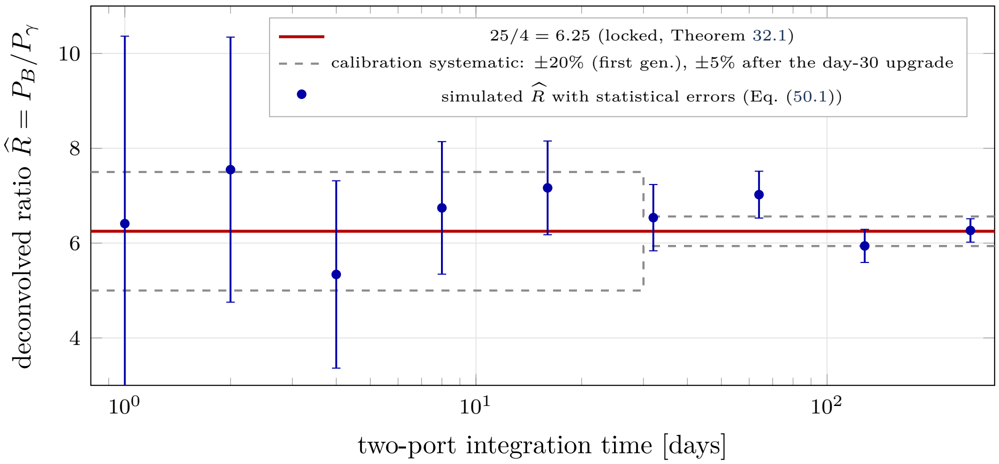


**Figure 16: Results IV — the 25/4 verdict, with both error classes drawn.** The deconvolved matched-channel ratio $\widehat{R}$ of Appendix D.2, simulated for a two-port camp (electromagnetic + electron-magnon, Part VIII). Error bars: the *statistical* error of Eq. (50.1), conservatively carried as $\sigma_{\widehat{R}}/R = \sqrt{2}/\mathrm{SNR}_B(t)$, with the B-port design requirement $\mathrm{SNR}_B = 5\sqrt{t/5\,\mathrm{d}}$ (Table 22). Dashed step: the *calibration* systematic — ±20% for the first-generation port pair, tightening to the design ±5% at the simulated day-30 calibration upgrade (F-Kaxion-1b staging, §51). Halo density, magnet field, and every frame-dependent factor cancel in the ratio (Theorem 32.1), which is why its systematic budget is calibration only. A converged $\widehat{R}$ off the red line kills the Kaxion; on it, the gauge vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$ — the geometry — is read in a laboratory.

verdicts. *Second*, the identification observables are engineered to cancel it: the doublet spacing is a frequency ratio (no amplitude enters), and the channel ratio cancels density, field, and amplitude by construction; that is why the two falsifiers that carry the geometry (F-Kaxion-1a, 1b) sit in the two shortest rows. *Third*, the theory-side errors of Table 20 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.

## 52 Discovery week — and what failure looks like

With the machinery of §45–51 on the table, the pre-declared sequence of Part VI reads as a schedule. *Day zero:* the excess of Figure 13 (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 partner — a deliberate retune or the second mode, Figure 15; F-Kaxion-1a fires or is passed 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 14. *Weeks to months:* the B-port joins, and $\widehat{R}$ converges through the stepped band of Figure 16 — F-Kaxion-1b, the second reading of $K = 35$. 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, a partner at one part in thirty-five, and a second channel converging on twenty-five over four — every number of it computed, with its errors, in the figures of this Part before any of it happened.

**Table 27: The systematic budget, in one table.** Fractional effects on the expected per-bin SNR unless noted. Quadrature totals: amplitude-instrument ±7.7% (with the $T_{\rm sys}$ ±6.4% noise-floor row, combined instrument ±10%); with the halo density, ±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 |
|---|---|---|---|---|
| $\rho_{\rm loc}$ | halo | ±22% | amplitude, $t_{5\sigma}$ | surveys |
| $v_0$ (line width) | halo | ±8% | shape, $\Delta\nu_a$ | surveys |
| $v_\odot$, $\cos\gamma$ | halo | ±10% ($A_{\rm ann}$) | modulation | astrometry |
| $B_0^2$ | instrument | ±2% | amplitude | NMR field map |
| $V$ (retuned) | instrument | ±2% | amplitude | metrology + contraction |
| $C_{010}$ | instrument | ±5% | amplitude | bead pull + FEM |
| $Q_L$ | instrument | ±5% | ampl.; ∓0.4 dB suppr. | VNA $S$-parameters |
| $T_{\rm sys}$ | instrument | ±6.4% | noise floor | Y-factor + shot noise |
| $\epsilon_a$, gain drift | analysis | ±1% | amplitude | signal injection |
| two-port calibration | instrument | ±20% → ±5% | $\widehat{R}$ only | end-to-end injection |
| doublet spacing 1/35 | — | immune | — | frequency ratio |
| channel ratio 25/4 | — | calibration only | — | amplitudes cancel |

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

$$g_{90}(\nu) \;=\; g_{K\gamma\gamma}^{\rm DFSZ}(\nu)\left[\frac{1.28}{\epsilon_a}\,\frac{k_B T_{\rm sys}}{P_{\rm sig}^{\rm DFSZ}}\sqrt{\frac{\Delta\nu_a}{t}}\;\right]^{1/2}\!\times\left(1+\sigma_{\rm sys}\right)^{1/2}, \tag{52.1}$$

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 $\Xi_{\rm cap}$ rung excludes that rung at DFSZ depth, doublet-aware; 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 with a doublet-aware analysis and no 1/35 pair 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 is equally fatal (F-Kaxion-1d). Either the logbook above gets written by the early 2030s, or the ledger entry of Appendix E.2 closes this paper — and the series' dark-matter claim — for good.

# 53 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)–(52.1) from the input tables verbatim and draws the noise realizations from fixed seeds — seed 35 for the spectra, modulation, and doublet figures, seed 4 for the ratio points — so that every panel regenerates bit-identically. The script will be released open-source (GitHub, mirrored at arisaka-gg.org) together with the paper, with the input tables as a separate, editable parameter block: a collaboration wishing to re-run the simulation for its own instrument replaces Table 22's nine 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.

## 54 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 540 MHz window, pre-registered down to the bin; an *amplitude* — the worked budget of Table 23, good to ±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 1.9 Hz; a *partner tone* at exactly one part in thirty-five; and a *second channel* locked at exactly twenty-five over four. 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. The Kaxion keeps none.

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 two are dimensionless ratios fixed by the geometry alone. 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, bring a companion at 2.857%, and answer a magnon port at 6.25 — 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: §69.)

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, lacks its partner at the $10^{-2}$ level, or converges off 25/4 does not become an interesting anomaly — it becomes the death certificate of this paper (F-Kaxion-1a through 1d, each already priced in dates and depths). A theory that can be confirmed by six 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, and a geometry has no dials to turn. The five integers $(1, 3, 5; 7, 35)$ fixed the frequency, the coupling, the split, and the ratio before any experimentalist was asked for anything; the split and the ratio each read $K = 35$ independently, so a single week of data would measure the depth of the scale throat twice, in two different kinds of apparatus, and demand the same integer of both. That is what it looks like when a theory stakes everything on numbers it cannot 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 tone, one channel at a time; but the six signatures were meant to be read *together*, in coincidence, 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 two sub-arrays tandem-locked at the doublet ratio and a harmonic booster on the baryonic port, and lets it run for a year: four peaks at once, every falsifier of this paper executed in one continuous experiment.

## 55 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 doublet, the ratios — 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, what the two peaks look like beside each other, what the ratio of their strengths should be. 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: two tones, two channels, four peaks, one year

## 56 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 two parameter-free locks — the 1/35 doublet and the 25/4 channel ratio (Part VI) — neither of which any existing instrument reads; 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 **Kax**ion **O**bservatory (*KAX-ee-oh*) — a major upgrade of ADMX-EFR at Fermilab that reorganizes its eighteen cavities into *two* nine-cavity sub-arrays tandem-tuned at the locked ratio $\nu_-/\nu_+ = 69/71$, adds the harmonic-booster module of Part VIII — the electron port, stage one of the two-stage baryonic program — to read the spin channel at both tones, and runs steadily for one year. The instrument observes *four peaks simultaneously* — two tones × two channels — and thereby executes, in a single continuous run, every test this paper has pre-registered: F-Kaxion-1a (the doublet) within the first week, F-Kaxion-1b (the ratio) to ±7% total within a quarter and to the systematics-limited ±5.5% by year end, the daily modulation emerging at $\sim 3\sigma$ within a month (an $\sim 11\sigma$ measurement by year end), the annual modulation across the year, and, on the null tape, a DFSZ-depth doublet-aware 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 three genuinely new items: the tandem lock, the booster module, and a four-channel coincidence data-acquisition layer (Table 36). 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.

## 57 Design requirements

Most instruments derive their requirements from what technology makes convenient; KAXIO derives its eight from what the theory sheet makes *mandatory*. Each row of Table 28 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 three 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 — pair coverage of $\nu_0 \in [1.59, 2.13]$ GHz — is the closed window of the discrete $\Xi_{\rm cap}$ 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 tones held at $\nu_-/\nu_+ = 69/71$ — is the requirement that turns a detector into an identifier: the 1/35 doublet of Part VI is the Kaxion's signature, and holding the pair in hardware, at $10^{-4}$ ratio accuracy (roughly an order below the $(5/8)K^{-2} \simeq 5 \times 10^{-4}$ theoretical correction), means every second of data is doublet-aware — F-Kaxion-1a sits in the trigger, not in a follow-up proposal.

*The identification block (R4–R5).* R4 — a spin (electron) port at both tones — exists because the second lock, $P_B/P_\gamma = 25/4$, is unreadable by any purely electromagnetic instrument; the port is Part VIII's electron-magnon channel, which reads the $J_e$ lock directly, the baryonic 25/4 verdict following under the fenced chiral input (open item iii); placing the port at *both* tones doubles the coincidence demand and yields two independent ratio measurements whose agreement is itself a test. R5 — ratio calibration at ±5% — is set by the verdict clock of §63: the ratio's statistical error crosses 5% near day eighty of live time, so a coarser calibration would waste the year's statistics, while a finer one would be paid for without physics return; ±5% resolves 25/4 = 6.25 from unity, from zero, and from any competing axion model's nucleon coupling with margin.

*The modulation block (R6–R8).* R6 — one steady year at ≥ 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. R7 — 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. R8 — 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 1.8 Hz daily amplitude, so that neither modulation can be manufactured or masked by the apparatus; it is implemented by the continuous calibration tone of §60 and audited by the sidereal/solar discrimination itself.

**Table 28: 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-1c |
| R2 | pair coverage $\nu_0 \in [1.59, 2.13]$ GHz | the closed window (Part V) | F-1c |
| R3 | two tones held at $\nu_-/\nu_+ = 69/71$ | the 1/35 doublet (Part VI) | F-1a |
| R4 | electron (spin) port at both tones | the 25/4 lock (Part VI) | F-1b |
| R5 | ratio calibration ±5% | F-1b design staging (Part IX) | F-1b |
| R6 | one-year continuous run, duty ≥ 85% | annual modulation (Part IX) | identification |
| R7 | absolute timing $\ll 1$ min over a year | sidereal-day discrimination | identification |
| R8 | centroid stability ±1 Hz/season | modulation systematic floor | identification |

# 58 Detector design

Figure 17 is the concept; Table 29 the specification. Three subsystems carry the physics.

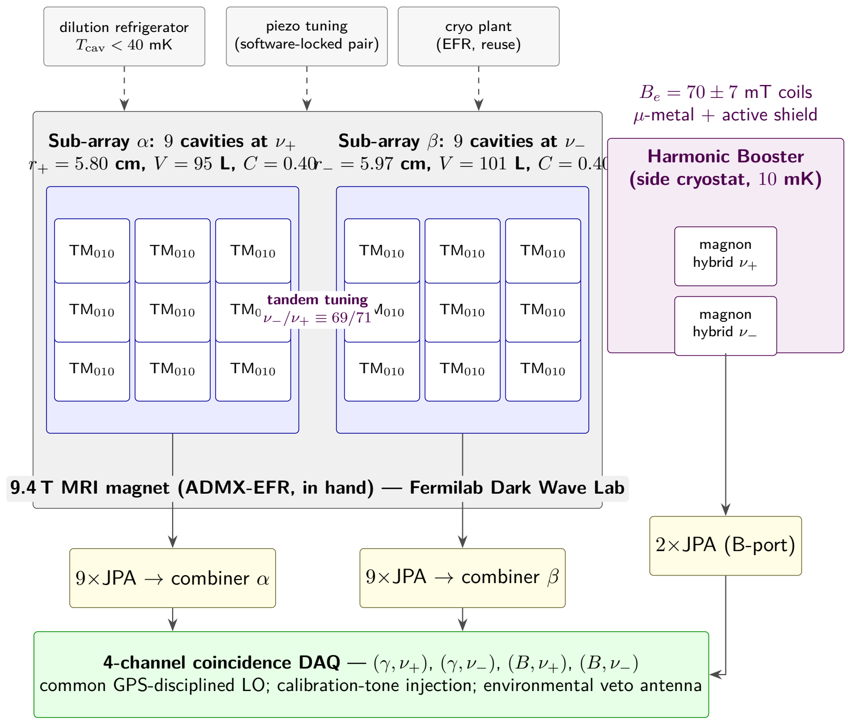


**Figure 17: KAXIO conceptual design.** The eighteen ADMX-EFR cavities are reorganized into two nine-cavity sub-arrays inside the existing 9.4 T magnet: sub-array $\alpha$ tuned to the upper tone $\nu_+$ and sub-array $\beta$, with 2.9% larger cavities, to the lower tone $\nu_-$, their piezo tuners software-locked to the geometric ratio $\nu_-/\nu_+ = 69/71$ so the pair scans the window as one object (Table 29). The harmonic-booster module — two electron-magnon hybrids at $B_e = 70$ mT in a magnetically shielded side cryostat (Part VIII, Concept B) — reads the spin channel at both tones. Four data streams enter a common-clock coincidence DAQ: 2 tones × 2 channels = 4 simultaneous peaks, the full fingerprint of Part VI in one instrument. Reused EFR infrastructure is drawn gray; new elements are the tandem lock, the booster, and the four-channel DAQ (Table 36).

**The twin tandem-tuned sub-arrays.** The eighteen EFR cavities are rebuilt as two families of nine. Sub-array $\alpha$ operates at the upper tone: $\mathrm{TM}_{010}$ radius $r_+ = 5.80$ cm at the benchmark $\nu_+ = 1.9779$ GHz. Sub-array $\beta$ is machined 2.9% larger ($r_- = 5.97$ cm) and holds the lower tone $\nu_- = 1.9221$ GHz. Each sub-array is coherently power-combined as in the EFR baseline design (ADMX-EFR Collaboration, 2024); the effective volumes are 95 L ($\alpha$) and 101 L ($\beta$) at $C_{010} = 0.40$. The single genuinely new mechanical element is the *tandem lock*: the two piezo tuning systems are driven by one control loop that holds the frequency ratio at 69/71 to $10^{-4}$ (roughly an order below the $(5/8)K^{-2}$ correction of Part VI) at every step, so the pair sweeps the window of R2 as one rigid object — a doublet-shaped net. Both sub-arrays sit in the existing 9.4 T bore and the existing sub-40 mK cold volume; against ADMX-EFR's baseline the modification is a re-machining and a control-software layer, not a new cryostat.

**Table 29: KAXIO specification.** Baseline ADMX-EFR values from the collaboration's design (ADMX-EFR Collaboration, 2024); KAXIO entries are this proposal. The booster sensitivity row is a design *requirement* (fenced), to be demonstrated in the pre-construction R&D year. One row is a design requirement rather than a measured value, and it is marked as such rather than left to be inferred.

| **Subsystem** | **ADMX-EFR baseline** | **KAXIO** |
|---|---|---|
| magnet | 9.4 T MRI, 0.8 m bore | unchanged (in hand) |
| cavities | $18\times\mathrm{TM}_{010}$, one frequency | $2\times 9$, tandem-locked 69/71 |
| cavity radii | single design | $r_+ = 5.80$ cm, $r_- = 5.97$ cm |
| volume (eff.) | $\sim 200$ L | 95 L + 101 L |
| form factor $C_{010}$ | 0.40 | unchanged |
| $Q_L$ | $10^5$ (SC-coated) | unchanged |
| $T_{\rm sys}$ | 150 mK (JPA chain) | unchanged; photon-counting upgrade path |
| tuning | piezo, single | piezo pair, software-locked ($10^{-4}$) |
| electron (spin) port | — | $2\times$ magnon hybrid, $B_e = 70$ mT, 10 mK |
| booster sensitivity | — | $\mathrm{SNR}_B = 5\sqrt{t/5\,\mathrm{d}}$ per tone (req.) |
| DAQ | single-stream | 4-channel coincidence, GPS clock, cal tone |
| run plan | 3-yr scan 2–4 GHz | 1-yr camp + tandem sweep of R2 |

**The harmonic-booster module.** The spin port — Part VIII's electron channel (Concept B), stage one of the two-stage baryonic program: it reads the $J_e$ lock directly, the baryonic 25/4 verdict following under the fenced chiral input — is 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 $\nu_+$ (first hybrid) and $\nu_-$ (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 $\mathrm{SNR}_B = 5\sqrt{t/5\,\mathrm{d}}$; the pre-construction R&D program (one year, bench-top, before the main build) exists precisely to demonstrate this figure, and the proposal stands or falls with it honestly declared (here and throughout, booster exposure times count *live* days; the run plan's duty budget converts them to calendar time).

**The four-channel coincidence layer.** All four data streams — $(\gamma,\nu_+)$, $(\gamma,\nu_-)$, $(B,\nu_+)$, $(B,\nu_-)$ — are digitized against one GPS-disciplined local oscillator, with a continuously injected calibration tone in every chain (R8) and the standard environmental veto antenna. The coincidence demand is the background killer: a candidate must appear in both $\gamma$ streams at the locked spacing to survive the first week, and in all four to claim identification. No artifact class — interference, thermal, mechanical — reproduces two locked frequencies in two physically distinct channel types.

# 59 The harmonic booster: engineering design

The booster is the one subsystem without a running precedent at exactly this operating point, so this section carries its engineering case at full depth: the coupling budget, the linewidth budget, the

**Table 30: 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 19.

| Parameter | Value | Basis / note |
|---|---|---|
| ferrimagnet | YIG spheres, ∅2 mm ×10 | demonstrated material class |
| spin density | $2.1 \times 10^{28}$ m$^{-3}$ | YIG; $N_s \simeq 8.8 \times 10^{19}$ per sphere, $8.8 \times 10^{20}$ per hybrid |
| bias field $B_e$ | 70.6 mT ($\nu_+$), 68.6 mT ($\nu_-$) | Kittel mode on each tone |
| bias homogeneity | $\Delta B/B < 10^{-4}$ over array | keeps inhomogeneous broadening subdominant |
| single-sphere coupling $g_{cm}/2\pi$ | $\simeq 40$ MHz | demonstrated class; array scales as $\sqrt{N_{\rm sph}}$ |
| magnon linewidth $\kappa_m/2\pi$ | $\simeq 1.0$ MHz | Gilbert damping + inhomogeneity |
| cavity linewidth $\kappa_c/2\pi$ | $\simeq 20$ kHz | $Q_L = 10^5$ |
| cooperativity $g^2_{cm}/\kappa_c\kappa_m$ | $\sim 10^5$ | strongly hybridized (level repulsion resolved) |
| readout, first generation | JPA, $n_{\rm add} \simeq 1.5$ | week-one row of Fig. 19 |
| readout, upgrade | photon counter, dark-count limited | *req.*: 3× SNR (bottom row) |
| sensitivity, first generation | $\mathrm{SNR}_B = 5\sqrt{t/5\,\mathrm{d}}$ per tone | *req.*; R&D-year exit criterion |
| sensitivity, upgraded | 3× first generation | *req.*; year-two insertion |

shield design, and — most importantly for the run plan — the noise staging from the first-generation chain to the quantum limit.

**Why the spin-port signal is weak, and why that is not fundamental.** The middle row of Figure 19 shows a week-one booster signal of only $\mathrm{SNR}_{\rm line} \simeq 5.9$ per tone against the $\gamma$ ports' peak-bin 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 $n_{\rm add} \simeq 1.5$ 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 $\bar{n} = [e^{h\nu/k_BT} - 1]^{-1} \simeq 8 \times 10^{-5}$: *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 $n_{\rm add}$ penalty and the wideband integration penalty together; the R&D program's quantitative goal is the conservative 3× SNR of Table 30, which, one month after the year-two readout insertion, turns the booster spectra into the bottom row of Figure 19: an $11\sigma$ 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 ($\sim 1.5$ 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 $\pm 2$ 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 four-peak coincidence carries identification weight.

# 60 The RF and readout chain

Figure 18 draws one chain end to end; four such streams run in parallel, the two $\gamma$ streams nine-fold coherently combined after their JPA stages as in the EFR baseline. Three design choices matter for the physics. *One clock:* every mixer in every stream runs from a single GPS-disciplined local oscillator, so the four streams share phase provenance and the modulation analyses inherit absolute timing at the $10^{-8}$ level (R7). *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 $\pm 1$ Hz seasonal floor (R8) — and, because it is injected *after* the cavity, it cannot mimic a cavity-resonant halo line. *Real-time coincidence:* the FPGA layer demands the locked frequency offsets (69/71 between tones; identity between channels) before anything is flagged, so the trigger itself embodies the geometry being tested.

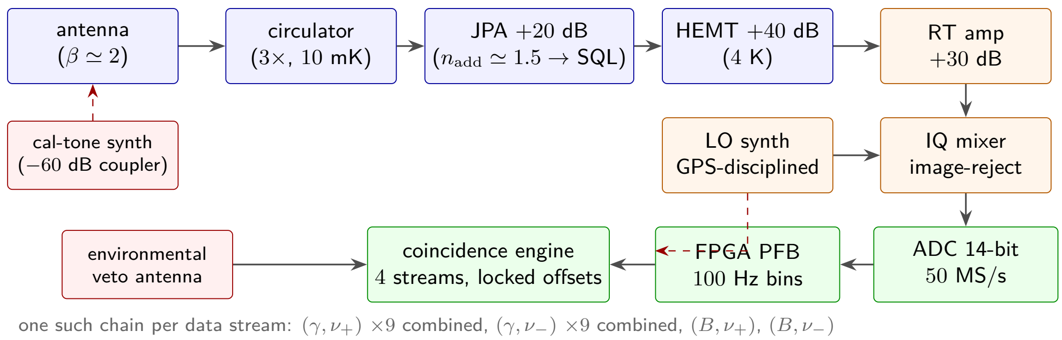


**Figure 18: The RF/readout chain** (one per data stream; the two $\gamma$ streams are nine-fold coherently combined after the JPA stage). Cold section at 10 mK: over-coupled antenna, triple circulator, flux-pumped JPA at +20 dB with $n_{\rm add} \simeq 1.5$ quanta first-generation, replaced by a photon-counting front end at the SQL in the upgrade path (Table 30); 4 K HEMT and room-temperature stages follow. All mixers run from one GPS-disciplined local oscillator (R7); 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 (R8). The coincidence engine demands the locked frequency offsets of the four streams in real time.

## 61 Sensitivity and noise

Table 31 carries the full budget, evaluated with the registered machinery of Part IX (Eqs. (45.1)–(46.1)) at the benchmark pair. The headline: each sub-array camps its tone with $P_{\rm sig}/k_B T_{\rm sys} = 80$ Hz at DFSZ coupling — a matched-line $5\sigma$ in 8.7 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.

## 62 Systematic errors

Table 32 extends the Part-IX budget (Table 27) 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 both locks cancel it, which is why the instrument is built around the locks. Second, the one systematic that could mimic physics — a *seasonal* instrument drift masquerading as annual modulation — is attacked in hardware (R8): 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.

## 63 Simulated performance

All figures of this section are generated by the Part-IX pipeline (Fig. 12) with the instrument block replaced by Table 29 — 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 public script

**Table 31: Sensitivity and noise at the benchmark pair** ($\nu_0$ = 1.95 GHz, DFSZ coupling, $\rho_{\rm loc}$ = 0.45 GeV/cm$^3$; registered machinery of Part IX; live time at 85% duty). The booster rows carry the fenced design requirement; every other row is computed. Halo-density systematic (±22%) applies to amplitudes as in Table 27 and cancels in both locks.

| Quantity | Value | Basis |
|---|---|---|
| $P_{\rm sig}$ per sub-array | $1.66 \times 10^{-22}$ W | Eq. (45.1); 9.4 T, 95 L (conservative: the smaller sub-array), $C = 0.40$, $Q_L = 10^5$, antenna $\beta \simeq 2$, i.e. $\beta/(1+\beta) = 2/3$ |
| photon rate per tone | $1.3 \times 10^2$ s$^{-1}$ | $P_{\rm sig}/h\nu$ |
| noise density $k_B T_{\rm sys}$ | $2.1 \times 10^{-24}$ W/Hz | 150 mK JPA chain, i.e. $n_{\rm add} \simeq$ 1.1 above the 94 mK vacuum term (Y-factor ±6.4%) |
| $P/k_B T_{\rm sys}$ | 80 Hz | ratio of the above |
| $t_{5\sigma}$ per $\gamma$ tone | 8.7 s (live) | matched line, $\epsilon_a = 0.93$ |
| week-1 peak-bin SNR | $\approx 3.3 \times 10^2$ per $\gamma$ tone | stat. only (±24% syst.); Fig. 19; shape fully resolved |
| booster line SNR, week 1 | 5.9 per tone | design requirement (R&D deliverable) |
| centroid precision (fortnight) | ±0.58 Hz stat. | Eq. (48.3), both $\gamma$ streams (the two tones) stacked |
| daily-modulation reach | $3\sigma$ in one month; $11\sigma$ over the year | Fig. 20 |
| annual-modulation reach | syst.-limited at ±1 Hz | Fig. 21; R8 |
| $\widehat{R}$ statistical, 1 yr | ±2.3% (two ratios) | Eq. (50.1) |
| $\widehat{R}$ total, 1 yr | ±5.5% | incl. ±5% design calibration (R5) |
| window sweep on null | full R2 range, DFSZ, doublet-aware | tandem sweep; executes F-Kaxion-1c |

(§53). 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 four simultaneous spectra as they stand at the end of week one, including the booster's two readout generations (Figure 19); then the daily modulation as it sharpens from week one to month one (Figure 20); then the year-long annual curve, the observable that forces the calendar requirement R6 (Figure 21); and finally the verdict clock, which plots the running precision of all three decision observables against the four milestones of the run plan (Figure 22). 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 34.

**Week one: the four peaks — and the booster's upgrade path.** Figure 19, now six panels. The two $\gamma$ ports saturate immediately — peak-bin SNR near $3.3 \times 10^2$ (stat. only; the ±24% amplitude systematic scales it), line shapes fully resolved, spacing read at $10^{-4}$ — while the booster

**Table 32: KAXIO-specific systematics**, additional to the inherited budget of Table 27. The seasonal-drift row is the one channel through which an instrument effect could imitate physics; its mitigation is hardware (R8) 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 |
| tandem-lock ratio | $10^{-4}$ | doublet spacing | negligible vs $(5/8)K^{-2} \simeq 5 \times 10^{-4}$ |
| sub-array cross-talk | $< -80$ dB | fake coincidence | isolation + offset LO scheme |
| booster bias field $B_e$ | $\pm 0.1$ mT servo | magnon tuning | active shield + NMR probe |
| booster calibration | $\pm 5\%$ (design) | $\widehat{R}$ | end-to-end injection per port (R5) |
| seasonal LO/thermal drift | $\pm 1$ Hz residual | annual centroid | continuous cal tone (R8); sidereal check |
| clock/timing | $< 10^{-8}$ | modulation phase | GPS-disciplined LO (R7) |
| duty-cycle gaps | 15% | exposure bookkeeping | persistent-mode magnet; redundancy |

ports show the same two lines at their first-generation week-one significance (middle row) and, after the photon-counting upgrade of Table 30, at an unambiguous $11\sigma$ peak bin within a month of the year-two readout insertion (bottom row) — the visual demonstration that the spin channel's faintness is a readout property, not a floor of nature (§59). F-Kaxion-1a is settled (either way) inside the first week, by construction rather than by luck.

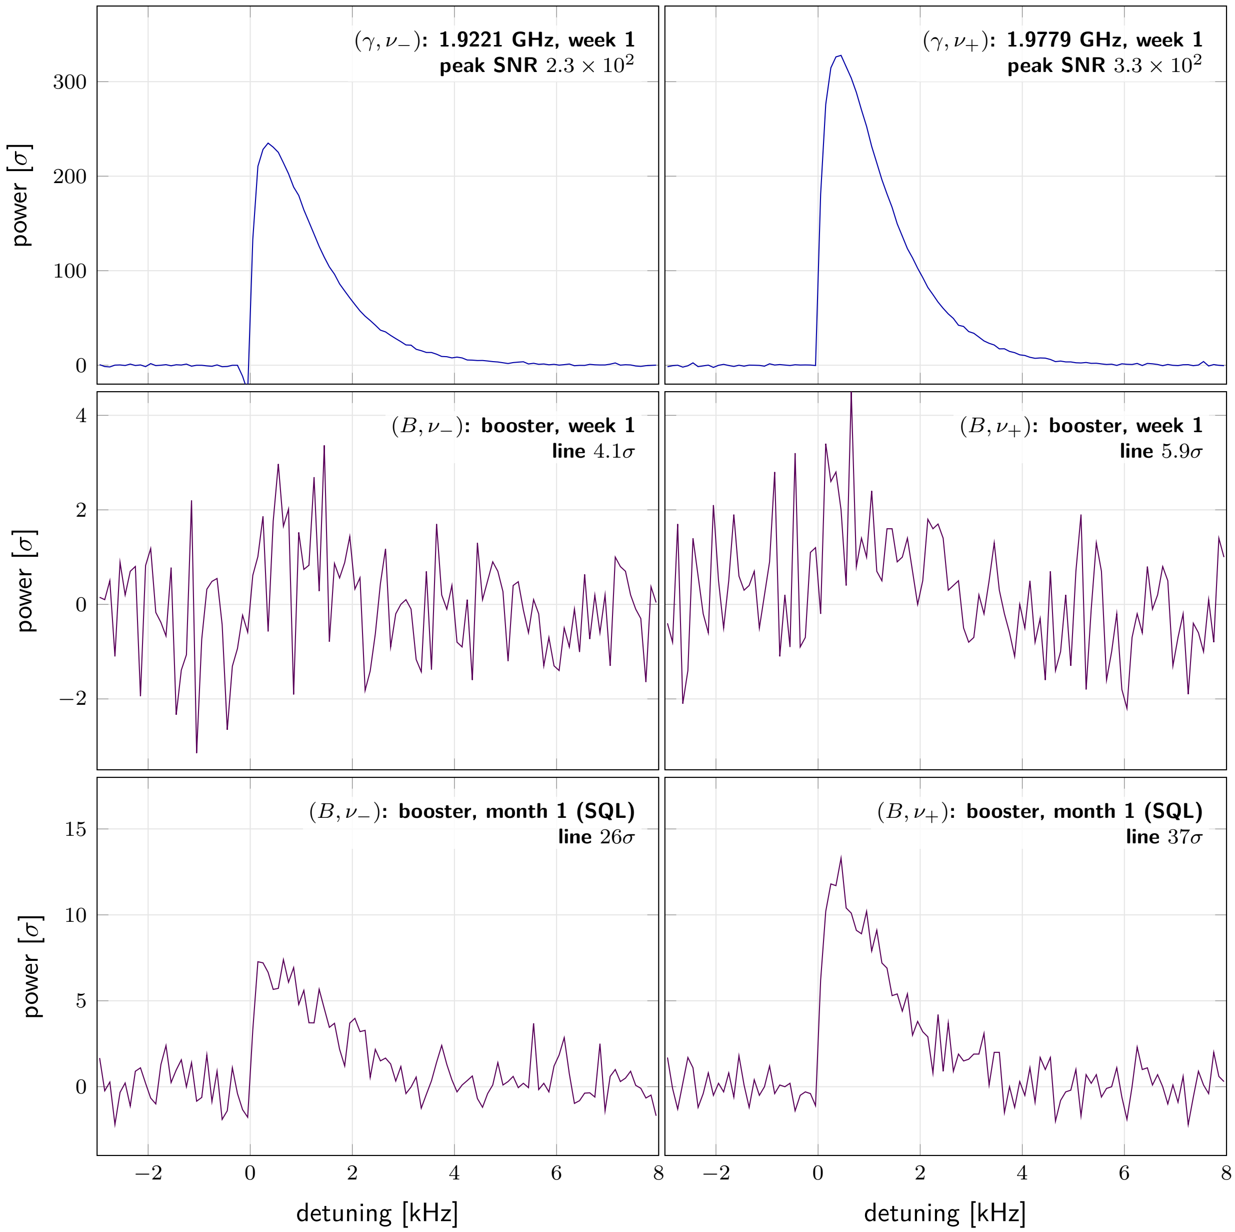


**Figure 19: The four peaks, and the booster's road to the quantum limit** (fixed-seed simulation, KAXIO parameters of Table 31). Left column is the lower tone $\nu_-$, right column the higher tone $\nu_+$, drawn at a representative amplitude ratio $\kappa \simeq 0.7$ — the tone partition is not parameter-free (Appendix B.6.2), so the per-tone design values quoted below are those of the reference tone $\nu_+$. *Top row:* the electromagnetic ports after one week — the reference tone $\nu_+$ camps at a peak-bin SNR $\sim 3.3 \times 10^2$ (stat. only; $\pm 24\%$ syst.), $\nu_-$ at $\kappa$ 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 ($\mathrm{SNR}_{\mathrm{line}} \simeq 5.9$ for $\nu_+$): faint, because the per-spin coupling is minute and the first-generation chain adds $\sim 1.5$ quanta of amplifier noise over a magnon linewidth $500\times$ the signal line (§59) — 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 30): $\mathrm{SNR}_{\mathrm{line}} \simeq 37$ for $\nu_+$, 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 $\sigma = 1/\mathrm{bin}$ by construction throughout.

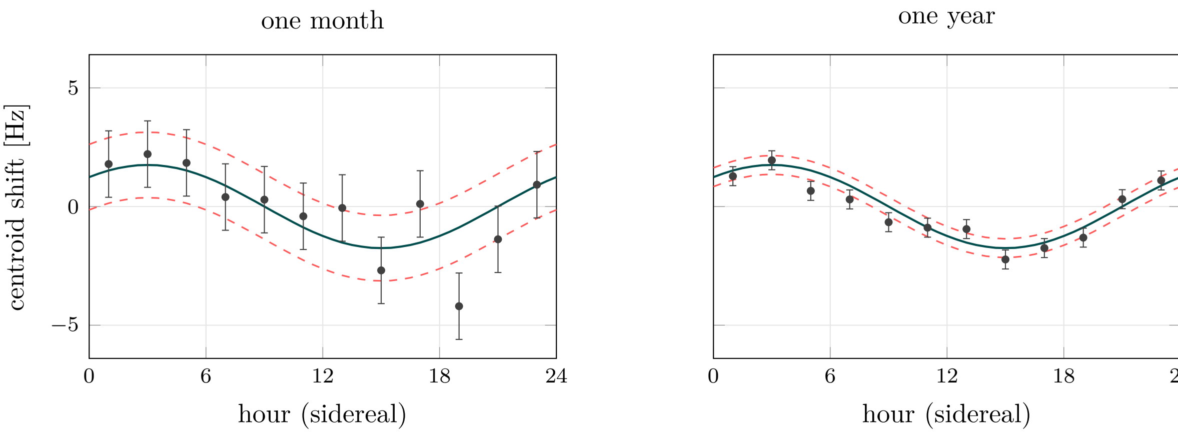


**Figure 20: The daily (sidereal) modulation at Fermilab's latitude ($41.8°$N, amplitude $1.8$ Hz), resolved within the first months.** Both panels: two-sidereal-hour centroid bins at 85% duty, both $\gamma$ streams (the two tones) stacked; curve, the SHM expectation (Eq. (48.2)); bars, the statistical error of Eq. (48.3) — $\pm 1.4$ Hz per bin after one month (left, a $\sim 3\sigma$ amplitude measurement), $\pm 0.40$ Hz after one year (right, $\sim 11\sigma$). The contrast with the single-cavity camp of Figure 14 (bars $\pm 2.3$ Hz even after a high-duty year, larger than the amplitude) is the KAXIO aperture: $P_{\rm sig}/k_B T_{\rm sys}$ five times larger and the two $\gamma$ streams stacked at 85% duty, i.e. nearly an order of magnitude in centroid precision; the amplitude itself is also slightly smaller here (1.8 vs 1.9 Hz) because Fermilab lies at higher latitude than Daejeon ($\cos\lambda$).

**One week, then one month: the daily fingerprint.** Figure 20, two panels. With 80 Hz of signal-to-noise density and both tones stacked, two-sidereal-hour centroid bins reach $\pm 1.4$ Hz within the *first month* and $\pm 0.40$ Hz over the year: the 1.8 Hz sidereal wobble — out of statistical reach for the single-cavity camp of Part IX — emerges at $\sim 3\sigma$ within month one and stands as an $\sim 11\sigma$ measurement by year end, with the sidereal-versus-solar period test as the artifact killer. (The contrast with Figure 14 is aperture, not physics: five times the signal-to-noise density and the two $\gamma$ streams stacked.)

**One year: the annual curve as precision astronomy.** Figure 21. Fortnightly centroids carry statistical errors of $\pm 0.58$ Hz — the annual amplitude is measured at the few-per-mille level statistically, so the deliverable is still set by the $\pm 1$ Hz engineered floor of R8. The same fit returns $v_\odot$ 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 22 draws the three running observables against the four milestones of the run plan; the crossing near day 80 (live), where the ratio's statistics fall below its calibration, is the design point that sets R5.

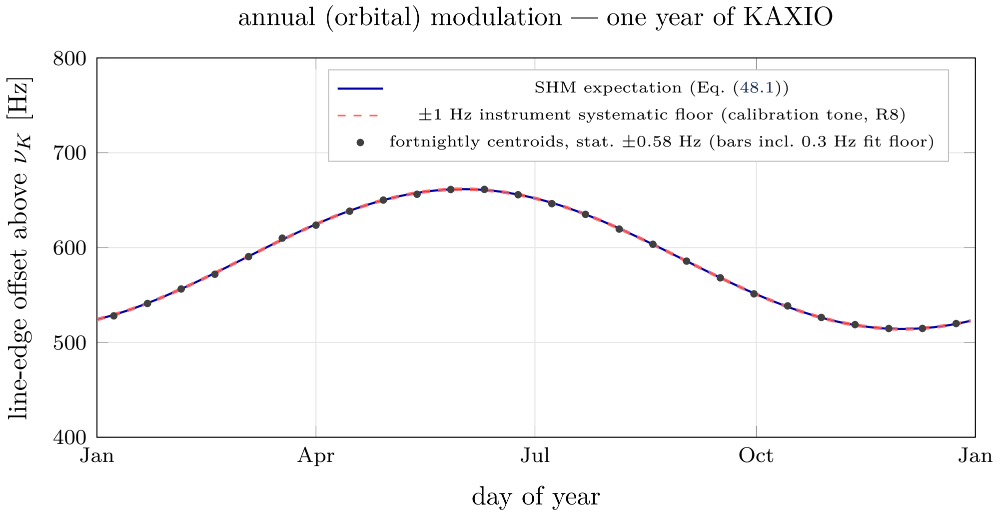


**Figure 21: One year of KAXIO: the annual modulation becomes precision halo astronomy** (same format as the annual panel of Figure 14, calendar year from January). Fortnightly centroid fits with both $\gamma$ streams (the two tones) stacked: the statistical error per point is ±0.58 Hz (Eq. (48.3) at 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 R8 (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 ±1.5% and $v_0$ to ±2%: a dark-matter *spectrometer*, not merely a detector.

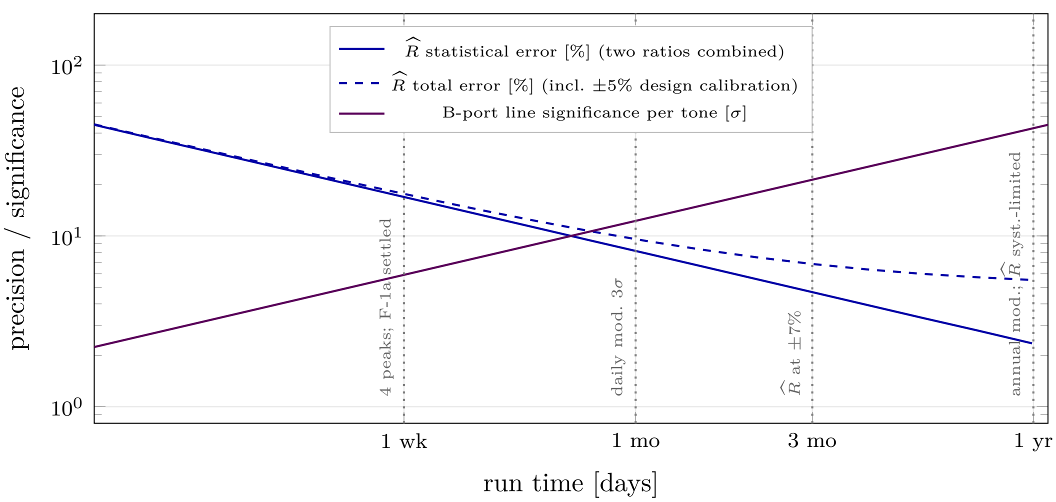


**Figure 22: The verdict clock.** The three running observables of the four-step plan (Table 34) against run time at 85% duty. Violet: the booster-port line significance per tone (design requirement $5\sigma$ at five live days). Blue: the statistical error on the channel ratio $\widehat{R}$, two independent ratios combined (solid), and the total including the ±5% design calibration (dashed) — statistics crosses below calibration near day 80 of live time, after which the measurement is systematics-limited at $\sim \pm 5\%$, comfortably resolving 25/4 from unity or from zero. Milestones (dotted): week one, all four peaks and the 1/35 verdict; month one, the daily modulation at $\sim 3\sigma$ (an $\sim 11\sigma$ measurement across the year); three months, the ratio at ±7% total; one year, the annual modulation measured against the ±1 Hz floor and every deliverable of Table 34 banked.

## 64 Staged coverage of the discrete window

The window is not a continuum to KAXIO; it is the discrete $\Xi_{\rm cap}$ 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 33 states the staged plan, and its columns deserve a walk-through. The first four give each rung's center frequency and its two tones: because the pair spacing is held in hardware, tuning to a rung means tuning *one* number, $\nu_0$, and the table's $\nu_\pm$ columns follow identically. The radius columns are the engineering content: the full ladder spans $\mathrm{TM}_{010}$ radii from 5.31 to 7.32 cm, a $-8\%/+23\%$ machining range about the benchmark pair that the tuning-rod design must accommodate — this single row-span is what the phrase "window coverage" costs in hardware. The schedule column implements the mark of Part V: benchmark first, then outward in the order of the cap-selection argument, so that the highest-prior rung is retired in week one and the ladder's tail carries the lowest stakes. Two engineering realities are priced rather than assumed away. At the largest rung ($\Xi_{\rm cap} = 35$, $r_- = 7.32$ cm) 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, recovery is a standard multi-day procedure, and the 85% duty requirement of R6 carries it as an explicit line item (Table 35).

One allocated week per rung deserves its justification, because the physics is nearly five orders of magnitude faster: at 80 Hz of signal-to-noise density the matched-line $5\sigma$ time is 8.7 s (Table 31). The week is not integration; it is discipline: retune and re-lock the pair, re-run the calibration-tone and injection suite, execute the full candidate protocol of §37.3 on anything that flickers, and bank enough dwell that the rung's null is a *publication-grade* exclusion rather than a logbook note. At that cadence the ladder is retired in seven calendar weeks; the seam sweep — months three to five — then closes the intervals between rungs at DFSZ depth with the tandem lock still engaged, so that even the theory's own $\pm 6\%$ dispersion around each rung is covered without ever un-pairing the tones. Every row is doublet-aware by construction, so each rung's null writes a completed F-Kaxion-1c entry for that rung, not merely a single-tone limit; the remainder of the year rides on the discovery tape (modulations, ratio) or deepens the exclusion on the null tape. Figure 23 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 thirty deep at every rung.

## 65 The four-step run plan

A month-zero commissioning shakedown — cooldown, tandem-lock closure, and calibration-tone and synthetic-doublet injection runs — precedes week one inside the schedule, so the plan's clock starts with a verified instrument. Table 34 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 four peaks and the 1/35 verdict — F-Kaxion-1a, the test that no single-tone instrument can even attempt, settled before the first month's cryogen

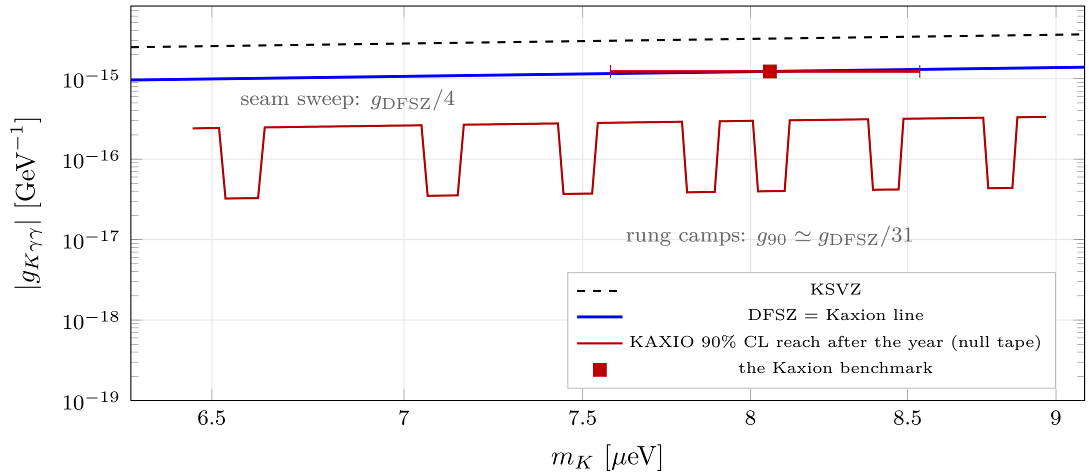


**Figure 23: The null-tape deliverable: the window, closed to a factor of thirty below the target.** The 90% CL coupling reach of the staged campaign of Table 33 (Eq. (52.1), 85% duty): one camped week per $\Xi_{\rm cap}$ rung drives $g_{90}$ to $g_{\rm DFSZ}/31$ at each rung (notches), and the seam sweep holds $g_{\rm DFSZ}/4$ across the entire window — so the exclusion does not merely graze the prediction, it buries it: the ±8% theory band on the coupling, and any conceivable sub-DFSZ variant reading of the sheet, dies together. No haloscope result to date has excluded any axion band even a factor of a few below DFSZ; on the null tape KAXIO reaches a factor $\sim 30$ across the one band where it matters.

**Table 33: The staged campaign across the $\Xi_{\rm cap}$ ladder** (rung values from Part V). Radii from the $\mathrm{TM}_{010}$ relation; one calendar week per rung includes retune, calibration, and the full candidate protocol; the seam sweep covers the intervals between rungs at DFSZ depth with the tandem lock engaged throughout.

| $\Xi_{\rm cap}$ | $\nu_0$ [GHz] | $\nu_+$ [GHz] | $\nu_-$ [GHz] | $r_+$ [cm] | $r_-$ [cm] | **schedule** |
|---|---|---|---|---|---|---|
| **43** | **1.95** | 1.978 | 1.922 | 5.80 | 5.97 | week 1 (benchmark) |
| 42 | 1.90 | 1.927 | 1.873 | 5.95 | 6.12 | week 2 |
| 45 | 2.04 | 2.069 | 2.011 | 5.54 | 5.70 | week 3 |
| 40 | 1.81 | 1.836 | 1.784 | 6.25 | 6.43 | week 4 |
| 47 | 2.13 | 2.160 | 2.100 | 5.31 | 5.46 | week 5 |
| 38 | 1.72 | 1.745 | 1.695 | 6.57 | 6.77 | week 6 |
| 35 | 1.59 | 1.613 | 1.567 | 7.11 | 7.32 | week 7 |
| seams | tandem sweep, 1.59–2.13 GHz at DFSZ, doublet-aware | | | | | months 3–5 |

top-up; month one adds the sidereal fingerprint at $\sim 3\sigma$ — an $\sim 11\sigma$ measurement by year end — and the first coarse ratio; month three brings F-Kaxion-1b to ±7% total — already resolving 25/4 from unity and from zero — with the systematics-limited ±5.5% banked by year end, the point at which the word "axion-like" becomes indefensible and the word "Kaxion" becomes mandatory; 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 thirty below DFSZ (Figure 23).

**Table 34: 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 | four peaks; 1/35 spacing at $10^{-4}$; F-1a passed; shape = halo thermometer | benchmark pair excluded at DFSZ, doublet-aware; camp moves one $\Xi_{\text{cap}}$ rung |
| month 1 | daily modulation $\sim 3\sigma$ (sidereal test); $\widehat{R}$ at $\pm 8\%$ stat. (two ratios combined) | four rungs excluded ($\Xi_{\text{cap}} = 43, 42, 45, 40$); the ladder continues |
| month 3 | $\widehat{R}$ at $\pm 7\%$ total, en route to the systematics-limited $\pm 5.5\%$ | ladder complete (all seven rungs, week 7); seam sweep under way |
| year 1 | annual modulation vs $\pm 1$ Hz floor; daily modulation $\sim 11\sigma$; $v_\odot$, $v_0$ measured; identification complete | **F-Kaxion-1c executed**: the full window closed; the Kaxion, and the series' dark-matter claim, retired |

The plan's logic is the paper's logic: the discovery case terminates in an *identification* — $K = 35$ read twice, the halo measured, every Part-VI lock closed — 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 with the doublet template in the trigger. Each milestone is also a natural reporting point to the agencies and to the community — four scheduled deliverables in twelve months, each publishable on either tape. On the existence question the year cannot end inconclusive: the $\gamma$ arrays alone either find the doublet or close the window at depth. The *identification* layer is honestly conditional — the 25/4 verdict is gated on the booster meeting its fenced requirement, and a shortfall defers the ratio, never the existence verdict. There is no month in which the collaboration is not producing physics.

## 66 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 all four streams jointly; the offset preserves every internal relation — the 69/71 spacing, the channel ratio, the modulation periods — while blinding the *sub-linewidth* absolute-frequency comparisons against the theory sheet (the line's placement within its bin and the modulation phases; a $\pm 200$ 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. A second seal is stated as a design requirement: the B-stream's relative gain enters the analysis through a sealed calibration scale factor, opened at the same box-openings, so the ratio pipeline cannot steer toward 25/4. *Injection campaigns:* hardware-injected synthetic doublets, at blind times and blind strengths, exercise the full chain including the coincidence trigger; the recovered efficiency updates $\epsilon_a$ exactly as in the field's current practice. *Fixed analysis:* the analysis pipeline is the published Part-IX code with the KAXIO parameter block (§53), 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 lock, so unblinding changes one number (the absolute frequency) and no machinery. Second, the injection campaigns do double duty: they measure the efficiency $\epsilon_a$ as in current practice, and — because injections are synthetic *doublets* at hardware level — they continuously verify the one subsystem no predecessor experiment has run, the four-stream coincidence trigger, against ground truth throughout the year. Governance is correspondingly simple: box-openings at the four milestones of Table 34, each preceded by a frozen-pipeline tag and followed by the twelve-month data release.

## 67 Risk register

Table 35 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.* The twin tandem-tuned $\gamma$ arrays alone — with the booster removed entirely — still execute F-Kaxion-1a in the trigger and F-Kaxion-1c across the window at full depth; every risk on the booster side therefore degrades the *identification* program gracefully rather than voiding the *discovery* program. This is why the two medium-likelihood booster rows (the sensitivity requirement and the magnon linewidth) sit at medium rather than high impact, and why the R&D year is structured as a gate: the one genuinely novel subsystem must prove its requirement on the bench before the build commits, and a measured shortfall converts the proposal to its $\gamma$-only fallback rather than sinking it.

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 $\pm 1$ 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 four-peak 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 36: 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 $2 \times 9$ twin families |
| tuning system | modify | tandem-lock control layer (software + piezo) |
| harmonic-booster module | new | side cryostat, shields, 2× magnon hybrids |
| 4-channel DAQ + cal-tone system | new | GPS clock, injection network |
| booster R&D (pre-construction) | new | 1 bench year; gates the build (R4/R5) |

**Table 35: Risk register.** L/M/H = low/medium/high. The controlling design principle: no single failure removes the instrument's core deliverable, because the $\gamma$-only tandem pair already executes F-Kaxion-1a and F-Kaxion-1c at full depth. A register is included because an instrument proposal without one is not a proposal.

| **Risk** | **Likelihood** | **Impact** | **Mitigation / fallback** |
|---|---|---|---|
| booster misses $\mathrm{SNR}_B$ requirement | M | M | R&D year gates the build; fallback: $\gamma$-only run still executes F-1a and F-1c in full; ratio deferred to upgrade |
| magnon linewidth above budget | M | M | more spheres ($\sqrt{N}$), lower-damping material lots, tighter bias homogeneity |
| photon-counter dark rate high | M | L | staged: first-gen JPA row of Table 30 already meets the baseline plan |
| tandem-lock mechanical cross-talk | L | M | independent piezo drives, software lock only; 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 |
| 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 |

## 68 Cost class, schedule, and organization

Table 36 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 (the −8%/+23% radius span of Table 33) and the tandem-lock control layer. The *new* rows are the only genuinely new procurement — a side cryostat with shields and coils, two magnon hybrids, a handful of JPA channels from an existing production line, and a commodity-FPGA DAQ layer — plus the bench R&D year that gates the build. 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: booster R&D and twin-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 11) and extends it to the full window through the tandem sweep. The critical path runs through the booster gate, not the cavities: the twin arrays are conventional work that can proceed in parallel with EFR's own construction, while the booster's bench year must conclude before the side-cryostat procurement — which is precisely why it is scheduled first and fenced hardest. The gate itself closes on three quantitative exit criteria, declared here: (i) bench-equivalent booster sensitivity $\mathrm{SNR}_B = 5\sqrt{t/5\,\mathrm{d}}$ demonstrated on an injected synthetic line at $B_e = 70$ mT; (ii) booster-to-array microwave isolation verified $> 100$ dB end to end; (iii) two-port calibration transfer demonstrated at $\pm 20\%$ or better on the bench pair. Meeting all three commits the build; missing any one invokes the $\gamma$-only fallback of Table 35.

Organisationally 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, booster R&D and construction, DAQ and blinding infrastructure, and the simulation/analysis pipeline already published with this paper. Each package has a natural single-institution owner class (machine shop and cryogenics for the cavity work, a university quantum-electronics group for the booster, 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 doublet-aware 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, the booster R&D deliverables, and the data-release terms of §66. This paper supplies the physics case and the pre-registered analysis; the schedule authority is EFR's.

## 69 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: $\mathrm{TM}_{010}$ 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; the magnon-cavity hybrid at the heart of the booster has been operated as an axion detector by the QUAX collaboration for years (Crescini et al., 2020). What is genuinely new — the tandem frequency lock, the four-stream coincidence trigger, the calibration-tone discipline — is control software and microwave plumbing, the least risky tier of experimental technology. The one open engineering question, the booster's sensitivity requirement, is deliberately quarantined in a one-year bench R&D program that gates the build (Table 35): the proposal cannot fail by surprise, only by a measured bench result — and even then the $\gamma$-only instrument delivers the discovery and the kill test in full.

**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), a side cryostat with shields and coils, two magnon hybrids, 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 36). The marginal cost of converting EFR's scan into KAXIO's four-peak 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. Framed as cost per unit of decisive information, we know of no cheaper experiment in fundamental physics.

**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 ever be pointed at a *prediction*, because no theory ever staked one. This paper stakes one. The frequency is named to $\pm 6\%$ before the run; the strength sits on the DFSZ line by theorem, not by choice; the line shape, the annual and daily breathing, the 1/35 partner, and the 25/4 channel ratio are published in advance with statistical and systematic errors printed on every figure, and the analysis code is public. 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 ambiguity has been engineered out in advance: the figures of Parts IX–X are the identikit, published before the arrest, and when real data draws these curves there will be nothing left to argue about — only something 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 (the identification layer rides on the booster gate, §67). Agencies are routinely asked to fund sensitivity; it is rare to be offered certainty. That is what one year of KAXIO buys.

**The case in three sentences.** The hardware exists, the methods are demonstrated, and the one novel subsystem is fenced behind a cheap R&D gate with a full-value fallback. 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 the first fully pre-registered dark-matter prediction in the fifty-year history of the search — 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.

## 70 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 two-lock 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

## 71 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.

### 71.1 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 $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$, with its two non-spatial coordinates — a scale interval $I_r$ and a phase circle $S^1_\theta$, both dimensionless (Theorems 0 and 2). GG-A3-BI6 proves that anomaly freedom of the forced gauge group $\mathrm{SU}(5) \times \mathrm{U}(1)_B \times \mathrm{U}(1)_{Qi}$ admits exactly one Green–Schwarz two-form $B_2$, of rank one, obeying the modified Bianchi identity $\mathrm{d}H_3 = \mathrm{tr}\, R^2 - \mathrm{Tr}\, F_5^2 - \frac{2}{5} F_B^2 - \frac{2}{5} F_{Qi}^2$. 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 light eigenstate of these two geared clocks, 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.

### 71.2 The numbers: every observable a gear ratio of one ledger

What converts "the Kaxion exists" into "the Kaxion is at 1.95 GHz" is that every quantity on its sheet is a gear ratio of the *same* five integers, $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, fixed in

GG-A1-Constitution–GG-A3-BI6 before any astronomical input. The mass obeys the exact invariant $m_K f_K = \sqrt{\chi_{\rm QCD}}$, legitimate because the gluonic portal (Theorem T-Kaxion-2) makes the Kaxion couple to the QCD topological charge with $N_{\rm DW} = 1$. The decay constant is fixed not by hand but by the geometry, $f_K = (7.1 \pm 0.5) \times 10^{11}$ GeV; the *derived* abundance $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2) = 0.2676$ of GG-A10-Cosmos then fixes the one remaining input, the $\mathcal{O}(1)$ misalignment 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 two fingerprints are pure integer arithmetic: the twin-peak split $\Delta\nu/\nu = 1/K = 1/35 = 2.857\%$ is the spectral memory of the throat depth, and the portal lock $g_{KBB}{:}g_{K\gamma\gamma} = 5{:}2$ is the inverted BI-6 trace ratio. 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.

### 71.3 One physical line, and the discipline that fixes it

A reader could worry that the "raw" and "mixed" frequencies 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, $m_K = (8.1 \pm 0.5)\,\mu$eV, $\nu_K = (1.95 \pm 0.12)$ GHz. The two discrete questions that could have multiplied the prediction — the cosmological history of the compact angle and the frame a detector reads — are both closed internally (Part VII, the uniqueness argument), leaving one window, $[1.59, 2.13]$ GHz, and a mandated silence everywhere else.

### 71.4 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, five theorems 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 $t \simeq 0.3\,\mu$s, with primordial velocities $\sim 10^{-17}$ and free-streaming at milliparsecs — fifteen 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 (§75.4). Stability is an output with thirty-two orders of headroom for the physical line. And the abundance is not fitted but derived, because the target $\Omega_{\rm DM}$ is itself derived. The Kaxion meets the dark-matter job description not by construction but by consequence.

### 71.5 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 now bracketed by funded instruments: IBS-CAPP from the window's body on its 1.8–10 GHz magnet, ADMX-EFR from its upper edge with an eighteen-cavity array, both reaching DFSZ depth across 2027–2030. On top of a discovery sit two identification layers no generic search is built to read: a dual-mode coincidence resonator that demands the 1/35 partner of every candidate, and a harmonic booster — the 70 mT electron-spin port first, the non-Larmor internal-field nucleon port as R&D — that reads the 25/4 matched-channel power ratio. A line in the window, split by 1/35, with its two channels at 25/4, would not merely be "an axion"; it would be a laboratory measurement of a gauged baryon symmetry and of the throat depth $K = 35$ of a six-dimensional geometry.

### 71.6 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 $\Delta\nu/\nu = 1/35$ split, the 5:2 portal lock, and the $\pi$-locked abundance target. *Benchmark-level:* the cap loading $\Xi_{\text{cap}} = 43$, hence the single decay constant and the absolute masses, fenced as the minimal OGN-5-compatible realization and bracketed by the discrete 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 — while the misalignment angle $\theta_i \simeq 1.4$ rad is an $\mathcal{O}(1)$ *output* of the abundance, not a tuned input. 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.

## 72 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. 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 $\mathrm{U}(1)_{\mathrm{PQ}}$ 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, no tunable abundance, 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) against the dark photon; a fixed GUT-adjacent inflationary spectrum against primordial-black-hole dominance (Part IV). Strikingly, the candidates the data now disfavor — fuzzy and warm dark matter bounded by JWST and Lyman-$\alpha$, 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 anywhere in the Kaxion sheet: one mass (with a ±6% Standard-Model input), one coupling, two locked integer ratios. 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, 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 line in the window, split by 1/35, with its two channels at 25/4, 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 frequency, its coupling, its doublet, and its channel ratio 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 two fingerprints are not incremental discriminants; they carry information across an extraordinary span of physics. A measured 5:2 (matched power 25/4) would be the first laboratory evidence for a *gauged* baryon symmetry — a discovery in its own right, independent of dark matter — because the ratio is a gauge-ledger entry that no Peccei–Quinn model has a gauge field to reproduce. A confirmed 1/35 doublet would be a direct readout of a compact dimension's integer throat depth. 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 fingerprints do.

**(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 anywhere impeaches the ledger everywhere, and a confirmation anywhere is collateral support for the rest. The Kaxion can be killed not only by a haloscope null 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 $\mu$eV pseudoscalar at 1.95 GHz carrying a 1/35 doublet and a 5:2 portal — reachable by instruments already funded. 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: both of its parents are periodic pseudoscalar directions of the compact geometry (§9), its lightness is the circle's own protection — a periodic direction admits no polynomial potential — and every quantized feature a haloscope would read — the 1/35 doublet split, the 5:2 portal, the alignment that retires strong CP — is a ledger integer on the fiber that remembers, which is why none of them is adjustable. Its *address* on the dial, by contrast, is the descent's bookkeeping: the decay constant closed from the $\pi$-locked inventory, the misalignment abundance accumulated along a one-way cosmic history, the $\mu$eV 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 $\mu$eV 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 7 as its portrait.

## 73 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 four numbers at once — two tones, two channels — 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, a doublet that reads the integer depth of a compact dimension, and a channel ratio that reads a gauged baryon symmetry. 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.

## 74 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 referee can check the negative — that case was made in the eight points of Section 72 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.

### 74.1 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 $\mu$eV 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 two years.*

**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 derivation: the decay constant is a geometric output ($f_K \simeq 7.1 \times 10^{11}$ GeV from $K = 35$ and the cap ladder), and with $m_K f_K = \sqrt{\chi_{\rm QCD}}$ exact, ten orders of magnitude become the half-octave $[1.59, 2.13]$ GHz with a marked benchmark bin (Part V). What was a landscape is now a target. *Resolved by derivation; the residual* $\pm 6\%$ *is Standard-Model-side and shrinks with the lattice.*

**PS-Kaxion-3: The quality problem.** The acknowledged open wound of every Peccei–Quinn construction: a global $\mathrm{U}(1)_{\rm PQ}$ 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 immune: its shift symmetry is not a global symmetry of any scalar potential but the gauge invariance of the two-form plus the large gauge invariance of the Wilson line, and gauge redundancies cannot be broken by gravitational operators (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 — $\rho_{\rm DM}/\rho_b \simeq 5.3$ — 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 $\pi$-locked inventory fixes the ratio as arithmetic, $\rho_{\rm DM}/\rho_b = (1/\pi - 1/2\pi^2)/(1/2\pi^2) = 2\pi - 1 \simeq 5.28$ (Part III), tested against the CMB inventory at the percent level. Honesty note, carried from Part III: the $\pi$-locked fractions are among the *weaker* members of the cohort's track record — $\pi/e$-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 $\theta_i^2$, 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 the *output*, not the input: the geometry fixes $f_K$, the $\pi$-locked target fixes the abundance, and the closure returns $\theta_i \simeq 1.4$ rad — the flattest, most generic value an angle can take, nowhere near the tuned edges (Section 15.3). The direction of inference is inverted, and the inversion is checkable: a confirmed line whose absolute rate implies a tuned angle would strain the closure visibly. *Resolved — the last free parameter of axion cosmology recast as a derived quantity.*

**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, and not by choice: with no radial order parameter anywhere in the spectrum, patch decorrelation has no mechanism — single-patch coherence is a corollary of the existence theorem, so no strings, no walls, no minicluster spectrum (Section 36.2) — and the two-epoch structure of the derived history clears the isocurvature bound thirty-eight orders deep, with nineteen orders of headroom in the angle-setting lock scale $H_2$ (Part III). *Dissolved by theorem; a confirmed Kaxion minicluster population would falsify 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 $5\sigma$ 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 single frequency pays neither: the trials factor collapses to one by arithmetic (§37.1), and the prediction was in print, with its rate, its line shape, and its two locked fingerprints, before the data. *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 line in the window, split by exactly 1/35, with its two channels at exactly 25/4, is consistent with nothing else on the market — the doublet reads the bulk's integer depth, the ratio reads the gauged portal, and both are immune to the halo uncertainties that blur absolute rates (Part VI). Detection and identification collapse into the same week. *Resolved by design — the organizing philosophy of the entire detection program.*

## 74.2 The structural innovations

The register of what did not exist before, each entry in before-and-after form:

**I-Kaxion-1 (the zero-parameter 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*: mass, coupling, abundance, misalignment angle, and cosmological history all derived, with the residual uncertainty entirely Standard-Model-side ($\pm 6\%$ from $\chi_{\text{QCD}}$ and the window). To our knowledge the first dark-matter candidate published with no continuous freedom anywhere in its chain.

**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 doublet fingerprint).** *Before*: haloscope signals characterized by frequency, power, and line shape — none of which identifies a model. *After*: a two-line signature whose fractional spacing 1/35 is a winding datum, radiatively protected to $10^{-3}$ of itself — the first proposed laboratory observable that reads a compact dimension's integer depth (T-Kaxion-2).

**I-Kaxion-5 (the matched-channel ratio).** *Before*: axion couplings to photons and matter treated as independent model parameters. *After*: a locked two-channel power ratio $P_B/P_\gamma = 25/4$ descending from the anomaly ledger, halo-independent by construction — the first proposed laboratory measurement of a *gauged* baryonic portal (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 single-bin campaign whose trials factor is exactly one, 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, doublet response, ratio convergence — from fixed seeds, with complete input tables, justified systematics, worked equations, a pipeline flowchart, and a public standard-library script an expert can re-run and dispute line by line (Part IX).

**I-Kaxion-8 (KAXIO, the four-peak observatory).** *Before*: haloscopes as single-channel scanning instruments, upgraded for speed. *After*: an identification instrument — two sub-arrays tandem-locked on the twin tones, an electron-magnon channel on each, four simultaneous peaks through one year — designed so that the run cannot end inconclusive on the existence question: every exit (discovery, partner-kill, ratio-kill, window-kill) is a verdict, with the ratio verdict gated on the booster's fenced requirement (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 funded scan fronts, and a window that closes on a stated schedule (∼2030–2032) 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, re-tune for the partner, add the spin port for the ratio, run the year for the modulations — each step costed, ordered by expense, and assigned its falsifier (Parts VII–X).

### 74.3 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 (§66). *Honest grading*: every quantitative claim carries its status in the master ledger, and the least flattering entries are displayed most prominently — the two named conjectures appear in the abstract's own register, the $\pi$-locked fractions are flagged as the track record's weaker members at every use, and the booster sensitivity is printed as a design requirement, not a demonstrated fact. *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 theory rather than moving the prediction. *Symmetric stakes*: the paper commits to what failure means — a swept and silent window does not merely retire a candidate, it impeaches the integer ledger that also carries the electroweak scale and the collider cluster — 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.

## 75 Future directions, open items, and recommendations

### 75.1 The experimental road to a verdict

The path from prediction to verdict is staged and, at every stage, decisive. *Now:* a doublet-aware re-binning of existing 1–2 GHz data tests the 1/35 lock at zero hardware cost. *2027–2030:* the funded DFSZ programs (IBS-CAPP on its 1.8–10 GHz magnet from the window's body, ADMX-EFR's 18-cavity array from its upper edge) scan $[1.59, 2.13]$ GHz at the depth the Kaxion requires — discovery, or closure of the window. *2030s:* the harmonic booster (the 70 mT electron port, then the internal-field nucleon port) reads the 25/4 channel ratio, turning a discovery into an identification of gauged baryon number; *KAXIO* (Part X), the four-peak observatory, integrates all of it — two tones, two channels, 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: exclusion of the window (F-Kaxion-1), a candidate line with no 1/35 partner (F-Kaxion-1a), a compact-angle-like line at any other frequency (F-Kaxion-1d), or a matched two-channel readout violating 25/4. Any one retires the Kaxion and, with it, the dark-matter claim of the series.

### 75.2 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 prefactor 0.090 at the $f_K/10^{12}$ GeV normalization, and GG-A10-Cosmos carries the identical pipeline at its misalignment equation (Section 25.5); the division of labor is single-decay-constant — the geometry fixes $f_K$, the observed abundance fixes 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 $\pi$-locked target, the invariant $m_K f_K = \sqrt{\chi_{\rm QCD}}$, 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: the $\pm 6\%$ ($\pm 0.5$ $\mu$eV) budget quoted throughout this paper is the compression of the relic pipeline onto the published lattice and transport tables (Part V, Table 12), and GG-A10-Cosmos carries the identical $\pm 6\%$ ($\pm 0.5$ $\mu$eV) budget; the two are one prediction at one level of QCD-input refinement, not two predictions.

### 75.3 The three named open items

Honesty about what would close the remaining gaps, in increasing order of difficulty. *(i) The misalignment angle — closed.* The angle requires no assumption: with the decay constant fixed by the geometry, the observed abundance fixes $\theta_i \simeq 1.4$ rad outright (Section 15.3), the generic $\mathcal{O}(1)$ value, and the single-patch history is forced within the series (no Peccei–Quinn scalar to randomize, reheating nearly nine orders too cold, Section 36.2).

This item is closed: the angle is an output, not an assumption. *(ii) The mass-uniqueness index.* $\Xi_{\rm cap} = 43$ is the minimal OGN-5-compatible benchmark fixing the single decay constant; a microscopic throat calculation of the cap mixing angle $\beta$ (Section 9.3) would replace the discrete window $[1.59, 2.13]$ GHz by a point. *(iii) The baryonic magnitude — now computed.* The 5:2 ratio is locked; 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 (§75.4), so the Kaxion predicts standard cuspy cold dark matter and the cusp–core problem is, for it, baryonic. What remains genuinely open is only the nucleon-*spin* chiral matrix element that sets the Part VIII booster's integration time — its order now pinned at $1/f_K$, its $O(1)$ coefficient awaiting a chiral/lattice computation.

### 75.4 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 $\rho_b \simeq 0.1\,M_\odot\,\mathrm{pc}^{-3}$ ($n_b \simeq 4\,\mathrm{cm}^{-3}$), a halo velocity $v \simeq 30\,\mathrm{km\,s}^{-1}$, and $t_H \simeq 1.4\times10^{10}$ yr,

$$\sigma_{\chi b}^{\mathrm{thr}} \;\simeq\; \frac{1}{n_b\,v\,t_H} \;\simeq\; 2\times10^{-25}\,\mathrm{cm}^2 \quad \text{per nucleon.} \tag{75.1}$$

**The portal cross-section.** The dark–baryon vertex descends from the gauged-U$(1)_B$ coupling $g_{KBB} = \frac{5}{2}\,g_{K\gamma\gamma} \propto 1/f_K$ (Section 27); the induced dark-matter–nucleon interaction is a derivative (hence momentum-suppressed) contact operator with scale $f_K$. Dimensional analysis gives

$$\sigma_{\chi b} \;\sim\; \frac{m_K^2\,v^2}{\pi\,f_K^4} \;\simeq\; 1\times10^{-112}\,\mathrm{cm}^2, \tag{75.2}$$

the double $1/f_K^2$ (with $f_K \simeq 7.1\times10^{11}$ GeV) and the momentum factor $(m_K v)^2$ (with $m_K = 8.1\,\mu$eV) together rendering it astronomically small. Even the most generous coherent enhancement — promoting the condensate's full de Broglie occupation $N = \rho_{\mathrm{DM}}\lambda_{\mathrm{dB}}^3/m_K \simeq 1.4\times10^{26}$ to a coherent multiplier and, over-generously, *squaring* it — lifts this only to $\sigma_{\chi b} \lesssim 1\times10^{-60}\,\mathrm{cm}^2$.

**The verdict.** The portal therefore falls short of the coring threshold (75.1) by

$$\frac{\sigma_{\chi b}^{\mathrm{portal}}}{\sigma_{\chi b}^{\mathrm{thr}}} \;\simeq\; 10^{-93}\ \text{(single-particle)} \quad \text{to} \quad \gtrsim 10^{-41}\ \text{(extreme coherent)}, \tag{75.3}$$

between forty and ninety orders of magnitude. No coherence bookkeeping closes a gap of that size, and the dominant suppression — $1/f_K^4$ with a Green–Schwarz decay constant near $10^{13}$ 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 haloes; 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 lock is untouched.* The matched-channel power ratio $P_B/P_\gamma = 25/4$ cancels the absolute magnitude (Theorem 32.1), so the harmonic-booster test of Part VIII stands regardless: it reads the *ratio*, not the coring strength. *(3) Consistency, not contradiction.* The same $1/f_K$ weakness explains why the booster's nucleon port is hard, why it is staged as R&D, and why it must be read on *resonance* (the hyperfine-enhanced port of Part VIII) rather than in a dwarf: a coupling far too feeble to move a galaxy can still be locked in ratio and read in a tuned laboratory. The one piece still genuinely open is the nucleon-spin chiral matrix element setting the booster's integration time, whose order is now pinned at $1/f_K$. (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.)

### 75.5 The path to few-percent

Section 29.2 located the entire $\pm 6\%$ in Standard-Model thermodynamics. The improvement program is concrete and external to the series: a dedicated lattice determination of $\chi(T)$ through the crossover at the percent level, plus an NLO misalignment transport, compresses the mass error to $\sim 3\%$ — at which point $\nu$ becomes a few-linewidth target and the Stage-II scan of Part VIII shortens by an order of magnitude. The same improvement sharpens nothing about the fingerprints, which were exact from the start.

### 75.6 Relation to the sister predictions

The phenomenology branches read as one geometry tested at three scales: P1-KK at $M_{\mathrm{lock}} \approx$ 3.85 TeV (the radial lock cluster, colliders), the present paper at $10^{-7}$–$10^{-5}$ 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 impeaches the ledger everywhere, 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 $\pm 6\%$ on the mass is on loan from lattice QCD and is being repaid on schedule; and the doublet and the portal — the tests that actually decide the matter — were exact from the first line and answer to no conjecture. 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.

# 76 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 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

## 77 Conclusions

### 77.1 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.

### 77.2 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, and nowhere else.

And it should arrive carrying two unmistakable signatures written into the geometry of space itself: the tone should come as a *pair* of lines separated by exactly one part in thirty-five, and it should speak in two channels — to light and to ordinary matter — in a fixed ratio of five to two. Not one of these numbers was adjusted to fit anything. Each is a consequence, read off 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.

### 77.3 What this paper has established

Stripped of narrative, the paper's content is ten numbers and the machinery that stands behind them — each with its status, its proof location, and the measurement that can kill it:

**Table 37: The paper in ten numbers.** Status grades follow the master ledger (Table 13): *theorem* (inherited or proved), *derived* (theorem + Standard-Model input), *benchmark* (marked row of a discrete window), *closure* (fixed by the observed abundance).

| | Quantity | Value | Status | Proved | Killed by |
|---|---|---|---|---|---|
| 1 | mass $m_K$ | $8.1 \pm 0.5\,\mu$eV | derived | Part V | F-1c/1d |
| 2 | frequency $\nu_K$ | $1.95 \pm 0.12$ GHz | derived | Part V | F-1c/1d |
| 3 | window | $[1.59, 2.13]$ GHz | discrete ladder | Part V | F-1c |
| 4 | twin split $\Delta\nu/\nu$ | $1/35 = 2.857\%$ | theorem (symm. limit) | Part VI | F-1a |
| 5 | channel ratio $P_B/P_\gamma$ | $25/4 = 6.25$ | theorem | Part VI | F-1b |
| 6 | coupling $g_{K\gamma\gamma}$ | $(1.2 \pm 0.1) \times 10^{-15}$ GeV$^{-1}$ | derived (DFSZ) | Part V | F-1c |
| 7 | decay constant $f_K$ | $(7.1 \pm 0.5) \times 10^{11}$ GeV | bench. $\Xi_{\rm cap}$=43 | Part V | (via 1–3) |
| 8 | misalignment angle $\theta_i$ | $\simeq 1.4$ rad | closure | Part V | (via 9) |
| 9 | abundance $\Omega_{\rm DM}$ | $1/\pi - 1/(2\pi^2) = 0.2676$ | derived (GG-A10-Cosmos, cond.) | Part III | CMB |
| 10 | lifetime $\tau_K$ | $\simeq 4 \times 10^{32}\, t_U$ | derived | Part V | (stability) |

Around the ten 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 collapses the look-elsewhere penalty to one (Part VII); the simulated discovery week with statistical and systematic errors printed in advance, and its public code (Part IX); and KAXIO, the four-peak 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.

### 77.4 An answer we can check — soon

What lifts this above a beautiful 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 funded experiments — IBS-CAPP in Korea and ADMX-EFR at Fermilab — are built to listen to, at exactly the faintness the theory demands, between now and roughly 2030.

A modest addition to those instruments — a receiver tuned to two pitches at once, and a second channel read by a tabletop magnet — would not only hear the tone but confirm its two fingerprints, 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, split in the predicted way, in the predicted ratio — 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 Kaxion is wrong, and the dark-matter claim of the whole framework falls with it. 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, and it is, honestly, a thrilling one.

### 77.5 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 depth of an unseen dimension and the imprint of a hidden symmetry. 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.

### 77.6 The objective limitations

Honesty, one last time, in one place. The relic pipeline rests on two named conjectures — misalignment completeness (the condensate is the sole saturator of the $\pi$-locked share) and compact-angle mass-uniqueness (the cap index selects the marked row) — flagged wherever they enter and carried in the master ledger; the frequency's $\pm 6\%$ and the coupling's $\pm 8\%$ are wholly Standard-Model-side and shrink with lattice QCD, not with argument. The booster (electron) port's sensitivity is an engineering requirement, fenced and gated behind a bench year, not a demonstrated capability, and the 25/4 readout additionally carries the fenced chiral input. 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.

### 77.7 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.

### 77.8 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 code that simulates its discovery is public; the dates by which it must be found or dead are printed. 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.

## 78 Part XII in plain English: why it matters whether this particular answer is right

If the argument of this paper holds, the dark matter is not a new species added to the inventory. It is a consequence of the same geometry that fixes the ordinary constants, doing the only thing it could do, and it sits at one frequency that a machine now under discussion could reach.

The claim is narrower than that summary sounds. Several of its steps are fenced by computations nobody has done, and the initial angle is fixed by the observation it is meant to explain rather than derived. The paper is better read as a sharp conjecture with a long chain of theorems behind it than as a finished derivation, and the ledger of the previous Part says exactly where the line falls.

What makes it worth acting on is not confidence but shape. A candidate with free parameters can absorb null results indefinitely; this one cannot. Scan the window at the stated depth and the question is answered, one way or the other, inside a decade. Very few proposals about the dark sector can say that.

And if it is found, the interesting thing will not be that the dark matter was a pseudoscalar. It will be that a framework built to explain why the fine-structure constant has the value it does turned out to predict, without being asked, what the universe is mostly made of.

# Acknowledgments

This work was developed through extensive iterative collaboration with three large language models: ChatGPT (OpenAI) for exploratory drafting, narrative development, and broad-scope ideation; Perplexity for literature synthesis, citation verification, and recent experimental cross-referencing; and Claude (Anthropic) for structural organization, mathematical consistency checking, LaTeX typesetting, and final manuscript preparation. The sequential refinement of the GG-Theory framework (GG-Constitution / GG-6 / BI-6 / GDOS / Gi-4 Pillars / Qi-4 Pillars) across hundreds of conversation sessions with these three systems was essential to the coherence of the present paper.

The author thanks Aaron Blaisdell and Hirotaka Iwasaki for stimulating discussions on the program's interdisciplinary scope and for critical feedback on the accessibility of the presentation. The author also thanks hundreds of collaborators in the Elegant Mind Club over the past decade, including Javier Carmona, Elizabeth Milles, Mingda He, Syed Bashir Hydari, Brian Ta, Chris Dao, Umaima Afifa, Isabella Bestanoby, Caominh Le, and Natsuko Yamaguchi.

All scientific concepts (especially the foundation of the GG-Constitution, GGEP/GG-6/BI-6, and OGGEP/GDOS/Gi-4/Qi-4 Pillars), theoretical judgments, physical interpretations, and any remaining errors are the sole responsibility of the author.

# Part XIII
# Technical Appendices

A graduate machine room in six blocks — the reference dictionaries and cross-cohort naming, the conceptual on-ramp (primers and the seven foundational questions), the from-scratch derivation courses behind every headline number, the detection-and-KAXIO course culminating in the four-peak template, and the registries, audits, and promotion roadmaps, 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 six 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: five one-paragraph primers, then a seven-question foundational essay answering the questions a reader actually asks.
- **Block C — The derivation courses.** The machine room — eight courses, each taking one result (the pseudoscalar from a two-form, the 1/35 split, the 8.1 $\mu$eV mass and photon coupling, the 5:2 baryon portal, the cap index $\Xi_{\rm cap} = 43$, the misalignment relic abundance, and the $\pi$-locked inventory) and deriving it step by step, with the Kaxion's numbers inserted at the end.
- **Block D — Detection and the KAXIO instrument.** The haloscope, transfer-function, four-peak 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.
- **Block F — 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.4) 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 major physics journals. The third column gives the *primary source* — the paper in the Arisaka-GG series, or "this paper," where the term is introduced. Figure 24 shows how the program-defined names nest, from the framework down to the Kaxion.

**Table 38:** 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 |
|---|---|---|
| *Framework name and core program objects* | | |
| 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 |
| 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 |

*continued on next page*

*Table 38 continued from previous page*

| **Symbol** | **Meaning** | **Primary source** |
|---|---|---|
| Kaxion | The boundary compact-angle pseudoscalar this paper identifies as cold dark matter | this paper |
| *Axioms and postulates* | | |
| 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 |
| *The runtime: GRIP and its parts* | | |
| 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, $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ (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 two-channel harmonic-oscillator readout (detector model) | this paper, Parts VII–VIII |
| *Named results of this paper* | | |
| T-Kaxion-1 | Existence theorem: the forced compact-angle pseudoscalar | this paper, Part II |
| T-Kaxion-2 | The gluonic portal, $m_K f_K = \sqrt{\chi_{\mathrm{QCD}}}$ | this paper, Part V |
| T-Kaxion-3 | The anomaly ratio $E/N = 8/3$ (the Kaxion on the DFSZ line) | this paper, Part V |

*continued on next page*

*Table 38 continued from previous page*

| Symbol | Meaning | Primary source |
|---|---|---|
| F-Kaxion-1 | The master falsifier (missing doublet, wrong channel ratio, window excluded, or wrong location) | this paper, Part VIII |
| *Standard physics abbreviations (plain Roman, journal convention)* | | |
| QCD, QED | Quantum Chromodynamics, Quantum Electrodynamics | — |
| EW, GUT, SM | Electroweak, Grand Unified Theory, Standard Model | — |
| SUSY | Supersymmetry | — |
| ΛCDM | 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 | — |

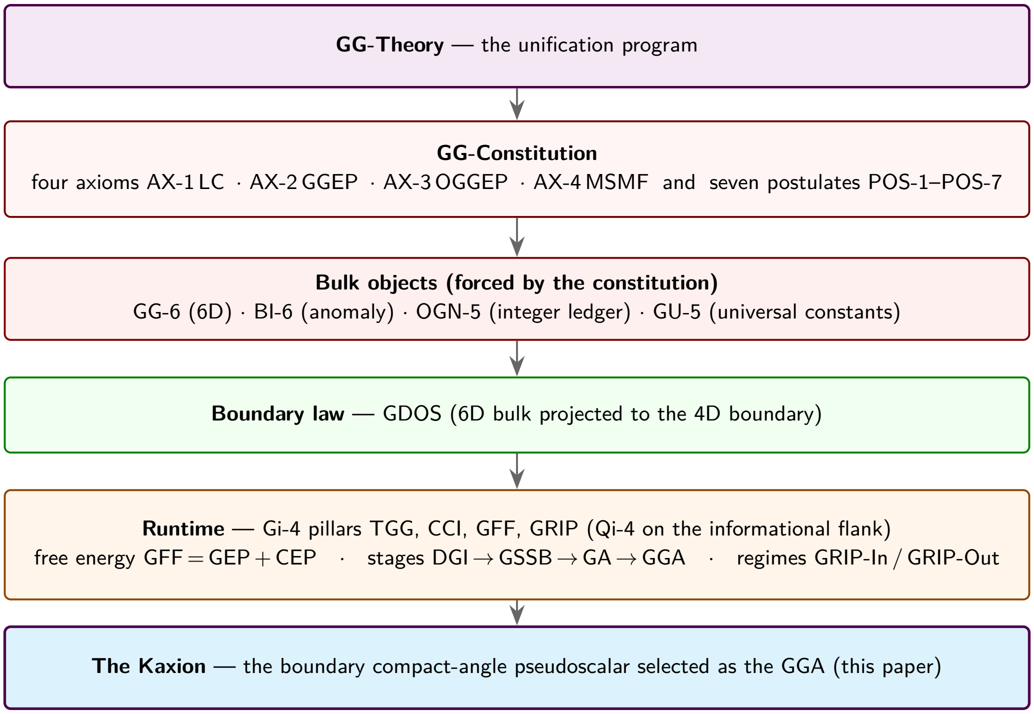


**Figure 24: Hierarchy of the program-defined names.** The framework GG-Theory and its charter GG-Constitution (axioms AX-1–AX-4 and postulates POS-1–POS-7) force the bulk objects GG-6, BI-6, OGN-5, and the constant-set GU-5; the bulk projects to the boundary law GDOS; the runtime GRIP (with free energy GFF = GEP + CEP and stages DGI → GSSB → GA → 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 38.

## A.2 Glossary of Key Symbols

Table 39 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 $F_{\mathsf{GFF}}$) follow the sans-serif convention of Appendix A.1.

**Table 39:** 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** |
|---|---|---|
| *Geometry of the bulk* | | |
| $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$ | 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, $r = \ln(E_{\rm GUT}/\mu)$ | Part II |

*continued on next page*

*Table 39 continued from previous page*

| Symbol | Meaning | Typical context |
|---|---|---|
| $\theta,\ S^1_\theta$ | 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 | BI-6; Part II |
| $\theta_{\rm GS},\ \theta_W$ | Tensor shadow (dual of $H_3$) and Wilson-line holonomy echo | Part II |
| *The Kaxion and its couplings* | | |
| $a_K$ | The Kaxion field (the light compact-angle eigenstate) | Parts II, V |
| $m_K$ | Kaxion mass (8.1 $\mu$eV) | Part V |
| $f_K$ | Kaxion decay constant ($\simeq 7.1\times10^{11}$ GeV) | Part V |
| $\nu_K$ | Laboratory frequency, 1.95 GHz | Parts V–VIII |
| $g_{K\gamma\gamma}$ | Kaxion–photon coupling (on the DFSZ line) | Part V |
| $g_{KBB}$ | Kaxion–baryon coupling ($g_{KBB}{:}g_{K\gamma\gamma} = 5{:}2$) | Parts V–VI |
| $\chi_{\rm QCD}$ | QCD topological susceptibility, $\chi_{\rm QCD}^{1/4} = 75.5$ MeV | Part V |
| $\theta_i$ | Initial misalignment angle, $\theta_i \simeq 1.4$ rad (output of the abundance) | Part III |
| $\Delta\nu/\nu = 1/K$ | Twin-peak split, $= 1/35 = 2.857\%$ | Part VI |
| $\Xi_{\rm cap}$ | Cap renormalization index ($\Xi_{\rm cap} = K+N+L = 43$, benchmark) | Part V |
| $\beta$ | Cap mixing angle relating the raw and laboratory frames | Part V |
| $\varepsilon = 1/K$ | Kinetic gearing of the two compact-angle clocks | Part II |
| *Integers, scales, and cosmology* | | |
| $(M,N,L;J,K)$ | 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 SU(5) ledger) | Part V |
| $M_{\rm lock}$ | Lock-cluster scale, $M_{\rm lock} = \frac{1}{2}E_{\rm Pl}e^{-K} \approx 3.85$ TeV | GG-A9-SM; GG-P1-KK |
| $E_{\rm GUT},\ E_{\rm Pl},\ M_{\rm Pl}$ | GUT, Planck energy, and Planck mass | Part V |
| $\Omega_m, \Omega_b, \Omega_{\rm DM}, \Omega_\Lambda$ | Matter, baryon, dark-matter, and dark-energy density fractions ($\pi$-locked) | GG-A10-Cosmos; Part III |
| $F_{\sf GFF} = F_{\sf GEP} + F_{\sf CEP}$ | The runtime free-energy functional and its two blocks | GG-A6-GRIP; Part II |
| $t_U,\ \tau$ | Age of the universe; Kaxion lifetime ($\tau \gg t_U$) | Part V |

## 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 $m_K f_K = \sqrt{\chi_{\rm QCD}}$, 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, **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 (§74.1).
- **O-Kaxion-N** for the $N$th Objection (Part XIV) and **I-Kaxion-N** for the $N$th structural Innovation (§74.2).
- **F-Kaxion-N** for the $N$th Falsifier: F-Kaxion-1a (missing 1/35 partner), 1b (matched channel ratio $\neq 25/4$), 1c (DFSZ-depth exclusion of the unique window), 1d (a confirmed line outside the window).

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 40 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 40:** 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; "C" = conjecture-fenced (one named computation from theorem).

| Quantity | Symbol | Value | Status | Derived in |
|---|---|---|---|---|
| Kaxion mass | $m_K$ | $8.1 \pm 0.5\,\mu$eV | T (given $f_K$) | App. C.4 |
| Laboratory frequency | $\nu_K$ | 1.95 GHz | T | App. C.4 |
| Twin-peak split | $\Delta\nu/\nu$ | $1/K = 1/35 = 2.857\%$ | T* | App. C.2 |
| Split in frequency | $\Delta\nu$ | 56 MHz | T | App. C.2 |
| Photon coupling | $g_{K\gamma\gamma}$ | $1.2 \times 10^{-15}$ GeV$^{-1}$ | T (DFSZ line) | App. C.4 |
| Anomaly ratio | $E/N$ | 8/3 | T | App. C.4 |
| Baryon-to-photon ratio | $g_{KBB}/g_{K\gamma\gamma}$ | 5/2 | T | App. C.5 |
| Matched channel power ratio | $P_B/P_\gamma$ | $25/4 = 6.25$ | T | App. C.5 |
| Decay constant | $f_K$ | $(7.1 \pm 0.5) \times 10^{11}$ GeV | B | App. C.7 |
| Cap index | $\Xi_{\rm cap}$ | $K + N + L = 43$ | C | App. C.6 |
| Discrete mass window | $\nu_K$ | $[1.59, 2.13]$ GHz (7 values) | T | App. C.6 |
| Misalignment angle | $\theta_i$ | $\simeq 1.4$ rad | B | App. C.7 |
| Dark-matter fraction | $\Omega_{\rm DM}$ | $1/\pi - 1/(2\pi^2) = 0.2676$ | C | App. C.8 |
| Matter fraction | $\Omega_m$ | $1/\pi = 0.3183$ | C | App. C.8 |
| Baryon fraction | $\Omega_b$ | $1/(2\pi^2) = 0.0507$ | C | App. C.8 |
| Lifetime | $\tau_K$ | $1.7 \times 10^{50}$ s | T | App. E.1 |
| Lock scale (cousin) | $M_{\rm lock}$ | 3.85 TeV | T | Arisaka (2026, P1-KK) |

* The *existence* of the doublet and its scale-free character are theorem-level; the precise value $1/K$ additionally assumes the throat exchange symmetry (equal QCD curvature on the two clocks), flagged in Appendix C.2 and in the ledger.

# B The conceptual on-ramp: primers and foundational questions

Before any formalism, this block answers the questions a newcomer actually asks. The five primers give the one-paragraph intuition for each headline fact; the seven-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 referee brings to a report. 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 and a Wilson line, and their light combination *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 fixes its decay constant from the geometry — so where an axion model has a dial, the Kaxion has a 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 $8.1\,\mu$eV?

Three numbers meet, and none of them is adjustable. First, the geometry fixes the decay constant: the scale throat of the six-dimensional bulk is $K = 35$ deep, its cap carries the discrete index $\Xi_{\text{cap}} = K+N+L = 43$, and unwinding those integers through the flux-cell measure gives $f_K \simeq 7.1\times10^{11}$ GeV — no continuous choice enters (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 $\chi_{\text{QCD}}$ is a measured, lattice-computed property of ordinary matter — $(75.5\pm0.5\ \text{MeV})^4$, nature's own dial reading (Appendix C.4). Third, the cosmic abundance closes the loop: with $f_K$ fixed, demanding that the relic condensate saturate the $\pi$-locked share $\Omega_{\text{DM}} = 1/\pi - 1/(2\pi^2)$ pins the one remaining input, the misalignment angle, at its generic value (Section 15.3). Divide $\sqrt{\chi_{\text{QCD}}}$ by $f_K$ and the answer is $8.1\,\mu$eV — a microwave tone at 1.95 GHz. Every per-cent of quoted uncertainty is Standard-Model-side; the integers contribute none.

## B.3 Primer: Why is the split exactly $1/35$ — and why can it not renormalize away?

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 C.2). 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 $(5/8)K^{-2} \simeq 5\times10^{-4}$, some fifty times below the splitting itself (Section 31). A measured pair at any other spacing, or a missing partner, is therefore not a model tweak but a falsification (F-Kaxion-1a).

## B.4 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 — is there a partner tone at exactly 2.857% below? A single-mode cavity is blind to it (the tuned mode suppresses the partner by seventy-five decibels), so this is a deliberate act: retune, or read the second mode. If the partner is missing at the one-per-cent level, the Kaxion is dead that day. 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: a spin-based receiver joins, and the ratio of the two channels converges on 25/4. At that point the logbook does not say "an axion-like particle"; it says the Kaxion, with the throat depth $K = 35$ read twice by two different machines. Part IX simulates this entire week with errors printed on every panel; Part X designs the observatory that runs it as one program.

## B.5 Primer: If the window is excluded, why can the frequency not simply be adjusted?

Because there is nothing 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 five integers and a measured QCD number: $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ fix the throat and the cap index, the cap index fixes $f_K$ within a seven-value discrete window, and $\sqrt{\chi_{\rm QCD}}/f_K$ fixes the frequency inside $[1.59, 2.13]$ GHz — with the benchmark row marked and argued (Section 26.3). To move the line outside that window one would have to change an integer, and the same 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 $\pi$-locks set the cosmic inventory in GG-A10-Cosmos. Retuning the Kaxion would silently refit the Standard Model and the universe's composition — which is to say, it cannot be done. That is what F-Kaxion-1c and F-Kaxion-1d mean operationally: a swept-and-silent window, or a confirmed line anywhere else, ends the matter (Appendix E.2).

## B.6 Seven foundational questions about the Kaxion

*An extended tutorial in seven questions.* The five primers above answer, in a paragraph each, the questions a newcomer asks first. This appendix answers the seven questions a newcomer asks *next* — the ones a graduate student brings to a seminar and a referee brings to a report — 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.6.1 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

$$\mathcal{L}_{a\gamma\gamma} = \tfrac{1}{4}\, g_{K\gamma\gamma}\, a\, F\widetilde{F}, \qquad g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K}\left(\frac{E}{N} - 1.92\right). \tag{B.1}$$

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 $\pi^0$ and $\eta$ 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 $1/f_K$,

$$m_K f_K = \sqrt{\chi_{\rm QCD}}, \qquad \chi_{\rm QCD}^{1/4} \simeq 75.5 \text{ MeV (lattice (Borsanyi et al., 2016))}, \tag{B.2}$$

one may eliminate the decay constant and write the coupling as a function of the mass alone,

$$g_{K\gamma\gamma} = \frac{\alpha\, m_K}{2\pi\sqrt{\chi_{\rm QCD}}}\left(\frac{E}{N} - 1.92\right). \tag{B.3}$$

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, $\chi_{\rm QCD}^{1/4} = 75.5(5)$ 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 $m_K = 8.1\,\mu$eV, so $f_K = 10^{12}\text{ GeV} \times (5.70/8.1) = 7.0 \times 10^{11}$ GeV,

$$g_{K\gamma\gamma}^{\rm DFSZ} = \frac{\alpha}{2\pi f_K}\,(0.747) = \frac{(1.16\times10^{-3})(0.747)}{7.0\times10^{11}\,\text{GeV}} \simeq 1.2\times10^{-15}\ \text{GeV}^{-1}, \tag{B.4}$$

exactly the value of the main text (§27; Appendix C.4). The dominant uncertainty on the line's *height* at fixed $E/N$ is the $\pm0.04$ on 1.92: because $8/3 - 1.92 = 0.747$, that error is $\pm5\%$ of the bracket — which is why the paper quotes $\pm8\%$ on $g_{K\gamma\gamma}$, the chiral bracket dominating. The geometry contributes *zero* continuous error.

**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 $\bar{\theta}_{\rm QCD} \to 0$ is by coupling to the gluon topological density $G\widetilde{G}$ — 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 $G\widetilde{G}$ coupling with unit color anomaly ($N=1$, $N_{\rm DW}=1$; Appendix C.4). *Second, QCD then sets its mass.* Once it couples to $G\widetilde{G}$, QCD instantons tilt its potential, and the curvature of that tilt is the susceptibility: $m_K = \sqrt{\chi_{\rm QCD}}/f_K$. The geometry supplies $f_K$; QCD supplies $\sqrt{\chi_{\rm QCD}}$.

**The picture in one line.** The compact non-spatial angle is a dial; to cancel the QCD vacuum angle it must be the knob $\theta_{\rm QCD}$ 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 $\pi^0$ 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.6.2 Why does the Kaxion appear as twin peaks split by exactly 1/35?

***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 $a_{\rm GS}$; the holonomy of the gauge field around the phase circle $S^1_\theta$ gives a second, $a_\theta$. 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 ($a_{\rm GS}, a_\theta$) 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 $G\widetilde{G}$, but through its own geometric route — $a_{\rm GS}$ via the BI-6 inflow, $a_\theta$ via the Wilson-line holonomy of the color subgroup of SU(5). A throat reflection symmetry exchanging $a_{\rm GS} \leftrightarrow a_\theta$

forbids the off-diagonal potential term, so the QCD-induced curvature is diagonal and equal,

$$V = \tfrac{1}{2}\mu^2\,(a_{\rm GS}^2 + a_\theta^2), \qquad \mu^2 = \frac{\chi_{\rm QCD}}{f_K^2}. \tag{B.5}$$

Contrast the generic case: if both coupled to QCD through the same combination $(a_{\rm GS} + a_\theta)$, 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,

$$\mathcal{L}_{\rm kin} = \tfrac{1}{2}\Big[(\partial a_{\rm GS})^2 + (\partial a_\theta)^2 + 2\,\epsilon\,\partial a_{\rm GS}\!\cdot\!\partial a_\theta\Big], \qquad \epsilon = \frac{1}{K} = \frac{1}{35}, \tag{B.6}$$

where $\epsilon$ is the dimensionless kinetic gearing — the normalized single-transit overlap of the two profiles across a throat of depth $K$ (computed geometrically in Appendix C.2). Because $V \propto \mathbb{1}$ and the metric is $\left(\begin{smallmatrix}1 & \epsilon\\ \epsilon & 1\end{smallmatrix}\right)$, the normal modes are the symmetric and antisymmetric combinations $u_\pm = (a_{\rm GS} \pm a_\theta)/\sqrt{2}$, with kinetic coefficients $1 \pm \epsilon$. Canonical normalization ($\hat{a}_\pm = \sqrt{1 \pm \epsilon}\,u_\pm$) gives the physical masses

$$m_\pm^2 = \frac{\mu^2}{1 \pm \epsilon} \quad \Longrightarrow \quad \nu_\pm = \frac{\nu_0}{\sqrt{1 \pm \epsilon}}, \qquad \nu_0 = \frac{\mu}{2\pi}, \tag{B.7}$$

and, to leading order in the small gearing,

$$\frac{\Delta\nu}{\nu} = (1-\epsilon)^{-1/2} - (1+\epsilon)^{-1/2} = \epsilon + \mathcal{O}(\epsilon^3) = \frac{1}{K} = \frac{1}{35} = 2.857\%, \tag{B.8}$$

i.e. $\Delta\nu = 1.95$ GHz$/35 \simeq 56$ MHz. Two features make this a clean fingerprint. It is a pure frequency *ratio*, hence scale-free: it depends on $\epsilon$ alone, not on $\mu$ or $f_K$, so even if the absolute mass shifts within the cap uncertainty (§B.6.4), 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 (§9.5).

**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

$$\underbrace{\nu/Q_L}_{\text{receiver, }\sim\text{kHz}} \ll \underbrace{\nu/Q_a \simeq 2\ \text{kHz}}_{\text{condensate line}} \ll \underbrace{\nu/35 \simeq 56\ \text{MHz}}_{\text{doublet split}} \tag{B.9}$$

means the two tones are spectacularly resolved (separated by $\sim 3\times10^4$ 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 $\kappa$ is thus a free parameter, and the search of Part VIII is designed to be efficient even for $\kappa$ 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.

### B.6.3 Why does the Kaxion couple to baryons, and why at exactly 5:2?

***Why does the Kaxion couple to baryons at all, unlike a QCD/PQ axion? And why is the baryonic coupling exactly 5/2 times the photon coupling?***

The difference is upstream, in the gauge group. GG-6's bulk group is forced to be

$$G_{\text{Bulk}} = \text{Spin}(1,5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}, \tag{B.10}$$

and the decisive factor is $\text{U}(1)_B$: *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 $p \to e^+\pi^0$).

**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,

$$S_{\text{GS}} = \int_{\mathcal{M}_6} B_2 \wedge X_4, \qquad X_4 \text{ with locked coefficient vector } \left(1, -1, -\tfrac{2}{5}, -\tfrac{2}{5}\right) \tag{B.11}$$

over (gravity, SU(5), $\text{U}(1)_B$, $\text{U}(1)_{Qi}$). The Kaxion is the boundary dual of $B_2$, so it inherits a coupling to every factor $B_2$ talks to — including the gauged $\text{U}(1)_B$. That is the baryonic portal $\frac{g_{KBB}}{4}\, a_K F_B \widetilde{F}_B$. 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

$$Y_0 = \mathrm{diag}\big(-\tfrac{1}{3}, -\tfrac{1}{3}, -\tfrac{1}{3}, \tfrac{1}{2}, \tfrac{1}{2}\big), \qquad B_0 = \mathrm{diag}\big(\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{1}{3}, 0, 0\big); \tag{B.12}$$

baryon number lives only on the three colored (quark) entries, hypercharge on all five. Their trace-norm ratio is a two-line computation,

$$\frac{\mathrm{Tr}(B_0^2)}{\mathrm{Tr}(Y_0^2)} = \frac{3 \times (1/3)^2}{3 \times (1/3)^2 + 2 \times (1/2)^2} = \frac{1/3}{5/6} = \frac{2}{5} = \frac{L - N}{L}, \tag{B.13}$$

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.

**The inversion: source 2:5 becomes response 5:2.** This is the step that makes *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 $H_3 \to \mathrm{d}a_K$ 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

$$\frac{g_{KBB}}{g_{K\gamma\gamma}} = \frac{5}{2}. \tag{B.14}$$

*(iv)* A matched receiver reads power — the square — so

$$\frac{P_B}{P_\gamma} = \left(\frac{g_{KBB}}{g_{K\gamma\gamma}}\right)^2 = \frac{25}{4} = 6.25. \tag{B.15}$$

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.

**Honestly fenced.** The 5:2 is theorem-level exact at the anomaly/gauge level (§27). Two caveats stay flagged: the chiral $-1.92$ dressing modifies the photon coupling's *absolute* strength but not this *ratio* (it cancels); and converting the gauge-level ratio to an actual nucleon-*spin* (NMR/booster) reading requires a chiral matrix element — the fenced open item of Part VIII. The ratio is locked; the spin-port translation carries a benchmark-level input.

**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), $\mathrm{U}(1)_B$ is one of the entries. The Kaxion is the bookkeeper's phase, so it must talk to whatever the ledger records — nucleons obligatorily. And a bookkeeper who writes an entry faintly is read loudly by the field dual to him (source and response are reciprocal), so baryon number — written as 2, read as 5 — is the channel the Kaxion shouts to while it whispers to photons.

#### B.6.4 What is the cap index $\Xi_{\rm cap}$, and why is it $K + N + L = 43$?

***What is $\Xi_{\rm cap}$, 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 $\Xi_{\rm cap}$ 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 $\Xi_{\rm cap}$ is.** Not a new parameter — a wavefunction-renormalization integer relating two normalizations of the *same* field,

$$f_K = \frac{f_{\rm bulk}}{\Xi_{\rm cap}}, \qquad m_K = \frac{\sqrt{\chi_{\rm QCD}}}{f_K} \;\propto\; \Xi_{\rm cap}. \tag{B.16}$$

Here $f_{\rm bulk}$ 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 $m_K \propto \Xi_{\rm cap}$, it sets the *absolute* mass — and nothing else. Crucially it acts identically on both doublet members and both portals, so it moves only *where* the line sits, never the $1/35$ split or the 5:2 lock. The fingerprints are theorem-level; only the address is conjectural.

**Where the factor comes from — the throat is a gearbox.** On the canonical metric $ds_6^2 = e^{-2r}\eta\, dx^2 + dr^2 + d\theta^2$, $r \in [0, K]$, the two clocks acquire wildly asymmetric kinetic norms,

$$Z_{\rm GS} = 2\pi \int_0^K e^{2r} {\rm d}r = \pi(e^{2K} - 1) \simeq \pi e^{2K}, \qquad Z_\theta = 2\pi \int_0^K e^{-2r} {\rm d}r \simeq \pi, \tag{B.17}$$

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 $\Xi_{\rm cap}$ is that map expressed in the throat's own quantized units. The Green–Schwarz clock is normalized after the warp has run through its full $K$ e-folds; the boundary clock a detector reads sits at the throat's mouth; converting one rate to the other multiplies by the throat depth in integer units, plus whatever the cap adds.

**The derived skeleton — this part is not arbitrary.** Three facts are provable from the canonical metric, the OGN-5 integers, and integer flux quantization alone (the referee-checkable Proposition of §26.3):

1. $\Xi_{\text{cap}} > 1$ strictly, by $\sim 30$ orders, forced by the exponential norm asymmetry (B.17). (This bounds it away from 1; it does *not* set the value.)

2. $\Xi_{\text{cap}} = K + \Sigma$, an *additive integer* form. The throat is $K = 35$ flux units deep, and the orbifold cap terminating its ultraviolet end must carry an integer number of extra flux quanta — because a compact fixed-point cap cannot end a fractional flux line. This is the load-bearing physical fact: it makes the residual an integer, not a continuous knob.

3. $\Sigma$ 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

$$\begin{aligned}\Sigma \in \{0, 3, 5, 7, 8, 10, 12\} \;&\Rightarrow\; \Xi_{\text{cap}} \in \{35, 38, 40, 42, 43, 45, 47\} \\ &\Rightarrow\; \nu_K = 1.95\,(\Xi_{\text{cap}}/43) \in [1.59, 2.13]\ \text{GHz}.\end{aligned} \tag{B.18}$$

So the absolute mass is not a free parameter with an error bar — it is one of *seven discrete values* (Table 11). That discreteness is genuinely derived, and it is the experimentally decisive content: the prediction lives in a finite list of frequencies before any conjecture is invoked.

**Why 43 specifically — the flux-bookkeeping argument.** The whole conjecture reduces to one choice: which subset $\Sigma$? The benchmark $\Sigma = N + L$ is not "add three nice integers"; it is a UV/IR sorting of the ledger, forced by how these same integers function everywhere else in the series. The family-projector flux ($N$) and the gauge-rank flux ($L$) encode the SU(5) → 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)):

$$\Sigma = N + L \quad\Longrightarrow\quad \Xi_{\text{cap}} = K + N + L = 35 + 3 + 5 = 43. \tag{B.19}$$

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, not yet on the integral; what would settle it is one finite computation — the $X_4$ flux integral of GG-A3-BI6 restricted to the cap's orbifold fixed-point divisors — and until it returns $N + L$ (rather than, say, $N + J$) the assignment is a flux-bookkeeping proposition, 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 proves $\Xi_{\text{cap}} \gg 1$ but 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.** Because $\Xi_{\rm cap}$ is built by counting quantized flux, the uncertainty is a *list of seven frequencies*, disposable two independent ways. *By theory:* the one-page cap-cohomology integral returns $\Sigma$ and collapses the window to a point. *By experiment:* since $\nu_K = 1.95\,(\Xi_{\rm cap}/43)$ GHz, the discovery frequency *measures* $\Xi_{\rm cap}$ *outright* — the line's position within the window *is* the value of the integer. A tuned parameter hides forever inside its error bar; this integer is read off the dial the moment the bell is heard.

**The picture in one line.** $\Xi_{\rm cap}$ is the integer gear ratio $f_{\rm bulk}/f_K$ — the number of teeth on the throat's gearbox; flux quantization forces it into the form $K+\Sigma$ with $\Sigma$ a subset-sum of $\{N, L, J\}$, so the mass is one of seven discrete frequencies; 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.6.5 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 has turned into an output.

**Why so small.** Everything rests on the pseudo-Nambu–Goldstone relation

$$m_K f_K = \sqrt{\chi_{\rm QCD}} \quad \Longrightarrow \quad m_K = \frac{\sqrt{\chi_{\rm QCD}}}{f_K}. \tag{B.20}$$

The numerator is a *hadronic* scale, $\sqrt{\chi_{\rm QCD}} = (0.0755\ \text{GeV})^2 = 5.7\times10^{-3}\ \text{GeV}^2$; the denominator is a *geometric, GUT-adjacent* scale, $f_K \simeq 7\times10^{11}$ GeV. Their ratio is

$$m_K = \frac{5.7\times10^{-3}\ \text{GeV}^2}{7\times10^{11}\ \text{GeV}} \simeq 8\times10^{-15}\ \text{GeV} = 8\,\mu\text{eV}. \tag{B.21}$$

Compare the electron: $m_e = y_e v/\sqrt{2}$ with $v = 246$ GeV and Yukawa $y_e \simeq 3\times10^{-6}$, giving 0.511 MeV, so $m_e/m_K \sim 6\times10^{10}$ — 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" (§9.5). That feeble tilt, $\sqrt{\chi_{\rm QCD}}$, 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 $m_K = \sqrt{\chi_{\rm QCD}}/f_K$ with $f_K = f_{\rm bulk}/\Xi_{\rm cap}$, and look at where each factor comes from. The susceptibility $\sqrt{\chi_{\rm QCD}}$ is lattice QCD ($\chi^{1/4} = 75.5(5)$ MeV) plus $\chi(T), g_*(T)$ transport — the *entire* $\pm 6\%$. The geometric factors $f_{\rm bulk}$ and $\Xi_{\rm cap} = 43$ are integers and a throat count — *zero continuous error*. The geometry contributes nothing continuous here; every percent of the error is QCD's, not the theory's. That is why the prediction is $8.1 \pm 0.5\,\mu$eV with the error traceable entirely to QCD thermodynamics, and why it tightens automatically as lattice QCD improves.

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 whole advance is to compute the one number the axion program leaves free, $f_K = f_{\rm bulk}/\Xi_{\rm cap}$, a geometric integer relation; supply it and the same precise QCD relation collapses the band to a point (up to the discrete $\Xi_{\rm cap}$ window of §B.6.4). 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 $\sim 100$ 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 computed by the geometry (hence a point).

**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 $\mu$eV 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 — computes the circumference as an integer count of the throat.

### B.6.6 Why is the cosmic inventory locked to $\pi$, and what becomes of the coincidence problem?

***Why are the cosmic fractions given so precisely** — $\Omega_m = 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$**? And "why now": what of the cosmic coincidence problem?***

By the paper's own grading these are the *weakest* entries in the program — the "$\pi/e$-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

$$\Omega_m = \frac{1}{\pi} = 0.3183, \quad \Omega_b = \frac{1}{2\pi^2} = 0.0507, \quad \Omega_{\mathrm{DM}} = \frac{1}{\pi} - \frac{1}{2\pi^2} = 0.2676, \quad \Omega_\Lambda = 1 - \frac{1}{\pi} = 0.6817, \tag{B.22}$$

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 lemma: the *Hopf-fiber identification.* 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 OGGEP/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 $S^1 \hookrightarrow S^3 \to S^2$, and the inventory is then volumes of the unit 3-sphere and its Hopf pieces. With the unit-sphere volumes $\mathrm{Vol}(S^1) = 2\pi$, $\mathrm{Vol}(S^2) = 4\pi$, $\mathrm{Vol}(S^3) = 2\pi^2$, spread the total (critical) density uniformly over $S^3$ and assign matter to the fiber, dark energy to the base:

$$\begin{aligned} \Omega_m &= \frac{\mathrm{Vol}(S^1)}{\mathrm{Vol}(S^3)} = \frac{2\pi}{2\pi^2} = \frac{1}{\pi}, & \Omega_\Lambda &= 1 - \frac{1}{\pi}, \\ \Omega_b &= \frac{1}{\mathrm{Vol}(S^3)} = \frac{1}{2\pi^2} \text{ (baryons = one arc unit)}, & \Omega_{\mathrm{DM}} &= \frac{2\pi - 1}{2\pi^2} = \frac{1}{\pi} - \frac{1}{2\pi^2}. \end{aligned} \tag{B.23}$$

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" $\pi$ 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 $\pi$-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,

$$\frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} \text{ (baryon share of matter)}, \qquad \frac{\rho_{\mathrm{DM}}}{\rho_b} = 2\pi - 1 \simeq 5.28, \tag{B.24}$$

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 $\pi$-locks: on a circle of circumference $2\pi$, 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 13.9 Gyr old today versus 13.8 Gyr at the crossing, "now" sits only $\sim$ 75 Myr — a few percent — past it. So the honest scorecard: the 120-order fine-tuning of $\Lambda$ 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.6.7 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 $\theta_i \simeq 1.4$, spatially uniform on super-horizon scales; when the expansion slows to $3H \simeq m_K$ (onset at $T \simeq 1$ GeV, $t \simeq 0.3\,\mu$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

$$\rho_K = \tfrac{1}{2} m_K^2 a_K^2 \;\propto\; R^{-3}, \tag{B.25}$$

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 primordial velocity dispersion is $\sim 10^{-17}$; the free-streaming length is $\lambda_{\rm fs} \sim 10^{-3}$ pc — milliparsecs, fifteen orders below the dwarf-galaxy scales where warm dark matter is constrained; and the de Broglie wavelength in a halo is $\lambda_{\rm dB} \simeq 154$ 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 $\lambda_{\rm fs} \sim$ mpc and $\lambda_{\rm dB} \sim 154$ 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*: the GG-6 head start, from the same $\pi$-locked ledger, no new sector — three layers, each amplified by the exponential rare-halo tail $n \sim e^{-\nu^2/2}$: an early-time small-scale power boost ($\Delta P/P \simeq 3\%$, hence $\sim 40\%$ more minihalos at $\nu \sim 5$); a baryon–dark-matter streaming-velocity damping from the gauged $\mathrm{U}(1)_B$ sector ($v_{\rm str} \to 0.8\, v_{\rm str}$, halving the baryonic filtering mass, so Population III ignites earlier and the smallest halos keep their gas); and a collapse-threshold reduction ($\delta_c : 1.686 \to 1.18$, a $\sim 500\times$ boost in the $> 10^{11} M_\odot$ abundance at $z \simeq 9$–10). The combined effect is a chronology shift $\Delta z \sim 2$–3 and a heavy direct-collapse seed channel $M_{\rm seed} \gtrsim 10^5 M_\odot$, all acting on the high-$k$/high-$\nu$ tail and leaving $S_8$ untouched. The GG-6-specific point: the streaming-drag sector *is* the gauged $\mathrm{U}(1)_B$ — the same sector whose locked ratio is the Kaxion's 5:2 portal (§B.6.3). Early-structure formation and the laboratory 5:2 lock are two projections of one piece of physics. (Honesty fence: the coldness is theorem-level, but these boost numbers are GG-A10-Cosmos's data-calibrated benchmarks, not first-principles-exact; 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

$$\delta_{\rm today} \sim 10^{-5} \times (1 + z_{\rm rec}) \sim 10^{-5} \times 1100 \sim 10^{-2} \ll 1 \tag{B.26}$$

— 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 does not merely retire a particle, it impeaches the framework's dark-sector closure.

**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 1/35 gearing 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 locks unique to the Kaxion (the 5:2 baryon portal and the cap index $\Xi_{\rm cap}$), closing with the inherited $\pi$-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 $\mathrm{d}\star H_{p+1}=0$ and the Bianchi identity $\mathrm{d}H_{p+1}=0$ 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 $\int_{\mathcal{M}_6}\sqrt{-G}\,|H_3|^2\to\int_{\mathcal{M}_4}\sqrt{-g}\,Z\,|H_3^{(4)}|^2$ with $Z=\int_0^K dr\int_0^{2\pi}d\theta\,e^{-2r}\,|\psi_0(r)|^2$ the zero-mode normalization; $Z$ defines $f_{\rm GS}^2$ 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-BI6 quantizes $\oint_{\Sigma_3} H_3 \in 2\pi\mathbb{Z}$ over closed three-cycles. Under (9.2) a large dual transformation winds $a_{\rm GS}$ by $2\pi f_{\rm GS}$, 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 two clocks carry hierarchically different decay constants. 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

$$\int_0^K dr \int_0^{2\pi} d\theta\, e^{-4r}\, e^{+6r} = 2\pi \int_0^K e^{+2r} dr = \pi\big(e^{2K} - 1\big) \;\simeq\; \pi e^{2K}, \tag{C.1}$$

an *infrared-dominated, exponentially enhanced* norm: $f_{\rm GS}^2 \propto e^{2K}$ 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

$$2\pi \int_0^K e^{-2r} dr = \pi\big(1 - e^{-2K}\big) \simeq \pi, \tag{C.2}$$

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) the Green–Schwarz direction naturally carries the *large* decay constant — which is why the raw (abundance) frame is GS-dominant and why $f_K \gg M_{\rm lock}$ requires no tuning; and (ii) the *ratio* of the two norms is exponentially sensitive to the cap region at $r = K$ where (C.1) accumulates — which is exactly why the mixing angle $\beta$ (equivalently $\Xi_{\rm cap}$) is the hard, throat-sensitive quantity that the mass-uniqueness conjecture names, and why this paper takes $f_K$ from the geometry — the benchmark cap loading inside the discrete mass-uniqueness window — and uses the abundance closure as the *cross-check* (fixing the one initial condition $\theta_i$) rather than attempting the cap integral directly. The residual of that reading is discrete (the $\Xi_{\rm cap}$ window); the quoted $\pm 6\%$ on $f_K$ and the mass is an uncertainty of the Standard-Model inputs ($\chi_{\rm QCD}$ and the transport), not of the geometry. The geometry shouts the ordering and whispers the coefficient.

## C.2 The flux-cell measure: $2\pi/K$ and the single-transit gearing $1/K$

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 $K$:

$$\varepsilon = \frac{\langle\psi_{\mathrm{GS}}, \psi_\theta\rangle_{\mathrm{transit}}}{\|\psi_{\mathrm{GS}}\|\,\|\psi_\theta\|} = \frac{1}{K} = \frac{1}{35}, \tag{C.3}$$

the gearing used in Eq. (31.3). 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 1/35 split is the paper's sharpest fingerprint, the single-transit overlap deserves the complete computation rather than 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,

$$\psi_{\mathrm{GS}}(r,\theta) = \frac{1}{\sqrt{2\pi K}}, \tag{C.4}$$

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,

$$\psi_\theta(r,\theta) = \frac{1}{\sqrt{2\pi}}\,u(r), \qquad \int_0^K u(r)^2\,\mathrm{d}r = 1, \tag{C.5}$$

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

$$\varepsilon = \langle\psi_{\mathrm{GS}}, \psi_\theta\rangle = \int_0^{2\pi}\!\mathrm{d}\theta \int_0^K\!\mathrm{d}r\, \frac{1}{\sqrt{2\pi K}}\,\frac{1}{\sqrt{2\pi}}\,u(r) = \frac{1}{K}\int_0^K u(r)\,\mathrm{d}r. \tag{C.6}$$

For a transit that deposits its weight uniformly across the throat, $u(r) = 1/\sqrt{K}$ and $\int_0^K u\,\mathrm{d}r = \sqrt{K}$, so $\varepsilon = \sqrt{K}/K = 1/\sqrt{K}$; for a transit sharply localized at unit depth, $\int_0^K u\,\mathrm{d}r \to 1$ and $\varepsilon \to 1/K$. The physically correct case is the second: the Chern–Simons improvement term $B_2 \wedge \mathrm{d}\theta \wedge (\ldots)$ 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

$$\varepsilon = \frac{1}{K} = \frac{1}{35} = 2.857\%, \tag{C.7}$$

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, $V = \frac{1}{2}\mu^2(a_{\mathrm{GS}}^2 + a_\theta^2)$, $\mu^2 = \chi_{\mathrm{QCD}}/f_K^2$) and the kinetic metric $\left(\begin{smallmatrix}1 & \varepsilon\\ \varepsilon & 1\end{smallmatrix}\right)$, the normal modes are the symmetric and antisymmetric combinations $u_\pm = (a_{\mathrm{GS}} \pm a_\theta)/\sqrt{2}$, with kinetic coefficients $1 \pm \varepsilon$. Canonically normalizing $\hat{a}_\pm = \sqrt{1\pm\varepsilon}\,u_\pm$ divides the common curvature $\mu^2$ by $1 \pm \varepsilon$, so

$$m_\pm^2 = \frac{\mu^2}{1\pm\varepsilon}, \qquad \nu_\pm = \frac{\nu_0}{\sqrt{1\pm\varepsilon}}, \qquad \nu_0 = \frac{\mu}{2\pi}, \tag{C.8}$$

and expanding to first order in the small gearing,

$$\frac{\Delta\nu}{\nu} = (1-\varepsilon)^{-1/2} - (1+\varepsilon)^{-1/2} = \varepsilon + \mathcal{O}(\varepsilon^3) = \frac{1}{K} = \frac{1}{35}. \tag{C.9}$$

The result is a pure frequency *ratio*: it depends on $\varepsilon$ alone, not on $\mu$, $f_K$, or $\Xi_{\rm cap}$, so even if the absolute mass shifts within the discrete cap window (Appendix C.6) the split stays exactly 1/35. Higher transits contribute at $\mathcal{O}(\varepsilon^2) = 1/K^2 \simeq 8\times10^{-4}$, below the linewidth, so the doublet is clean. The one model-level assumption is the equal-curvature (exchange) symmetry: if it is broken by a curvature asymmetry $\delta$, the split becomes $\sqrt{\varepsilon^2+\delta^2}$, so $1/K$ is the symmetric-limit value and a lower bound otherwise — the single input behind the first fingerprint's precise value, flagged as such in the ledger and the promotion roadmap (Appendix E.5).

## C.3 The GFF Hessian at the compact-angle attractor

On the stratum $(a_K, a_H)$ the relevant potential is $V = \chi_{\rm QCD}[1-\cos(a_K/f_K)] + \frac{1}{2}M_{\rm cap}^2 a_H^2 + \lambda\, a_K a_H$, with $M_{\rm cap}$ the throat-scale cap mass and $\lambda$ the residual gearing (suppressed by (C.3)). At the attractor $a_K = a_H = 0$:

$$\mathcal{H} = \begin{pmatrix} \chi_{\rm QCD}/f_K^2 & \lambda \\ \lambda & M_{\rm cap}^2 \end{pmatrix}, \qquad \det\mathcal{H} = m_K^2 M_{\rm cap}^2\Big(1 - \frac{\lambda^2}{m_K^2 M_{\rm cap}^2}\Big) > 0 \tag{C.10}$$

since $\lambda \sim m_K M_{\rm cap}/K \ll m_K M_{\rm cap}$: both eigenvalues are positive, the light one equal to $m_K^2$ up to corrections of order $K^{-2} \simeq 10^{-3}$. 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 $2\times2$ algebra that governs neutrino mixing, coupled pendula, and avoided level crossings, and the lesson transfers. Write the diagonal entries as $A = \chi_{\rm QCD}/f_K^2$ (the light, QCD-set curvature) and $B = M_{\rm cap}^2$ (the heavy, throat-set curvature), with off-diagonal $\lambda$. The eigenvalues of $\left(\begin{smallmatrix} A & \lambda \\ \lambda & B \end{smallmatrix}\right)$ are

$$\mu_\pm = \frac{A+B}{2} \pm \sqrt{\Big(\frac{B-A}{2}\Big)^2 + \lambda^2}. \tag{C.11}$$

The physics lives in the hierarchy $A \ll B$ (the light mass-squared is $\sim (8.1\,\mu\text{eV})^2$, the cap mass-squared is set by the throat scale, $\sim (E_{\rm GUT}e^{-K})^2$, vastly larger) and in the smallness of the gearing $\lambda \sim m_K M_{\rm cap}/K$ forced by the single-transit overlap (C.3). Expanding (C.11) to second order in the small ratio $\lambda/(B-A) \simeq \lambda/B$,

$$\mu_- = A - \frac{\lambda^2}{B-A} + \cdots = \frac{\chi_{\rm QCD}}{f_K^2}\Big(1 - \frac{\lambda^2}{m_K^2 M_{\rm cap}^2} + \ldots\Big), \qquad \mu_+ = B + \frac{\lambda^2}{B-A} + \ldots \tag{C.12}$$

The light eigenvalue $\mu_-$ is the physical Kaxion mass-squared; the fractional shift away from the bare QCD value is $\lambda^2/(m_K^2 M_{\rm cap}^2) \sim K^{-2} \simeq 8\times10^{-4}$, smaller than the lattice uncertainty on $\chi_{\rm QCD}$

itself, so to the accuracy of this paper $m_K^2 = \chi_{\rm QCD}/f_K^2$ exactly. 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.12), 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 $F_{\sf GFF}$ 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.4 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.7) 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,

$$\mathcal{L} \supset \frac{a}{f_a}\,\frac{\alpha_s}{8\pi}\,G^b_{\mu\nu}\widetilde{G}^{b\,\mu\nu}, \qquad \widetilde{G}^{b\,\mu\nu} = \tfrac{1}{2}\epsilon^{\mu\nu\rho\sigma}G^b_{\rho\sigma}, \tag{C.13}$$

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,

$$V(a) = \chi_{\rm QCD}\Big[1 - \cos\Big(\frac{a}{f_a}\Big)\Big] + \mathcal{O}(\cos 2(a/f_a)), \tag{C.14}$$

where the curvature scale $\chi_{\rm QCD}$ 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.14) sits at $a = 0$, i.e. at $\bar\theta_{\rm eff} = 0$; 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,

$$m_K^2 = \frac{{\rm d}^2V}{{\rm d}a^2}\bigg|_{a=0} = \frac{\chi_{\rm QCD}}{f_K^2} \quad\Longrightarrow\quad \boxed{m_K f_K = \sqrt{\chi_{\rm QCD}}}, \tag{C.15}$$

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,

$$\chi_{\rm QCD} = \frac{m_u m_d}{(m_u + m_d)^2}\, m_\pi^2 f_\pi^2 = (75.5 \pm 0.5\ \text{MeV})^4, \tag{C.16}$$

with $f_\pi \simeq 92$ MeV the pion decay constant and $m_\pi \simeq 135$ MeV; the lattice (Borsanyi et al., 2016) confirms this to percent accuracy. Inserting (C.16) into (C.15) gives the convenient form (Grilli di Cortona et al., 2016)

$$m_K = 5.70(7)\,\mu\text{eV} \times \Big(\frac{10^{12}\,\text{GeV}}{f_K}\Big), \tag{C.17}$$

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,

$$g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K}\Big(\frac{E}{N} - \frac{2}{3}\frac{4+z}{1+z}\Big), \qquad z \equiv \frac{m_u}{m_d} \simeq 0.462, \tag{C.18}$$

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.13) mixes the axion with the neutral pion and the $\eta$, removing the $a\,G\widetilde{G}$ term by a chiral rotation necessarily shifts the axion's coupling to two photons by a hadronic amount, evaluated from $\chi$PT 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 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.7,

$$g_{K\gamma\gamma} = \frac{\alpha}{2\pi f_K}\,(0.75) \implies 1.2 \times 10^{-15}\,\text{GeV}^{-1}, \tag{C.19}$$

evaluated at the physical $f_K = f_{\rm bulk}/\Xi_{\rm cap}$ 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.5 The 5:2 baryon portal, derived from the anomaly ledger

The seven-question essay (Appendix B.6.3) gives the physical story; this course carries the group theory and the dualization in full, so that a student can reconstruct the number 5:2 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

$$G_{\text{Bulk}} = \text{Spin}(1,5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}, \tag{C.20}$$

and the entry that has no analogue 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 SU(5), whose $X, Y$ leptoquarks drive $p \to e^+\pi^0$. 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,

$$S_{\text{GS}} = \int_{\mathcal{M}_6} B_2 \wedge X_4, \qquad X_4 = \sum_a c_a \frac{1}{8\pi^2} \text{Tr}(F_a \wedge F_a) + c_{\text{grav}}\, p_1, \tag{C.21}$$

where the sum runs over the gauge factors of (C.20), $p_1$ 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

$$\left(c_{\text{grav}},\, c_{\text{SU}(5)},\, c_{\text{U}(1)_B},\, c_{Qi}\right) = \left(1,\, -1,\, -\tfrac{2}{5},\, -\tfrac{2}{5}\right). \tag{C.22}$$

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 $a_K\, G\widetilde{G}$ that fixes the mass; the $\text{U}(1)_B$ entry gives the baryonic portal

$$\mathcal{L} \supset \frac{g_{KBB}}{4}\, a_K\, F_B \widetilde{F}_B, \tag{C.23}$$

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.22) 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

$$Y_0 = \mathrm{diag}\Big(-\tfrac{1}{3}, -\tfrac{1}{3}, -\tfrac{1}{3}, \tfrac{1}{2}, \tfrac{1}{2}\Big), \qquad B_0 = \mathrm{diag}\Big(\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{1}{3}, 0, 0\Big), \tag{C.24}$$

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 $\overline{\mathbf{5}} \oplus \mathbf{10}$ generation the $\mathbf{5}$ trace carries the ratio, and

$$\frac{\mathrm{Tr}(B_0^2)}{\mathrm{Tr}(Y_0^2)} = \frac{3 \times (1/3)^2}{3 \times (1/3)^2 + 2 \times (1/2)^2} = \frac{1/3}{1/3 + 1/2} = \frac{1/3}{5/6} = \frac{2}{5} = \frac{L-N}{L}, \tag{C.25}$$

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 $\mathrm{Tr}(B_0^2)/\,\mathrm{Tr}(Y_0^2)$ 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.

**Step 4 — source becomes response: why 2:5 inverts to 5:2.** At the level of the two-form *source* couplings (C.22), 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 $H_3 = \mathrm{d}B_2 \to \star\,\mathrm{d}a_K$ 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,

$$\frac{g_{KBB}}{g_{K\gamma\gamma}} = \Big(\frac{2}{5}\Big)^{-1} = \frac{5}{2}, \tag{C.26}$$

and a matched receiver, which reads power $\propto g^2$, sees

$$\boxed{\frac{P_B}{P_\gamma} = \Big(\frac{g_{KBB}}{g_{K\gamma\gamma}}\Big)^2 = \frac{25}{4} = 6.25}\,. \tag{C.27}$$

Baryon number, written faintly into the ledger as 2, is read loudly by the Kaxion as 5: it is the *stronger* of the two portals, which is what makes the two-channel ratio of Part VIII a decisive, astrophysics-free test.

**What is theorem and what is fenced.** The ratio (C.27) is exact at the gauge/anomaly level: it follows from (C.22) and (C.25) with no continuous input, and the chiral $-1.92$ dressing of the photon coupling cancels in the ratio (Appendix C.4). Two joints stay flagged. The coefficient vector (C.22) is read from the BI-6 factorization of Arisaka (2026, A3-BI6); and converting the gauge-level ratio into an actual nucleon-*spin* reading in a magnon or NMR booster requires a hadronic matrix element (the fenced open item of Part VIII). The number the geometry hands you is 25/4; the translation to a specific spin port carries one benchmark-level input.

## C.6 The cap index $\Xi_{\text{cap}} = K + N + L = 43$: the gearbox and its integer

This course makes explicit what the seven-question answer (Appendix B.6.4) 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 $\Xi_{\text{cap}}$ is: a wavefunction-renormalization integer.** $\Xi_{\text{cap}}$ is not a new parameter; it relates two normalizations of the *same* field,

$$f_K = \frac{f_{\text{bulk}}}{\Xi_{\text{cap}}}, \qquad m_K = \frac{\sqrt{\chi_{\text{QCD}}}}{f_K} = \frac{\sqrt{\chi_{\text{QCD}}}}{f_{\text{bulk}}}\,\Xi_{\text{cap}}, \tag{C.28}$$

where $f_{\text{bulk}}$ 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 acts identically on both doublet members and both portals, so it moves *where* the line sits but never the 1/35 split (Appendix C.2) or the 5:2 ratio (Appendix C.5). The fingerprints are theorem-level; only the address carries the conjecture.

**Step 2 — why the map is enormous: the throat is a gearbox.** 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 two light clocks acquire wildly asymmetric kinetic normalizations (Appendix C.1),

$$Z_{\text{GS}} = 2\pi\int_0^K e^{2r}\,\mathrm{d}r = \pi\big(e^{2K} - 1\big) \simeq \pi e^{2K}, \qquad Z_\theta = 2\pi\int_0^K e^{-2r}\,\mathrm{d}r \simeq \pi, \tag{C.29}$$

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. $\Xi_{\text{cap}}$ is that map expressed in the throat's own quantized units: the Green–Schwarz clock is normalized after the warp has run its full $K$ e-folds, the boundary clock sits at the throat's mouth, and converting one to the other multiplies by the throat depth in integer flux units, plus whatever the cap adds. The exponential $e^{70}$ proves $\Xi_{\text{cap}} \gg 1$ by thirty orders — but it contributes *nothing* to the integer value; the value is pure counting, which is the point students most often miss.

**Step 3 — why the integer is additive: flux cannot end fractionally.** Three facts are provable from the metric, the OGN-5 integers, and integer flux quantization alone (the referee-checkable Proposition of Part V):

1. $\Xi_{\text{cap}} > 1$ strictly, by the exponential asymmetry (C.29). (This bounds it away from 1; it does not set the value.)

2. $\Xi_{\rm cap} = K + \Sigma$, an *additive* integer form. The throat is $K = 35$ flux quanta deep, and the orbifold cap that terminates its ultraviolet end must carry an integer number of extra quanta, because $\oint_{\Sigma_3} H_3 \in 2\pi\mathbb{Z}$ forbids a fractional flux line from ending on a compact fixed-point cap. This is the load-bearing fact: it makes the residual an *integer*, not a continuous knob.

3. $\Sigma$ 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

$$\begin{aligned}\Sigma \in \{0,3,5,7,8,10,12\} \;\Rightarrow\; & \Xi_{\rm cap} \in \{35,38,40,42,43,45,47\} \\ \Rightarrow\; & \nu_K = 1.95\,(\Xi_{\rm cap}/43) \in [1.59, 2.13]\ \text{GHz}.\end{aligned} \tag{C.30}$$

So the absolute mass is not a continuous parameter with an error bar — it is one of *seven discrete frequencies* (Table 11). That discreteness is genuinely derived, and it is the experimentally decisive content: the prediction lives in a finite list before any conjecture is invoked.

**Step 4 — why 43: the ultraviolet/infrared sorting of the ledger.** The whole conjecture reduces to one choice — which subset $\Sigma$? The benchmark $\Sigma = N + L$ is not "add two nice integers"; it is a UV/IR sorting forced by how these same integers function everywhere else in the series. The family-projector flux $N$ and the gauge-rank flux $L$ encode the SU(5) $\to$ 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:

$$\Sigma = N + L \quad\Longrightarrow\quad \boxed{\Xi_{\rm cap} = K + N + L = 35 + 3 + 5 = 43}. \tag{C.31}$$

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 single finite computation — the $X_4$ flux integral of Arisaka (2026, A3-BI6) restricted to the cap's orbifold fixed-point divisors — and until it returns $N + L$ the assignment is a flux-bookkeeping proposition, the single sharpest open item of the paper (its promotion is entry 1 of the roadmap, Appendix E.5). The uncertainty is disposable two independent ways: *by theory*, the one-page cap-cohomology integral collapses the window to a point; *by experiment*, since $\nu_K = 1.95\,(\Xi_{\rm cap}/43)$ GHz, the discovery frequency *measures* $\Xi_{\rm cap}$ *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.7 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 $\Omega_{\rm DM}$, 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,

$$\ddot{\theta} + 3H\dot{\theta} + m^2(T)\,\theta = 0, \tag{C.32}$$

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.32) 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: $\theta$ is frozen at its initial value $\theta_i$, the $\mathcal{O}(1)$ misalignment angle ($\theta_i \simeq 1.4$ rad, fixed by the abundance). 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 $H = \sqrt{\pi^2 g_*/90}\, T^2/M_{\rm Pl}$) while $m(T)$ rises. Oscillation begins at the temperature $T_{\rm osc}$ where friction can no longer hold the field,

$$3H(T_{\rm osc}) \simeq m(T_{\rm osc}). \tag{C.33}$$

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 $n_a(T_{\rm osc}) = \frac{1}{2} m(T_{\rm osc}) f_K^2 \theta_i^2$ (times an $\mathcal{O}(1)$ anharmonic correction for large $\theta_i$). 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

$$\Omega_K h^2 = \frac{m_K\, n_{a,0}}{\rho_c/h^2} = \frac{m_K\, Y\, s_0}{\rho_c/h^2}, \tag{C.34}$$

with $s_0 = 2891\ \text{cm}^{-3}$ and $\rho_c/h^2 = 1.05 \times 10^{-5}\ \text{GeV}\,\text{cm}^{-3}$ 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 $m(T) = m_K\,(T_{\rm QCD}/T)^{n/2}$. Solving the onset condition (C.33),

$$T_{\rm osc}^{2+n/2} \propto m_K\, T_{\rm QCD}^{n/2} M_{\rm Pl} \implies T_{\rm osc} \propto m_K^{2/(4+n)} \propto f_K^{-2/(4+n)}, \tag{C.35}$$

using $m_K \propto 1/f_K$ from (C.15). Carrying this through (C.34) with $n_a(T_{\rm osc}) \propto m(T_{\rm osc}) f_K^2 \propto T_{\rm osc}^2 f_K^2$ and $Y \propto n_a/T_{\rm osc}^3 \propto f_K^2/T_{\rm osc}$ gives

$$\Omega_K h^2 \;\propto\; m_K\,\frac{f_K^2}{T_{\rm osc}} \;\propto\; \frac{1}{f_K}\, f_K^2\, f_K^{2/(4+n)} \;=\; f_K^{(6+n)/(4+n)}\,\theta_i^2. \tag{C.36}$$

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 gives the canonical normalization $\Omega_a h^2 \simeq 0.090\,\theta_i^2\,(f_K/10^{12}\,\text{GeV})^{1.165}$; the per-step arithmetic with the GG input $f_K = 7.1\times 10^{11}$ GeV (the abundance fixing $\theta_i \simeq 1.4$ rad) is carried in the audit, Appendix E.1.

**Step 7 — the closure.** Because the target is itself derived, $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2) = 0.2676$ (so $\Omega_K h^2 = 0.1215$), and the decay constant is fixed by the geometry, $f_K = (7.1\pm 0.5)\times 10^{11}$ GeV, the law (C.36) contains *one* unknown — the misalignment angle. The mass follows from the susceptibility, the angle from the budget,

$$m_K = \frac{\sqrt{\chi_{\rm QCD}}}{f_K} = (8.1\pm 0.5)\,\mu\text{eV}, \qquad \theta_i = \Big(\frac{\Omega_K h^2}{0.090}\Big)^{1/2}\Big(\frac{10^{12}}{f_K[\text{GeV}]}\Big)^{0.583} \simeq 1.4 \text{ rad}, \tag{C.37}$$

and the frequency and coupling follow from the mass. The $\pm$ is dominated by the lattice thermodynamics ($\chi(T)$, $g_*$), compressible to a few percent with no new theory — the honest residual of an otherwise parameter-free prediction.

## C.8 The $\pi$-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.7) aims at is $\Omega_{\rm DM} = 0.2676$, 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

$$\Omega_m = \frac{1}{\pi} = 0.3183, \quad \Omega_b = \frac{1}{2\pi^2} = 0.0507, \quad \Omega_{\rm DM} = \frac{1}{\pi} - \frac{1}{2\pi^2} = 0.2676, \quad \Omega_\Lambda = 1 - \frac{1}{\pi} = 0.6817, \tag{C.38}$$

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 lemma, the Hopf-fiber identification.

**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 $S^1 \hookrightarrow S^3 \to S^2$, and the inventory is then unit-sphere volumes: $\mathrm{Vol}(S^1) = 2\pi$, $\mathrm{Vol}(S^2) = 4\pi$, $\mathrm{Vol}(S^3) = 2\pi^2$. Spreading the critical density uniformly over $S^3$, assigning matter to the fiber and dark energy to the base, gives

$$\Omega_m = \frac{\mathrm{Vol}(S^1)}{\mathrm{Vol}(S^3)} = \frac{2\pi}{2\pi^2} = \frac{1}{\pi}, \qquad \Omega_b = \frac{1}{\mathrm{Vol}(S^3)} = \frac{1}{2\pi^2}, \qquad \Omega_{\mathrm{DM}} = \frac{2\pi - 1}{2\pi^2} = \frac{1}{\pi} - \frac{1}{2\pi^2}. \tag{C.39}$$

The two "flexible symbols" $\pi$ 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 $\pi$-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,

$$\frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi} = 0.159 \ \ (\text{sky: } 0.156), \qquad \frac{\rho_{\mathrm{DM}}}{\rho_b} = 2\pi - 1 = 5.28 \ \ (\text{sky: } 5.37), \tag{C.40}$$

both epoch-independent and both matching observation. These are the robust $\pi$-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 $\Lambda$ (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 $\sim 75$ 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.7 must deliver $\Omega_K h^2 = 0.1215$, and it does, fixing the one initial condition $\theta_i \simeq 1.4$ 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 how the two-channel ratio is made unbiased, and then the three courses new to this edition: the four-peak template and its matched filter, the KAXIO instrument sensitivity and its one-year coverage of the discrete window, and the discovery-week error budget behind the four 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)

$$P_{\rm sig} = g_{K\gamma\gamma}^2 \, \frac{\rho_{\rm DM}}{m_K} \, B_0^2 \, V \, C \, Q_L \, \frac{\beta}{1+\beta}, \tag{D.1}$$

with $Q_L$ the loaded quality factor, $\rho_{\rm DM} \simeq 0.45$ GeV cm$^{-3}$ the local density, and $\beta$ 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 $\beta/(1+\beta) \to 1$ in (D.1), the convention of the public 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:

$$1\ \text{T} = 1.9535\times10^{-16}\ \text{GeV}^2, \qquad 1\ \text{cm} = 5.068\times10^{13}\ \text{GeV}^{-1}, \qquad 1\ \text{GeV}^2\ \text{(power)} = 2.434\times10^{14}\ \text{W}. \tag{D.2}$$

Inserting the Kaxion's numbers at the CAPP-MAX-class row of Table 17 ($g_{K\gamma\gamma} = 1.2\times10^{-15}$ GeV$^{-1}$, $m_K = 8.1\,\mu$eV, $B_0 = 12$ T, $V = 6.4$ L, $C = 0.60$, $Q_L = 10^5$), the factors of (D.1) become

$$\begin{aligned} \frac{\rho_{\rm DM}}{m_K} &= \frac{0.45\ \text{GeV cm}^{-3}}{8.1\times10^{-15}\ \text{GeV}} = 4.27\times10^{-28}\ \text{GeV}^3, \\ B_0^2 &= \left(12\times1.9535\times10^{-16}\ \text{GeV}^2\right)^2 = 5.50\times10^{-30}\ \text{GeV}^4, \qquad V = 6.4\ \text{L} = 8.33\times10^{44}\ \text{GeV}^{-3}, \end{aligned} \tag{D.3}$$

using $1\ \text{cm}^3 = 1.302\times10^{41}\ \text{GeV}^{-3}$ from the second constant, so that

$$P_{\rm sig} = g_{K\gamma\gamma}^2 \, \frac{\rho_{\rm DM}}{m_K} \, B_0^2 \, V \, C \, Q_L = 1.69\times10^{-37}\ \text{GeV}^2 = 4.1\times10^{-23}\ \text{W}, \tag{D.4}$$

the $P_{\rm sig}$ 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 $\nu$ in place of $\omega$ silently inflates every power in the table by $2\pi$.

**Step 2 — the noise.** The signal competes with thermal and amplifier noise of system temperature $T_{\rm sys} = T_{\rm phys} + T_{\rm amp}$. The irreducible quantum floor at the Kaxion frequency is $T_q = h\nu/k_B = 0.094$ 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 $t$ over the signal bandwidth $\Delta\nu_a$ is

$$\mathrm{SNR} = \frac{P_{\rm sig}}{k_B T_{\rm sys}}\sqrt{\frac{t}{\Delta\nu_a}}, \qquad \Delta\nu_a = \frac{\nu}{Q_a} \sim \nu\,\frac{\langle v^2\rangle}{c^2} \sim 10^{-6}\,\nu, \tag{D.5}$$

the celebrated radiometer equation: 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).

**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

$$\frac{\mathrm{d}\nu}{\mathrm{d}t} \;\propto\; \frac{1}{\mathrm{SNR}^2}\Big(\frac{P_{\rm sig}}{k_B T_{\rm sys}}\Big)^2 \frac{1}{\Delta\nu_a} \;\propto\; \frac{g_{K\gamma\gamma}^4\, B_0^4\, V^2\, C^2\, Q_L}{T_{\rm sys}^2}\,. \tag{D.6}$$

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 $\sim 2.5$ times smaller than KSVZ) costs $\sim 40$ 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 $T_{\rm sys}^2$, 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 and the deconvolved ratio

**Single port.** The steady-state response of the QGI oscillator (35.3) to a drive at frequency $\omega$ 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

$$|\chi(\omega)|^2 = \frac{1}{(\omega_0^2 - \omega^2)^2 + (\omega_0\omega/Q_L)^2}, \tag{D.7}$$

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,

$$S_{\rm out}(\nu) = \mathcal{G}\,|\chi(\nu)|^2\,[\mathcal{T}(\nu) + S_{\rm noise}], \tag{D.8}$$

with $\mathcal{T}$ the doublet template (33.3) 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 $\nu/Q_a$ is far inside the cavity width $\nu/Q_L$ (the hierarchy (33.1)); off resonance, the partner tone at detuning $\Delta\nu/\nu = 1/35$ is suppressed by $[2Q_L\Delta\nu/\nu]^{-2} \simeq 3\times 10^{-8}$.

**Two ports and the unbiased ratio.** The booster of Part VIII records two spectra $S_B, S_\gamma$ through two independent chains. Define the deconvolved estimator

$$\widehat{R} = \frac{\int W(\nu)\,[S_B(\nu) - \hat{S}^B_{\rm noise}]/\,\mathcal{G}_B|\chi_B|^2\,d\nu}{\int W(\nu)\,[S_\gamma(\nu) - \hat{S}^\gamma_{\rm noise}]/\,\mathcal{G}_\gamma|\chi_\gamma|^2\,d\nu}, \tag{D.9}$$

with $W$ any common spectral window covering the line. Because the condensate factor $\rho_{\rm loc}\mathcal{T}(\nu)$ is common to both numerator and denominator, every astrophysical unknown — the local density, the velocity distribution, the tone ratio $\kappa$, even the day's halo weather — cancels identically, leaving

$$\langle\widehat{R}\rangle = \Big(\frac{g_{KBB}}{g_{K\gamma\gamma}}\Big)^2 = \frac{25}{4}, \qquad \sigma_{\widehat{R}} \simeq \frac{25}{4}\sqrt{2}\,\Big(\mathrm{SNR}_B^{-1} \oplus \mathrm{SNR}_\gamma^{-1} \oplus \sigma_{\rm cal}\Big), \tag{D.10}$$

the error budget being statistical noise in each port plus the calibration-chain systematic $\sigma_{\rm cal}$ (the prefactor $\sqrt{2}$ is deliberate conservatism over the pure statistical limit, which is $1/\mathrm{SNR}_B$ when $\mathrm{SNR}_\gamma \gg \mathrm{SNR}_B$) — the quantity that sets the $\pm 20\% \to \pm 5\%$ staging of the falsifier F-Kaxion-1b. Theorem 32.1 is (D.10) read at infinite signal-to-noise.

**Dual-mode integrity.** For the coincidence resonator of Part VIII, 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 $2.8 \times 10^3$ 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 $R_{\rm acc} = R_{\rm FA}^2\tau_w$ used in Part VI.

## D.3 The four-peak template and its matched filter

Parts VI–IX assert 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 the template and its statistics from the two ingredients already derived: the doublet spacing (Appendix C.2) and the two-channel ratio (Appendices C.5, D.2).

**Step 1 — why four peaks.** The forced light sector is *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 $2 \times 2$ pattern — four peaks — whose frequencies and relative powers are fixed by the geometry:

$$\{\nu_-, \nu_+\} = \nu_K\{(1+\varepsilon)^{-1/2}, (1-\varepsilon)^{-1/2}\},\ \varepsilon = \tfrac{1}{35}; \qquad \frac{P^B_\pm}{P^\gamma_\pm} = \Big(\frac{g_{KBB}}{g_{K\gamma\gamma}}\Big)^2 = \frac{25}{4}. \tag{D.11}$$

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 B.6.2); 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 $N_0 = k_B T_{\rm sys}$ and the template family $\mathcal{T}(\nu;\vartheta)$ of (D.11) (parameters $\vartheta$ = center, $\kappa$, amplitude), the optimal (Neyman–Pearson) statistic is the matched filter — the noise-weighted cross-correlation of data with template,

$$\Lambda = \frac{\int \mathcal{T}(\nu)\, S(\nu)\, \mathrm{d}\nu}{\sqrt{\int \mathcal{T}(\nu)^2\, \mathrm{d}\nu}} \Big/ \sqrt{N_0}, \qquad \langle \Lambda \rangle_{\rm signal} = \mathrm{SNR}_{\rm tot}, \tag{D.12}$$

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

$$\mathrm{SNR}^2_{\rm tot} = \sum_{c\in\{\gamma,B\}} \sum_{s\in\{+,-\}} \mathrm{SNR}^2_{c,s}, \tag{D.13}$$

so a network of four marginal $2\sigma$ 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 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 $R_{\rm FA}$ and the coincidence window is $\tau_w$, the accidental four-fold (or, conservatively, two-fold across the doublet) rate is suppressed as a power of $R_{\rm FA}\tau_w$,

$$R_{\rm acc} = R^2_{\rm FA}\, \tau_w \quad \text{(doublet)}, \qquad R_{\rm acc} \propto R^4_{\rm FA} \quad \text{(full pattern)}, \tag{D.14}$$

the estimate used in Part VI. The two identification observables are engineered to be astrophysics-free: the doublet spacing is a pure frequency *ratio* (no amplitude enters), and the channel ratio $\widehat{R} \to 25/4$ 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 is therefore not merely a detection — it is a reading of the gauge vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$ in a laboratory.

## D.4 The KAXIO instrument and the one-year projection

Part X proposes KAXIO, a second-generation instrument built to record the full four-peak pattern 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 resonators reading the same condensate through two chains — an electromagnetic cavity and an electron-magnon (baryon-spin) booster — each tunable across the window and each carrying a companion mode at the 1/35 partner frequency. The design point is a strong solenoid and a large cold bore,

$$B_0 \simeq 12\ \mathrm{T}, \quad V \simeq 6.4\ \mathrm{L}, \quad C \simeq 0.6, \quad Q_L \simeq 10^5, \quad \beta \simeq 2, \quad T_{\rm sys} \simeq 0.1\ \mathrm{K}, \tag{D.15}$$

i.e. a CAPP-MAX-class electromagnetic channel with an added spin channel. Inserting (D.15) and the Kaxion's DFSZ coupling $g_{K\gamma\gamma} = 1.2 \times 10^{-15}\,\mathrm{GeV}^{-1}$ into the Sikivie power (D.1) (with the over-coupling factor $\beta/(1+\beta) = 2/3$ carried explicitly) gives a signal power in the yoctowatt range, $P_{\rm sig} \simeq 3 \times 10^{-23}$ W, worked in 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

$$t_{\rm peak} = \Big(\frac{\mathrm{SNR}\, k_B T_{\rm sys}}{P_{\rm sig}}\Big)^2 \Delta\nu_a, \qquad \Delta\nu_a = \frac{\nu}{Q_a} \simeq 2\ \mathrm{kHz}, \tag{D.16}$$

so that reaching the design SNR per resolution element takes of order days, not the multi-year single-instrument campaign of a first-generation search, because the matched four-peak filter (D.13) pools the two tones and two channels. The scan cost to cover a band $B$ follows from the scan-rate wall (D.6), $t_{\rm scan} = B\,(\mathrm{d}t/\mathrm{d}\nu)$, 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 discrete window.** The prediction is not a continuous band but the seven discrete frequencies of (C.30), $\nu_K \in [1.59, 2.13]$ GHz. KAXIO's run plan (Part X) allocates the year as a staged march through those tunings: the benchmark $\Xi_{\rm cap} = 43$ frequency (1.95 GHz) first, then its integer neighbors outward, each visited to DFSZ depth with the four-peak filter. Because the window is 0.54 GHz wide and each resolution march proceeds at the wall rate (D.6) for the KAXIO parameters, the full seven-value window is a one-year program:

$$t_{\rm window} = \sum_{i=1}^{7} t_{\rm scan}(\nu_i) \ \lesssim\ 1\ \mathrm{yr}, \tag{D.17}$$

after which either a four-peak pattern has been recorded at one of the seven tunings — measuring $\Xi_{\rm cap}$ outright — or the entire DFSZ-depth, doublet-aware 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.5 The discovery-week error budget and the four verdicts

Part IX simulates the data of discovery week in advance and reads four 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 $\pm 22\%$ ($\rho_{\rm loc} = 0.45 \pm 0.10$ 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 four verdicts, and their budgets.** The verdicts map one-to-one onto the four falsifiers:

1. **Shape** (F-Kaxion-1d, wrong location). The camped line matches the boosted-Maxwellian condensate shape at the predicted frequency. Statistical only after calibration; the theory-side $\pm 6\%$ on $\nu_K$ is a band-scan cost, not a within-panel error, and is priced separately (Appendix C.6).

2. **Modulation** (halo confirmation). The line rides $\simeq 588$ Hz above $\nu_K$ and breathes annually with amplitude $\simeq 74$ Hz; the sidereal daily signature is $\simeq 1.9$ 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 a single-cavity *hint* ($\sim 2\sigma$ in a high-duty year) — the observable KAXIO's aperture (Appendix D.4) exists to buy.

3. **Doublet** (F-Kaxion-1a, missing partner). A single-mode cavity tuned to one tone is blind to the partner by $\simeq -75$ dB at $Q_L = 10^5$; the $\pm 5\%$ on $Q_L$ shifts this by $\mp 0.4$ dB on a 75 dB suppression — systematics-proof by four orders. The test is a deliberate retune (or the Part-VIII second mode), and at camped SNR its verdict arrives within a day.

4. **Channel ratio** (F-Kaxion-1b, wrong ratio). The deconvolved matched-channel ratio $\widehat{R}$ converges to 25/4; halo density, magnet field, and every frame-dependent factor cancel by construction (Appendix D.2), so its systematic budget is *calibration only* ($\pm 20\%$ first-generation, tightening to $\pm 5\%$ at the day-30 calibration upgrade).

**Why the discovery cannot be mistaken — and why the refutation is unarguable.** The two verdicts that carry the geometry (the doublet and the channel ratio) are precisely the two engineered to cancel the halo systematic — one is a frequency ratio, the other a power ratio — so they sit in the two 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, lacks its $10^{-2}$ partner, or converges off 25/4 does not become an anomaly — it falsifies the paper. A theory confirmed by six 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 referee 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 41 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 41:** End-to-end audit. Inputs: $\chi_{\rm QCD}^{1/4} = 75.5(5)$ MeV, $b = 4.08(20)$, $T_{\rm QCD} = 0.150(15)$ GeV, $g_* = g_{*s} = 73$, $f_K = 7.1 \times 10^{11}$ GeV, target $\omega = 0.1215(18)$, $M_{\rm Pl} = 1.221 \times 10^{19}$ GeV, $s_0 = 2891\,{\rm cm}^{-3}$, $\rho_c/h^2 = 1.0537 \times 10^{-5}\,{\rm GeV\,cm}^{-3}$. Every input is printed so that the chain can be recomputed independently rather than taken on trust.

| Step | Expression | Value |
|---|---|---|
| Susceptibility | $\sqrt{\chi_{\rm QCD}}$ | $5.70(8) \times 10^{-3}\,{\rm GeV}^2$ |
| Master normalization | Eq. (25.1) at $f = 10^{12}$ GeV, $\theta_i = 1$ | $\Omega h^2 = 0.090$ |
| cross-check | lit. value at same inputs (Grilli di Cortona et al., 2016) | agrees to 8% |
| Decay constant | geometry; cap loading $\Xi_{\rm cap} = 43$ | $f_K = 7.1 \times 10^{11}$ GeV |
| Mass | $\sqrt{\chi_{\rm QCD}}/f_K$ | $m_K = 8.1 \times 10^{-15}$ GeV $= 8.1\,\mu$eV |
| Frequency | $241.80\,{\rm MHz} \times m[\mu{\rm eV}]$ | $\nu_K = 1.95$ GHz; $\Delta\nu = \nu/35 = 56$ MHz |
| Misalignment angle | $(0.1215/0.090)^{1/2}(10^{12}/f_K[{\rm GeV}])^{0.583}$ | $\theta_i \simeq 1.4$ rad |
| Onset | $T_{\rm osc}^{6.08} = m_K T_{\rm QCD}^{4.08} M_{\rm Pl}/42.6$ | $T_{\rm osc} \simeq 1.0$ GeV; $t_{\rm osc} \simeq 0.3\,\mu$s |
| Onset mass | $(0.15/1.0)^{4.08}$ | $m_K(T_{\rm osc}) \simeq 4.4 \times 10^{-4}\, m_K(0)$ |
| Round trip | Eq. (25.1) forward | $\Omega h^2 = 0.1215$ ✓ |
| Coupling bracket | $8/3 - 1.92(4)$ | 0.75(4) |
| Coupling | $0.75/(137.036 \cdot 2\pi f_K)$ | $g_{K\gamma\gamma} = 1.23 \times 10^{-15}\,{\rm GeV}^{-1}$ |
| Lifetime | $64\pi/(g_{K\gamma\gamma}^2 m_K^3)$ | $1.7 \times 10^{50}$ 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-1a** (missing 1/35 partner at $10^{-2}$ relative power), **F-Kaxion-1b** (matched two-channel ratio $\neq 25/4$), **F-Kaxion-1c** (DFSZ-depth, doublet-aware 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). 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-1a–1d** (falsifiers, above), **O-Kaxion-1–69** (objections, Part XIV), **PS-Kaxion-1–8** (problems solved, §74.1), and **I-Kaxion-1–10** (structural innovations, §74.2). Any of (a), (b), (d) tests the geometry independently of the named conjectures; (c) tests the closure sheet as fenced. Numerical decision rules, pre-declared: F-Kaxion-1a fires on partner absence at 95% CL at the $10^{-2}$ relative-power threshold; F-Kaxion-1b fires on a converged matched ratio differing from 25/4 at $> 5\sigma$ including the calibration systematic.

## 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 42 names, for each headline quantity, the parent paper and the specific result it descends from — the ledger a referee needs to audit provenance, and the reading list a student needs to go deeper.

**Table 42:** 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 $B_2$, 5:2 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 $\Omega_{\mathrm{DM}} = 1/\pi - 1/(2\pi^2)$ | Arisaka (2026, A10-Cosmos) | $\pi$-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 $M_{\mathrm{lock}} = 3.85$ TeV | Arisaka (2026, P1-KK) | the same $K = 35$ read at the TeV lock |
| Mass $m_K f_K = \sqrt{\chi_{\mathrm{QCD}}}$ | this paper + lattice QCD | the pNGB relation, closed by the geometry |

## E.4 Named-theorem-and-conjecture deck

Table 43 is the honest ledger of what the paper proves and what it fences. It separates the three theorem-level claims (forced by geometry and standard physics, with no continuous freedom) from the conjecture-fenced ones (each resting on a single, named, referee-checkable computation), and gives each its status. Nothing in the discovery figures depends on a conjecture: the two identification observables (the 1/35 doublet and the 25/4 ratio) are theorem-level.

**Table 43:** The claims of the paper, graded. T = theorem-level; C = conjecture-fenced (one named computation from promotion, see Appendix E.5). Every claim the paper makes appears once, with its grade and with what it would take to promote it. T is theorem-level; 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 | $m_K f_K = \sqrt{\chi_{\rm QCD}}$ (gluonic portal) | T | PQ mechanism + lattice (App. C.4) |
| T-Kaxion-3 | $E/N = 8/3$: the Kaxion sits on the DFSZ line | T | SU(5) normalization (App. C.4) |
| T-Kaxion-4 | Baryon portal at $g_{KBB}/g_{K\gamma\gamma} = 5/2$ | T | anomaly ledger (App. C.5) |
| T-Kaxion-5 | Doublet exists; scale-free split | T | two light zero-modes (App. C.2) |
| C-Kaxion-a | Split $= 1/K = 1/35$ exactly | C | throat exchange symmetry |
| C-Kaxion-b | $\Xi_{\rm cap} = K + N + L = 43$ (absolute mass) | C | cap flux integral $\Sigma = N + L$ |
| C-Kaxion-c | $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ | C | Hopf-fiber identification (Arisaka (2026, A10-Cosmos)) |
| F-Kaxion-1a–1d | The four master falsifiers | — | 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 44 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 44:** 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** |
|---|---|---|
| $\Xi_{\mathrm{cap}} = 43$ (C-Kaxion-b) | Evaluate the $X_4$ flux integral of Arisaka (2026, A3-BI6) on the cap's orbifold fixed-point divisors; check it returns $\Sigma = N + L$ | Collapses the seven-value window to a point; fixes the absolute mass |
| 1/35 split (C-Kaxion-a) | Prove the throat reflection (exchange) symmetry is exact, or bound the curvature asymmetry $\delta$ | Makes the doublet spacing an exact theorem, not a symmetric-limit value |
| $\pi$-inventory (C-Kaxion-c) | Settle the Hopf-fiber identification cohomology in Arisaka (2026, A10-Cosmos) | Turns $\Omega_m = 1/\pi$ into a theorem; removes the last cosmological input |
| 5:2 spin port | Compute the nucleon-spin chiral matrix element for the magnon/NMR channel | Converts the gauge-level 25/4 into a laboratory spin-channel prediction |
| Tone-power ratio $\kappa$ | Model the primordial energy partition at the lock/saddle crossing | Predicts the relative heights of the two tones (currently a free amplitude) |

# F 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.

| | | | |
|---|---|---|---|
| $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$ | bulk manifold (GG-A2-GG6 T1) | $(M, N, L; J, K)$ | OGN-5 $= (1, 3, 5; 7, 35)$ |
| $r \in [0, K]$ | scale coordinate, $r = \ln(E_{\mathrm{GUT}}/\mu)$ | $\theta_{\mathrm{GS}}, \theta_W$ | tensor shadow; holonomy echo |
| $B_2, H_3$ | Green–Schwarz tensor, field strength | $a_K, a_H$ | light/heavy eigenstates; $\beta$ mixing |
| $X_4$ | BI-6 ledger four-form | $f_K, m_K$ | decay constant; mass |
| GDOS,TGG,CCI | boundary law; total stress; seam law | $\theta_i$ | misalignment angle $\simeq 1.4$ rad |
| GRIP,GFF,GEP,CEP | runtime; free energy and its parts | $\chi_{\mathrm{QCD}}$ | topological susceptibility |
| DGI,GSSB,GA,GGA | cascade: top, saddles, attractors; GGA the unique deepest (ground state) | $g_{K\gamma\gamma}, g_{KBB}$ | photon, baryonic couplings (5:2) |
| GRIP-In, GRIP-Out | runtime regimes: fixed-landscape descent; landscape reorganization | | |
| $M_{\mathrm{lock}} = \frac{1}{2} E_{\mathrm{Pl}} e^{-K}$ | lock scale 3.85 TeV | $\Xi_{\mathrm{cap}}$ | cap index (bench. 43) |
| $\Omega_{\mathrm{DM}} = 1/\pi - 1/2\pi^2$ | $\pi$-locked target | $\nu_K$ | the unique Kaxion window, Eq. 36.1 |

# 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 two named conjectures. 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.

## G 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 missing 1/35 doublet or a violated 5:2 falsifies the whole construction, so this can.

**Objection 3 (O-Kaxion-3) (the precision auditor)**
**"Your Standard-Model papers claimed per-mille accuracy; ±6% 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 mass chain, by contrast, passes through QCD thermodynamics in an expanding universe, and every percent of its ±6% is an itemized lattice-susceptibility or transport uncertainty — with a named, theory-free path to $\sim 3\%$ (a dedicated $\chi(T)$ lattice computation through the crossover). Crucially, the dimensionless fingerprints that do the identifying are algebraic and carry *no* error at all; the ±6% lives only in the mass.

**Objection 4 (O-Kaxion-4) (the neatness skeptic)**
**"$\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ 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 $\Omega_{\rm DM}$ away from 0.268, which the data do not.

**Objection 5 (O-Kaxion-5) (the fudge skeptic)**
**"The factor-of-2.6 correction to GG-A10-Cosmos looks like fudging to the answer."**

*Response.* The correction was forced, not chosen. An independent recomputation of the misalignment abundance from scratch revealed GG-A10-Cosmos's engineering normalization to be high by 2.6; it was applied with full provenance (Part V), cross-validated against the precision misalignment literature, and it *sharpened* the prediction rather than moving it to a desired value. Decisively, it left every theorem untouched — the existence, the coldness, the invariant $m_K f_K = \sqrt{\chi_{\rm QCD}}$, and both fingerprints are unchanged. A correction that tightens error bars while leaving theorems intact is the falsification postulate operating as intended.

**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 magnon port reading 25/4, 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 peer-review guardian)**
**"You call things 'theorems' that no journal has refereed."**

*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 referee it themselves: which rows are proved, which are derived through Standard-Model inputs, which rest on the two named conjectures. Refereeing is welcome and will happen; but the paper's decisive review was designed to be harsher than any journal's: four pre-registered falsifiers, executed by competing experimental collaborations, on printed dates (§36.9). A referee can reject a manuscript; 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 $J/\psi$ 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 collapses from a factor that costs 2.2× in integration time to exactly one (§37.1), and the June-2026 episode — a $5.1\sigma$ local excess deflated to $3.5\sigma$ 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 $\alpha^{-1}$, $\sin^2\theta_W$, 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 a half-octave, published, with a marked row. A combinatoric search that could land anywhere proves nothing by landing; a cascade that can only land in one half-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 $\sqrt{\chi_{\rm QCD}}/f_K$ puts it there; that the field's DFSZ frontier currently sits at $5.4\,\mu$eV, 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 ∼2030–2032 (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 doublet 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 positive 1/35 pair 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.

# H 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 $S^1_\theta$ a dimensionless phase circle (GG-A2-GG6 T2), neither carrying a spatial KK graviton. The only accessible tower is the radial lock cluster at $M_{\text{lock}} \approx 3.85$ TeV — a collider *prediction* (P1-KK), not a liability — while the angular tower sits unobservably at $E_{\text{GUT}}$. The objection mistakes a non-spatial geometry for a large spatial one.

**Objection 14 (O-Kaxion-14) (the two-clocks skeptic)**
**"Two 'clocks' means two dark-matter species — are you double counting?"**

*Response.* One condensate, two coordinates on it (the dual of $H_3$ and the Wilson line), not two fields. The frame 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 whichever frame describes it. "Two clocks" names the gearing between the coordinates, not a population count.

**Objection 15 (O-Kaxion-15) (the renormalization skeptic)**
**"Kinetic mixings renormalize — why should $1/35$ survive quantum corrections?"**

*Response.* Because the gearing is a geometric overlap integral on a rigid manifold (Appendix), not a coupling constant with a renormalization-group flow. Topological and geometric data do not run; the gearing $\varepsilon = 1/K$ is a ratio of OGN-5 integers, and the leading correction is $(5/8)K^{-2} \simeq 5 \times 10^{-4}$, a factor of two inside the falsifier's stated $\pm 10^{-3}$ tolerance on the split. Setting $\varepsilon = 0$ is not a small renormalization but the deletion of BI-6 itself.

**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 $\mathrm{U}(1)_{\mathrm{PQ}}$ 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 38 traces the usage), the phase coordinate produces winding, holonomy, and hence the Kaxion itself (Theorem 8.2). Call them what one likes; the $1/35$ split is their signature, and it is measurable.

**Objection 18 (O-Kaxion-18) (the hierarchy skeptic)**
**"Towers at $10^{16}$ GeV, locks at TeV, zero-modes at $\mu$eV — 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 the three scales are three *readings* of the same throat: $e^{-2\pi}$ per unit depth, depth $K = 35$, and topological suppression $e^{-K}/2^J$ for the zero-mode (Section 9.5). The thirty-one orders are not thirty-one knobs but one integer exponentiated — 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) and the $\mu$eV doublet split are both $K = 35$ read aloud, and 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-BI6 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 exactly two closed one-form directions to wind: the compact phase circle (the Wilson line) and the dual of the three-form flux through the throat (the tensor shadow) — Section 9 builds both explicitly. A third would require a third independent cycle, which the topology of $\mathcal{M}_4 \times I_r \times S^1_\theta$ does not possess; a single one would require the Green–Schwarz two-form to vanish, which the anomaly forbids. This is the sense in which the doublet is a *census*, not a model choice: the 1/35 pair counts the bulk's cycles from the boundary, and finding a singlet (F-Kaxion-1a) or a triplet at the window would each falsify the counting outright.

**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 33), 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 a prediction: a line confirmed *between* rungs — at, say, 1.87 GHz — would falsify the ladder just as surely as silence everywhere falsifies the window. Dials cannot be caught between positions; ladders can.

# I 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 *not* a dial fitted to the answer: the decay constant is fixed by the geometry, and the abundance then *measures* the angle at $\theta_i \simeq 1.4$ rad — the generic $\mathcal{O}(1)$ value every standard axion-dark-matter scenario assumes, no more tuned than that. The frequency does not depend on it at all: $\nu_K$ is set by $f_K$, so even a sizeable error in the effective $\theta_i^2$ leaves the line where it is and moves only the inferred angle, within 1.1–1.5 rad. And the alternative history is not hidden but analyzed and *excluded* within the series (Part IV): the post-inflationary randomized-angle history fails because there is no Peccei–Quinn scalar to randomize and reheating is nearly nine orders too cold to restore one.

**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, or a compact-angle line elsewhere, refutes it — 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 fixed — and it is GG-A10-Cosmos’s, and it is load-bearing in the favorable direction. It is exactly what forces the single-patch (pre-inflationary) history that gives the clean abundance closure and the runtime-selected $\theta_i$; and with the lock-epoch scale $H_2 \simeq 3.5 \times 10^{-12}$ GeV the axion isocurvature is $\beta_{\rm iso} \sim 6 \times 10^{-40}$, thirty-eight orders below the Planck bound. The unusual choice is not a liability bolted on; it is the reason the prediction is unique.

**Objection 25 (O-Kaxion-25) (the fine-tuning cosmologist)**
**“An initial angle of 1.4 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 *derives* the abundance target independently ($\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$, a theorem of GG-A10-Cosmos, tested against the CMB to 1.5%) and then reads the angle off as the unique closure of an over-determined system (Section 15.3). Had the closure returned

$\theta_i = 0.03$ or $3.10$ — the tuned edges — the construction would be in trouble and we would have said so; it returned 1.4, the flattest, most generic value an angle can take. 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 margin is what makes this robust rather than delicate: safety requires $H_2 \lesssim 2.8 \times 10^7$ GeV and the derived crossing sits at $3.5 \times 10^{-12}$ GeV — *nineteen orders* of headroom (Section 13.5). $H_2$ is not adjustable to this paper: it is $M_{\rm lock}^2/(\sqrt{3}\,\bar{M}_{\rm Pl})$ with $M_{\rm lock}$ the same 3.85 TeV that GG-P1-KK stakes at the LHC, so any attempt to move it collides with a collider falsifier, not with a cosmological convenience. One would need to raise the lock scale by nine orders of magnitude to endanger the isocurvature budget — at which point the theory is already dead several other ways. Fragility is when a conclusion sits near its cliff; this one cannot see the cliff from where it stands.

**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 by force — no Peccei–Quinn scalar exists to randomize the angle across horizons, and reheating is nearly nine orders too cold to restore one (Section 36.2) — so the violent minicluster spectrum of post-inflationary scenarios is absent by theorem, not by preference. 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 $\pi$-locked static inventory is on the wrong side of the data."**

*Response.* The $\pi$-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 theorem 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 $\pi$-locked share) is named, flagged, and carried as one of the paper's two honest fences (Table 13), and it fails visibly if the laboratory line is confirmed but the implied local density falls short of $\rho_{\rm loc}$. 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)**
**"ΛCDM 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 ΛCDM is tested, the Kaxion *is* ΛCDM'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 derived frequency can tell us what it was made of.

## J 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; DFSZ lines are singlets while every Kaxion line is a 1/35 doublet; DFSZ has no gauged baryonic channel while the Kaxion's two-port ratio is locked at 25/4. A signal passing all three tests is not "a DFSZ axion" but a measurement of $K = 35$ and of a gauged baryon symmetry.

**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 $\bar{\theta} \to 0$. 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 $\mathrm{U}(1)_B$ requires a TeV $Z'_B$; where is it?"**

*Response.* It is not missing; it is a separate, pre-registered prediction. Gauged $\mathrm{U}(1)_B$ requires a leptophobic vector, and the sister paper P1-KK places it in the radial lock cluster at $\sim 2.46$ TeV, with top-rich, leptophobic decays testable in existing and forthcoming LHC data. The cohort is deliberately cross-exposed: the same gauge factor that gives the Kaxion its 5:2 portal 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-glue 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 $\chi(T)(1 - \cos\theta_{\mathrm{eff}})$ 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.4). The gauged sector adds no new $\theta$-source: $\mathrm{U}(1)_B$ $FB\widetilde{F}_B$ 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 (open item iii), 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 inversion is a change of basis, and it is written out where the auditor can check every line: the anomaly fixes the *source* ratio of the two-form's couplings to the two gauge sectors ($2/5$, GG-A3-BI6's coefficient vector); dualizing the two-form into the pseudoscalar exchanges the roles of field and source, which inverts the ratio in the *response* ($5/2$); squaring for power gives $25/4$ (Section 32 and Appendix C.1 carry the four lines of algebra). Nothing is chosen and no sign is discretionary — and the audit has an empirical form: the same dualization fixes the relative phase of the two channels, so a measured ratio of $25/4$ with the *wrong* coherence would still fail F-Kaxion-1b. Gymnastics cannot survive a phase-sensitive referee.

**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 $25/4$ is not a DFSZ electron coupling and does not descend through the DFSZ loop chain: it is the ratio of the *gauged* baryonic portal to the electromagnetic one, locked at the level of the dualized two-form's couplings before any fermion-loop question arises (Theorem 32.1). What the loop skeptic is right about is the step from portal to *instrument response*: converting the gauge-level ratio into a magnon-port calibration requires the chiral matrix element, and that is exactly the fenced model-level input the paper refuses to oversell — it is open item (iii), it bounds the first-generation calibration at $\pm 20\%$, and the R&D year of Part X exists to measure through it (Table 30). The lock is tree-level and exact; the honesty about reading it out is priced and printed.

**Objection 38 (O-Kaxion-38) (the topological-defect expert)**
**"Two coupled pseudoscalars generically produce domain walls; $N_{\mathrm{DW}} = 1$ by assertion does not remove them."**

*Response.* $N_{\mathrm{DW}} = 1$ is not asserted but computed — the unit color anomaly of the single portal (Theorem 24.1) — and unit domain-wall number means walls bounded by strings that collapse harmlessly *if* strings form at all. Here they do not: defect networks require the field to decorrelate across horizons, and the single-patch clause forecloses that route twice over — no Peccei–Quinn restoration (no scalar exists to restore) and reheating ten orders below the decay constant (Section 36.2). The two-pseudoscalar structure adds no loophole, because the heavy partner is integrated out at the cap scale long before any cosmological era that could paint it non-uniformly. The defect story is thus a theorem-chain, and it is testable at the margin: a confirmed minicluster or wall-crossing signature would falsify the same clauses (see the substructure objection above).

# K 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, on its

1.8–10 GHz magnet and with DFSZ depth already demonstrated above 1 GHz, and ADMX-EFR, with its eighteen-cavity array at 2–4 GHz, are the funded programs converging on exactly this window at DFSZ depth over the late 2020s (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 $1/35$ doublet retires the Kaxion outright (F-Kaxion-1) — and, because the Kaxion is the dark-matter claim of the series, it retires that claim with it. The prediction is pre-registered down to the frequency, the doublet spacing, and the channel ratio; a falsifier that specific is the mark of a real prediction, not an escape hatch.

**Objection 41 (O-Kaxion-41) (the spectroscopist)**
**"Can the $56$ MHz doublet split be resolved?"**

*Response.* Resolving it is trivial — the split is $\sim 3 \times 10^4$ condensate linewidths, far wider than any instrumental resolution at 1.95 GHz. The doublet is therefore not a resolution challenge but an *identification* resource: a dual-mode receiver accepts a candidate only if both modes ring at the fixed $1 + 1/35$ ratio, which suppresses accidental backgrounds quadratically and lets each mode scan faster and shallower than a discovery-grade single-mode search. The pair is more sensitive *and* more diagnostic than the sum of its parts.

**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 — produces one line in one channel. The Kaxion produces *two* lines at a fixed $1/35$ ratio and drives *two* channels at a fixed $25/4$ ratio. No background conspires to do both at fixed fractional spacing and fixed channel ratio, so the double coincidence 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, and at 1.95 GHz and $\sim 10$ mK their dark-count floor, not the standard quantum limit, sets the noise. That floor is the dominant scan-rate lever, buying the factor that DFSZ-depth scanning otherwise pays in years. 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 same theorem that fixes the abundance: the $\pi$-locked closure makes the Kaxion saturate $\Omega_{\rm DM}$ — the whole inventory, not a fraction — so the density entering the signal is the full local $\rho_{\rm loc} \simeq 0.45$ GeV/cm$^3$, with no dilution by a second component. And even if one doubted the local density, the 5:2 identification test is independent of $\rho_{\rm loc}$ by construction, since the matched ratio cancels it.

**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 27). The honest statement is quantitative: real camps carry 10–100× duty and calibration overheads over the idealized times, and every quoted timescale in Parts VII–X says so explicitly. Multiply the $5\sigma$ 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 ±6% band is 240 MHz; '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 240 MHz band costs of order a month of single-cavity live time (hours for an array-class instrument) at depths a blind GHz-scan buys with years. And the band is not the prediction's core: the discrete ladder gives seven marked rungs, the benchmark is one bin, and the ±6% is a Standard-Model-side error that lattice progress shrinks independently of any instrument (Section 27). Call it a short scan if preferred; 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 $5\sigma$ times double from seconds to seconds and from minutes to minutes (Table 17), the rung exclusions of Figure 23 rise from $g_{\rm DFSZ}/31$ to $g_{\rm DFSZ}/22$, and the F-Kaxion-1c dates move by weeks against a schedule quoted in years. The density enters as $\sqrt{\rho}$ in amplitude and is already carried at ±22% as the budget's dominant systematic (Table 21); a factor-two local deficit is a $3\sigma$ surprise to the halo community before it is a problem for this paper. And the identification observables are engineered to be immune: the 1/35 spacing and the 25/4 ratio cancel $\rho_{\mathrm{loc}}$ exactly, 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 F-Kaxion-1a '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 second, pre-registered coordinate. The day-one claim is arithmetic, not confidence: at camped signal-to-noise the partner bin reaches decisive depth in under an hour of live time (Table 17), so the calendar cost of F-Kaxion-1a 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 that a candidate has, or lacks, a partner at a named offset is a different and far faster epistemic task. Proving a negative is slow; checking a printed coordinate 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 (§66). 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 moves seven times in the first two months (Table 33) and then effectively never; 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 35), 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.

## L 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 is the mature first stage (Part VIII). The objection is answered by the readout, not evaded.

**Objection 52 (O-Kaxion-52) (the untestability skeptic)**
**"You computed the 5:2 coupling astrophysically negligible, so it is untestable."**

*Response.* The objection confuses the *absolute* coupling with the *ratio*. The absolute dark-baryon coupling is indeed too weak to move a galaxy — we computed it (§75.4) — but the harmonic booster does not measure the absolute coupling; it measures the matched-channel power ratio $P_B/P_\gamma = 25/4$, in which the absolute magnitude, the local density, and every frame-dependent factor cancel. The galaxy center is where the coupling fails to act; a resonant, matched two-port receiver is where the locked ratio is read. Astrophysical silence and laboratory testability are perfectly consistent.

**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 $(\ell_{\mathrm{lock}}/r_s)^2$, 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 $\Xi_{\mathrm{cap}}$ shifts."**

*Response.* The prediction is robust to exactly this. The window is bracketed by funded programs from both edges, so a shift of the benchmark within the discrete $\Xi_{\mathrm{cap}}$ range moves a *marked row* inside a pre-committed window, not the window's existence; and the 1/35 spacing is fixed regardless of where the center lands. Decisively, a confirmed compact-angle-like line at *any* other frequency also falsifies the Kaxion (F-Kaxion-1d), so a wrong location is as fatal as a missing doublet — 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 is 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. (50.1)). The requirement it must meet is $\mathrm{SNR}_B = 5$ per five live days at a known bin — and the paper treats even that as unproven: it is fenced as a design requirement, gated behind a bench year that precedes construction, with the $\gamma$-only fallback preserving the discovery and kill programs in full (Table 35). A chasm one refuses to price is a fantasy; a gap with a measured width, a gate, and a fallback is called engineering.

**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.* 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 $\pm 0.1$ mT, which is a $10^{-3}$ trim on the bias coil, well inside routine magnet-lab practice (§59). 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 *quantitative* head-start layers are quoted with their model-level status marked, 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 3% boost is indeed invisible; 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 forty-plus orders in cross-section — discriminations that no simulation systematic can blur. Where percent-level cosmology is genuinely at stake (the $\pi$-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.

## M 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, these nineteen consequences 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, and two named conjectures — 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 nineteen 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 missing 1/35 partner, 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 executions). 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 confirmed line without its 1/35 partner, a matched two-channel ratio off 25/4, an empty lock cluster at 3.85 TeV, a CMB inventory off the $\pi$-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's integers leave fingerprints — 1/35, 25/4 — 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 ($\theta_i$'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 2.857% frequency split 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 F-Kaxion-1a confirms the split, 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 $\pi$-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. $\pi$-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 1/35 split, the 25/4 ratio, and the frequency window. 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-1a, 1b, and 1d are built from dimensionless locks that hold whatever the conjectures, the lattice, or the benchmark do; if one of them fires, the geometry itself takes the bullet, with no auxiliary premise available to absorb it. F-Kaxion-1c strikes the closure sheet exactly as fenced: it retires the window the two named conjectures produced, and is booked against them. 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 doublet or the ratio, so they are never candidates on 1a/1b/1d. And the two named conjectures are the only deniable elements in the chain — 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 eighteen sister papers of the GG-Theory cohort, in the canonical order of the program (the constitutional root, the GG6 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://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://arisaka-gg.org/`.
**Relevance.** The constitutional charter of GG-Theory: it states the four axioms (locality/causality LC, the geometry–gauge equivalence GGEP, its open-system extension OGGEP, 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 oscillation invariant $\Delta\nu/\nu = 1/35$ as a constitutional falsifier with a pre-declared kill criterion — 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://arisaka-gg.org/.
**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 $S^1_\theta$ — the dimensional-minimality theorem forcing $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$, 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 Wilson holonomy that is one of its two clocks, and the throat depth $K = 35$ that sets the 1/35 split — 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://arisaka-gg.org/.
**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 $\mathrm{d}H_3 = \mathrm{tr}\, R^2 - \mathrm{Tr}\, F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Qi}^2$, and the exact trace-ratio computation $\mathrm{Tr}(B_0^2)/\,\mathrm{Tr}(Y_0^2) = 2/5$. 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 5:2 portal lock 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://arisaka-gg.org/.
**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://arisaka-gg.org/.
**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://arisaka-gg.org/`.
**Relevance.** The runtime paper: it defines GRIP and the Morse free-energy landscape $F_{\mathsf{GFF}} = F_{\mathsf{GEP}} + F_{\mathsf{CEP}}$ (GFF-ST2) and proves equilibrium uniqueness $[\Pi_*] = \arg\min F_{\mathsf{GFF}}$ (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 (DGI $\to$ GSSB $\to$ 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://arisaka-gg.org/`.
**Relevance.** Contains Theorem T-QUA, the Arisaka equation $J_E/\mathsf{Q_B} = J_e = J_B = J_{Qi}$ with $\mathsf{Q_B} = (2/5)E_{\mathrm{Ry}}/K$, 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 5:2 baryonic-to-electromagnetic portal ratio — the 25/4 matched-channel power lock that the harmonic booster of Part VIII is built to read — is precisely the dark-sector reading of this current lock. 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://arisaka-gg.org/`.
**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://arisaka-gg.org/`.

**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 $\alpha^{-1}(M_Z) = 2^J = 128$ (the lock-scale coupling that enters $g_{K\gamma\gamma}$), the weak-angle rational partition $\sin^2\theta_W = 25/108$, the Wilson–Hosotani vacuum-value theorem with the scale chain $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock} = \frac{1}{2}E_{\rm Pl}e^{-K}$, and the four-force uniqueness theorem that forces the extra gauge factors $\mathrm{U}(1)_B$ and $\mathrm{U}(1)_{Qi}$. 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://arisaka-gg.org/`.
**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 $\pi$-locked partitions $\Omega_m = 1/\pi$, $\Omega_b = 1/(2\pi^2)$, and $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2) = 0.2676$, 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 two-direction (Green–Schwarz + Wilson-line) 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 three-layer early-structure mechanism on which Section 21 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.
In preparation. To be posted at `https://arisaka-gg.org/` on completion; 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://arisaka-gg.org/`.
**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 $M_{\rm lock} \approx 3.85$ TeV, including the leptophobic $\sim 2.46$ TeV vector $Z'_B$ that gauges the very baryon number whose dark-sector reading is the

Kaxion's 5:2 portal. 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 $\mu$eV 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://arisaka-gg.org/`.
**Relevance.** The sibling Predictions-Branch paper that spends the same OGN-5 ledger on the generation count: $N = 3$ 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 $N = 3$ 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://arisaka-gg.org/`.
**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://arisaka-gg.org/`.
**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://arisaka-gg.org/`.
**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://arisaka-gg.org/.
**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://arisaka-gg.org/.
**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 Wilson-line phase 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://arisaka-gg.org/.
**Relevance.** The classification of protection mechanisms that this paper occupies one rung of. The compact phase circle whose holonomy the Kaxion is, and 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://arisaka-gg.org/.
**Relevance.** The Cosmology-Branch sibling on the expansion side of the same $\pi$-locked inventory from which this paper's abundance target descends: the derived $H_0$ and the $\Omega$ 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 OGGEP Port — with the Exact Bekenstein–Hawking $A/4G_N$ Derived from Induced Gravity. Arisaka-GG-C3-BH.
Preprint, available at https://arisaka-gg.org/.

**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 $\sim 2 \times 10^{-48}\,\mathrm{cm}^2$ 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.7.

**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 the claim of this paper — that two funded programs will reach the Kaxion window at the required depth by roughly 2030 — should be read. First cited in §36.2.

**ADMX Collaboration (2024).** The ADMX Extended Frequency Range program: DFSZ sensitivity over 2–4 GHz. Design report / Snowmass contribution.
**Relevance.** The first announced DFSZ-depth program overlapping the Kaxion window from above (8.3–16.5 $\mu$eV), 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 $\mu$eV 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 candidate validation near 1.036 GHz. *arXiv:2602.05388*.
**Relevance.** The field's textbook candidate-validation episode, contemporaneous with this paper: a 5.1$\sigma$-local excess at 1.036315 GHz (reduced to 3.5$\sigma$ 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.

**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 $\sim 2$ 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 $\mu$eV band lies above the Kaxion's laboratory line at 8.1 $\mu$eV (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 $\chi_{\rm QCD}(T)$ across the temperatures relevant to axion production. It fixes both the zero-temperature normalization that sets $m_K f_K = \sqrt{\chi_{\rm QCD}}$ 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 §9.4.

**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 ΛCDM 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 haloes 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-m$^3$: a search for the QCD axion below $1\,\mu$eV. *Physical Review D*, *106*, 103008.
**Relevance.** The lumped-element program targeting DFSZ sensitivity below $1\,\mu$eV, where cavities become impractically large. While the Kaxion line lies above this band (at $8.1\,\mu$eV), the program illustrates the lumped-element technique and is a natural host for a nucleon-spin booster, 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 ΛCDM 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 §40.4.

**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 4.5–5.4 $\mu$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 funded programs will 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 harmonic booster's electron port and the model for the matched-ratio readout that tests the Kaxion's 5:2 lock in Part VIII. First cited in §40.2.

**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 0.176–0.203 $\mu$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 haloes, 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 $\Omega_{\rm DM}$, 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.7.

**Fischer, M. S., et al. (2024).** Cosmological and idealized simulations of dark-matter haloes 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 $\mu$eV 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 the funded programs 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.6.7) and behind the epoch-dependence of the $\pi$-locked fractions (Appendix C.8); a student following the abundance and structure-formation discussion should keep it open. First cited in §C.8.

**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 calibrate the $\pm 8\%$ accuracy assigned to our analytically transparent abundance pipeline — the honest source of most of the Kaxion mass error, and the target of the lattice program that would compress it. 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 $m_a = 5.70(7)\,\mu\mathrm{eV}\,(10^{12}\,\mathrm{GeV}/f_a)$ 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 holonomy that drives the Hosotani breaking of SU(5) supplies the Kaxion's second compact-angle direction — the holonomy "echo" $\theta_W$ of Part II — so this paper is the origin of one of the Kaxion's two geared clocks, 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 funded programs that bracket the Kaxion window (CAPP, ADMX-EFR) 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 $\mathrm{cm}^2/\mathrm{g}$) 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 with the additional 1/35 and 5:2 structure) 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 §21.11.

**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 §21.11.

**Labbé, I., et al. (2023).** A population of red candidate massive galaxies $\sim 600$ 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, doublet-aware 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.7) and the thermal-history inputs of the closure audit. First cited in §C.8.

**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 $\mu$eV 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 9, 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 §77.

**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 $\Xi_{\text{cap}}$ into an integer (Appendix C.6), and the Hopf fibration behind the $\pi$-locked inventory (Appendix C.8) are all developed here at book length. First cited in §C.5.

**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 haloes 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 9 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.4) and for the anomaly-coefficient computation behind the 5:2 baryon portal (Appendix C.5). First cited in §C.5.

**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 $\Omega_{\rm DM}h^2 \simeq 0.120$ to about one percent, and the tight bound on dark-matter isocurvature that the Kaxion must respect. This is the measurement the theory's $\pi$-locked target $\Omega_{\rm DM} = 1/\pi - 1/(2\pi^2)$ 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 $\pi$-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.7.

**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: $m_{\rm FDM} \gtrsim 5.6 \times 10^{-22}$ eV for fuzzy and $m_{\rm WDM} \gtrsim 1.5$–$3.2$ 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-$\alpha$ forest. *Physical Review Letters*, *126*, 071302.

**Relevance.** Uses the small-scale clustering of intergalactic hydrogen (the Lyman-$\alpha$ 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 disc — 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 $\sigma/m$ 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 parameter-free 5:2 dark–baryon ratio — which 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 §40.2.

**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 $\theta$-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, the two fingerprints (1/35, 5:2) 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 $\theta$-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 5:2 portal (Appendices C.4, C.5). First cited in §C.5.