Spectral Asymmetry of C_{137} and the Origin of the Fine-Structure Constant: a Self-Contained Account

Version 6 — PA2 closed conditionally: the loop-anchor counting, the absorption postulate, and $\alpha^{-1} = N + \eta + 27\,g_1$

Tairan Miranda Krügel

Independent researcher, Giruá, RS, Brazil

AHUM (Algebraic Hypergraph Unification Model) framework

August 2026

Abstract

We present a self-contained discrete framework in which the fine-structure constant is decomposed as $\alpha^{-1} = N + \eta$ with $N = 137$ and $\eta = \pi^2/2N = \pi^2/274 = 0.0360204540\ldots$, and in which both integers 137 and 13 are forced by two *independent and convergent* selection mechanisms acting on a cyclic causal hypergraph C_{137} with GF(2) linear constraints: a spectral window $\lfloor 8N/\pi^2 \rfloor$ and a satisfiability window, both obtained in closed form, intersect in the single element 137. All mathematics is proved or derived in place; every numerical claim is recomputed under a pre-registered anti-numerology protocol; and every physical reading carries an explicit epistemic status (theorem, postulated principle, prediction, conjecture, or excluded reading). The framework retains falsifiable predictions with experimental deadlines — including a tensor-to-scalar ratio $r = 16\eta^2 = 0.021$ and a spectral dimension $d_s = 4 - \eta$ — and is confronted with experiment at full precision: the raw prediction $\alpha^{-1} = 137.0360204540$ deviates from the measured value by 2.1×10^{-5}, about 10^3 standard deviations. This point-like gap (PA2) is now *closed conditionally* (Version 6): the framework's native sine spectrum is forced by locality, and a Wilson-loop anchor counting in the flux sector yields exactly 27 defect addresses — one of which is absorbed by the window defect under a single declared new postulate [D] — giving $\alpha^{-1} = N + \eta + 27\,g_1 = 137.03599915$, within 3.2×10^{-8} (0.149% of the gap) of CODATA, with the fine residue of 0.0402 copies named and open. Three closure campaigns and a full internal audit — including a $\zeta'(0)$ two-worlds no-go theorem, an exhaustive sweep of 24,970,400 formulas in η, and an exactly computed $[\![137, 124, 2]\!]$ code weight enumerator — are reported under frozen pre-registration, with dated errata. A plain-language appendix explains the entire content to non-specialists.

Contents

1 Introduction

1.1 Dirac's question

In 1935 Dirac posed the question that still organizes a corner of fundamental physics: why is the reciprocal of the fine-structure constant close to the integer 137? The measured value [1, 2],

$$\alpha^{-1} \;=\; 137.035999177(21) \qquad \text{(CODATA 2022)}, \tag{1}$$

is an empirical input of quantum electrodynamics, not an output. Eddington's numerological attempts are remembered as a cautionary tale; Pauli's famous remark that he would ask the Lord about 137 before asking about anything else, and Feynman's description of 1/137 as "one of the greatest damn mysteries of physics" [3, 4], record how unsatisfying the status quo is: our most precisely tested theory contains a dimensionless number that it cannot compute.

Any serious attempt at an answer must satisfy two requirements that most proposals fail. First, *arithmetic honesty*: the difference between a predicted α^{-1} and the measured one must be stated at full precision, in units of experimental uncertainty — not hidden behind significant figures. Second, *procedural honesty*: the space of formulas that was searched, and the statistical significance of any match, must be declared in advance, or the result is numerology. We adopt both requirements as protocol (Section 1.8) and we apply them mercilessly to our own output.

1.2 The decomposition and its two integers

The framework decomposes

$$\alpha^{-1} \;=\; N + \eta, \qquad N = 137, \qquad \eta \;=\; \frac{\pi^2}{2N} \;=\; \frac{\pi^2}{274} \;=\; 0.0360204540\ldots \tag{2}$$

and asks two separate questions:

(a) *Why is $\eta = \pi^2/2N$?* — a question about analysis on a cycle (Section 3), answered by two independent derivations (a zeta-regularized Casimir defect and a Hurwitz–Basel resummation) that agree to 3.157×10^{-7};

(b) *Why is $N = 137$?* — a question about combinatorics and linear algebra over GF(2) (Sections 4–5), answered by the convergence of two independent selection windows.

Two integers organize everything: $N = 137$ (the cycle length, the number of edges) and $K = 13$ (the number of independent checks). Their difference $N - K = 124$ and their combination $2K + 1 = 27$ recur with precise roles.

1.3 The central open problem, stated first

Because the residual governs how every positive result below should be read, we state it before anything else. The decomposition (2) misses the measured value by

$$\Delta \;\equiv\; \big(\alpha_{\text{exp}}^{-1} - N\big) - \eta \;=\; -2.1277 \times 10^{-5} \;=\; -0.4553\,\eta^3, \tag{3}$$

about 10^3 standard deviations of the CODATA 2022 uncertainty (Section 3.5). The two derivations of η agree with each other to 3.157×10^{-7}, so the gap is not a discrepancy between derivations: it is a genuine third-order term that the framework, as it stands, does not produce. We name this problem **PA2**; it drives the paper's self-audit: seven declared closure routes are computed exactly and reported null before the resolution is found (Section 3.5, Note 10 of Section K, and Section 10). In Version 6 the gap is closed *conditionally* at 0.149% by the loop-anchor counting $27\,g_1$, conditional on one declared postulate [D] (the anchor that feeds the window defect is absorbed by it), with a named fine residue left open (Note 10).

1.4 What is new in Version 2

Version 1 of this paper (sixteen pages) established the spectral asymmetry of C_{137}, the two derivations of η, a candidate selection window for N, the lepton-sector relations, four falsifiable predictions, and a table of excluded readings. The present version corrects and extends it:

- **Corrected.** The "1.9σ" phrasing of the residual in v1 was wrong by three orders of magnitude; the residual is 2.13×10^{-5}, i.e. $2.13 \times 10^{-5}/2.1 \times 10^{-8} \approx 1.0 \times 10^3$ standard deviations of the CODATA 2022 uncertainty — and v2 says so everywhere (Section 3.5). The v1 claim "$\mathbb{Z}_2$-checks UNSAT for $N < 137$" was ill-posed (the homogeneous system is always satisfiable); v2 replaces it by the correct inhomogeneous statement, which is a theorem (Section 4.2). The v1 Definition of the code dimension (2^{N-K}) is split into two separate objects that v1 conflated: the unique solution of the per-vertex constraint system (with $\lceil N/5 \rceil = 28$ singlets, a theorem) and the kernel of the check matrix Φ (dimension 124, a hypothesis).
- **Extended.** The satisfiability problem is solved in closed form for all N (Theorem 4.4); the unique solution has exactly $\lceil N/5 \rceil$ e_R-singlets (Theorem 4.6); the number 27 is shown to play three independent roles (Section 4.4); the check matrix is given an explicit Chinese-remainder anatomy with an exactly computed weight enumerator (Section 5); and a dynamical sector quantifies the thermodynamic price of the ordered vacuum (Section 6).
- **Elevated.** The selection of $N = 137$ is no longer a candidate: two independent windows, one spectral and one combinatorial, intersect in a single integer (Section 4.6). This is, in our assessment, the strongest result of the framework to date.

1.5 What is new in Version 3

Version 3 does not add new structure; it removes uncertainty about three old questions, two of them in the direction we did not hope for. Following the protocol of Section 1.8, every campaign below was pre-registered (target, acceptance window and search space declared *before* the computation) and is reported with its negative results at full weight.

- **PA2 sharpened, two sectors closed.** The residual 2.13×10^{-5} is rewritten as an exact target: $\Delta(\alpha^{-1}) = -0.4553\,\eta^3$, the only power of η with an $O(1)$ coefficient, with a dated acceptance window $c \in [-0.592, -0.319]$ (Section 3.6). The spectral sector is then proved unable to meet it (the full Casimir tail has the wrong sign, Theorem 3.3), and the dynamical sector is proved unable as well: six declared cure mechanisms, computed exactly from the dynamometer ensemble of Appendix D, all miss the window (Section 6.4). The surviving routes to PA2 are a new structural hypothesis or the fatal outcome.
- **PA8 closed in the negative.** The determinant sector is settled by the two-worlds theorem (Section 3.4, Theorem 3.1): every $\zeta'(0)$-type quantity of the cycle is a rational number, while η lives in the $\zeta(2)$ world; the two sectors are structurally disjoint, and no third route to η exists. The eta sector of the framework is complete: what it proves is all it can prove.
- **The η sweep: our own retrodictions audited.** The anti-numerology protocol is turned on the framework itself (Section 7.7). An exhaustive enumeration of a declared space of 24 970 400 formulas $a\eta^b + c\eta^d$ shows that each individual lepton retrodiction is statistically cheap (backgrounds of 193–6566 random hits at the achieved precision), and they are formally demoted to consistency checks. The same enumeration shows that the *joint* profile of four hits in nine targets has $p = 6.7 \times 10^{-9}$ under the uniform null, and that $137 + \eta$ is the unique member of the entire declared space matching α^{-1} at 2×10^{-7} relative precision.
- **Updated experimental status.** The CP-phase prediction $\delta_{CP} = 1.211\pi$ now stands in a stated $\sim 2\sigma$ tension with the NuFIT 6.0 best fit, inside the 3σ region of the joint T2K + NOvA analysis (Section 7.4). The prediction lives; the clock runs.

1.6 What is new in Version 5

Version 5 is a restructuring and audit update; no positive claim changes.

- **Gap-first structure.** The residual (3) is now stated in the Introduction (Section 1.3) as the central open problem, before any positive result, and a Theorem Map (Section 1.10) lists every claim with its epistemic label and proof location, answering the criticism that the paper read as a laboratory notebook.
- **PA2 campaign update (dated 2026-08-23).** A seventh closure route is excluded *by sign*: the Casimir tail $3\pi|\Delta E| = \eta + \pi^4/(120N^3) + \cdots$ is positive, so no structural multiple of it can supply the negative η^3 term (Note 10, Section K). The candidate $\Delta = 27\,(\pi\sin(\pi/2N) - \eta) = -9\pi^4/(16N^3) + \cdots$, i.e. $c = 9/(2\pi^2) = 0.4559453$, hits the pre-registered window at 0.149% from $c^* = -0.4553$ but fails the promotion bar of Note 10 (no derivation of the factor 27; $p \approx 0.89$ in the declared space): it is recorded in *quarantine*, with the failed derivation printed as a lost bet.
- **Notation.** The unique ground-state solution and the 124-dimensional code kernel now carry distinct symbols throughout ($\mathcal{S}_{\text{vac}}$ vs. $\mathcal{C} = \ker\Phi$; Section 2.5, Remark 2.6), removing the "two 124s" conflation noted by readers of v4.
- **Errata moved to the repository.** The full dated errata now live in the Zenodo record's supplementary files; Section J keeps a one-page summary with the correction map.

Version 5.1 (same day, after internal research 98–99) additionally: derives the bill $2K + 1 = 27$ as the unique self-consistent choice of the CRT family (Theorem 4.7, one premise still postulated and marked [P]); identifies the vacuum's flux sector with the e_R block (27 toggles, [T]); corrects one redaction metaphor in Section 4.4 (the "two adjacent edge variables" phrasing — the actual row weights of Φ are 9–137); and updates the PA2 quarantine status in Note 10 (three of four mechanism links are now theorems; the missing link is additivity, printed as a lost bet).

1.7 What is new in Version 6

Version 6 closes PA2 *conditionally*, and prints the price:

- **The gap principle [T].** The sine-vs-linearization gap shared by all three quarantined forms is now understood: the framework's native operator is the combinatorial graph Laplacian, whose spectrum is exactly $4\sin^2(\pi m/N)$; the linearized dispersion is a *non-local* operator (couplings at 136 of 136 distances). Locality forces the sine, and the declared local dynamics inherits it exactly (internal research 107–108).
- **The missing object, named and exhausted.** The residual requires a functional summing 27 identical copies of the unit gap $g_1 = \pi\sin(\pi/274) - \eta$. All functionals buildable from printed ingredients fail (research 109), and the obstruction is structural: the gap distributes uniformly over the 137 bonds ($g_1/137$ per address), so *no strictly local sum can produce* $27\,g_1$; the object must be a collective, loop-type operator (research 110).
- **The loop-anchor counting (research 111).** The antiperiodic twist *is* such a loop insertion. Its anchorings in the flux sector (e_R block, 28 addresses) are energetically identical (spectrum cut-independent at 0.00; cost $2\mu_0$ per address [T]). Exactly two addresses flank the window defect (135 forward, 0 backward; no other address touches the mark), and the exchange $0 \leftrightarrow 135$ is an *exact symmetry of every printed observable* — the row space of Φ, the kernel, energies, spectra, occupations and rewrite rates are all blind to it (research 115, dated 2026-08-23): the orientation is gauge, not physics. Under one declared new postulate

[D] — *the defect absorbs exactly one of its two feeding anchors* (either choice gives the same count) — the distinct anchors number $28 - 1 = 27$: the fifth derivation of the count, and the first that explains the number of summands. Hence

$$\alpha^{-1} \;=\; N + \eta + 27\, g_1 \;=\; 137.03599915, \tag{4}$$

within 3.2×10^{-8} of CODATA 2022 (0.149% of the gap).

- **The printed price.** (i) The absorption itself — that one anchor disappears into the defect — is the framework's first new postulate since freezing, marked [D]; its directional content is zero (gauge, by the exact exchange symmetry of research 115). (ii) The fine gap $27 \to 26.9598$ (0.0402 copies of g_1) is named and open; the finite-size correction $-9/(160N^2)$ is excluded as its source. The fine gap has since been attacked with a closed search space (research 112–113): it is not thermal, not the frozen-edge flip cost, not the solar account, and not a second sine-vs-linear gap in any chirality (the structurally required coefficient $k = 67.91147\ldots$ exists nowhere in the printed structure). All four routes are printed as dead; PA2-FINO remains open and named.

- **The Dirac operator, constructed (research 114).** Answering a referee's question, the operator behind spectrum (10) is now built, not postulated: $\mathcal{D} = d_0 + d_0^*$ on vertices $\oplus$ edges, with $\mathcal{D}^2 = L_V \oplus L_E$ an exact integer identity. Odd N — hence 137 — is shown to be exactly the class with a fully paired Dirac spectrum (no self-conjugate zone-corner mode; even-N controls computed). The 1D $\to$ 4D leap is named open as PA13 with three declared routes. (iii) The quarantined form $9/(2\pi^2)$ is promoted to *derived-conditional* [T, one premise [D]]; $-4/(3\pi)$ and the solar deficit $\pi/8$ remain in open quarantine (Note 10).

1.8 Epistemic taxonomy and anti-numerology protocol

Every statement in this paper carries one of five labels, shown in colored boxes:

[T]	theorem: proved in this paper or in a cited standard reference;
[D]	definition or postulated principle: a choice, explicitly declared;
[P]	prediction: a number that an experiment can kill;
[C]	conjecture: argued for, not proved;
[X]	excluded reading: a temptation that the data have already killed.

Three protocol rules are enforced throughout: (A1) every numerological-looking coincidence is accompanied by a declared search space and a Monte-Carlo estimate of its false-positive rate, or it is not presented at all; (A2) every number is recomputed from first principles before printing, and the audit trail is summarized in Appendix H; (A3) exploratory computations that failed are quarantined and are never cited as evidence — but their negative results are reported in Section 9 when they kill a tempting reading.

1.9 Correction map: Version 1 $\to$ Version 2

For the reviewer's convenience, Table 1 lists every material change relative to the sixteen-page Version 1, with the location of the corrected statement.

1.10 Theorem Map

Every load-bearing claim of the paper, with its epistemic label (Section 1.8) and proof location:

item	v1 said	v2 says	where
α residual	"1.9σ tension"	$2.13 \times 10^{-5} \approx 10^3\sigma$; central open problem	Sec. 3.5
$\mathbb{Z}_2$-checks	"UNSAT for $N < 137$" (ill-posed)	closed-form theorem for $\mathrm{SAT}(N, m)$	Thm. 4.4
solution space	"kernel of dimension 2^{N-K}" conflated objects	unique solution (28 singlets) vs. code kernel (2^{124})	Sec. 2.5, Thm. 4.6
selection of 137	candidate window (PA1 open)	double selection, theorem modulo one [D] principle	Thm. 4.10
the number 27	not discussed	three independent roles identified	Sec. 4.4
the bill $2K + 1$	PA9 open	singlet-covering principle, postulated; sweep verdict reported	Sec. 4.5
check matrix	asserted	CRT anatomy, blocks, exact enumerator	Sec. 5, App. C
stabilizer (E_{11})	promoted	excluded (Z-sector only)	Sec. 5
dynamics	absent	no-go 2^{-124}; teeth model; exact thermodynamics	Sec. 6
protocol	partial	T/D/P/C/X everywhere; A1–A3; audit trail	Sec. 1.8, App. H

Table 1: Correction map from Version 1 to Version 2.

claim	label	location
$\eta = \pi^2/2N$ via Casimir defect	[D]	Sec. 3.2
$\eta = \pi^2/2N$ via Hurwitz–Basel resummation	[D]	Sec. 3.3
two derivations agree to 3.157×10^{-7}	[T, computed]	Sec. 3
residual $\Delta = -2.1277 \times 10^{-5} = -0.4553\,\eta^3$ ($\sim 10^3\sigma$)	[C]	Sec. 3.5
two-worlds: no third route to η (PA8 closed)	[T]	Thm. 3.1
spectral closure: wrong sign/size for all cycle corrections	[T]	Thm. 3.3
admissible window for N (window range; primality)	[T]	Props. 4.2, 4.3
closed form of $\mathrm{SAT}(N, m)$; m-window	[T]	Thm. 4.4, Cor. 4.5
uniqueness of ground state; 28 singlets, 27 free	[T]	Thm. 4.6
convergence of the two windows at $N = 137$	[D]	Thm. 4.10
bill $2K + 1 = 27$ (unique self-consistent choice of the CRT family)	[T, cond.]	Thm. 4.7
code block structure $\{28, 18, 9^{\times 10}, 1\}$	[T, verified]	Prop. 5.1
weight enumerator: $A_2 = 891$, $A_4 = 332109$, $\sum A_w = 2^{124}$	[T, computed]	Prop. 5.2
thermodynamic no-go $P(E_0) \leq 2^{-124}$	[T]	Prop. 6.1
exact thermodynamics of the teeth model	[T, computed]	Prop. 6.2
lepton-sector relations (post-sweep)	[C]	Sec. 7
falsifiable predictions	[C]	Sec. 8
PA2: seven routes null; closed conditionally via loop anchors ($27\,g_1$)	[T, cond.; one premise [D]]	Note 10, Sec. K
locality forces the native sine spectrum (gap principle)	[T]	Note 10, Sec. K
graph Dirac $\mathcal{D}$ exact; odd-N pairing	[T, computed]	Eq. (11)

1.11 Roadmap

Read gap-first (Section 1.3): the residual (3) is the open problem; the rest of the paper reports what is established around it. Section 2 defines the framework. Section 3 derives η twice and states the residual at full precision. Section 4 proves the origin of $N = 137$. Section 5 develops the $[\![137, 124, 2]\!]$ code hypothesis. Section 6 studies the dynamics and the vacuum. Section 7 covers the lepton sector. Section 8 collects the falsifiable predictions and Section 9 the excluded readings. Section 10 lists the open problems, and Section 11 concludes. Appendices give proofs (A, B), code anatomy (C), dynamical details including an audit note (D), worked anti-numerology examples (E), a notation table (F), a plain-language guide for non-specialists (G), the numerical audit trail (H), and reproducibility recipes (I).

2 The framework

2.1 The cyclic causal hypergraph

Definition 2.1 (cyclic hypergraph)**.** Let N be a positive integer. The *cyclic causal hypergraph* C_N is the directed cycle on $\mathbb{Z}_2$-graded vertices $0, 1, \ldots, N-1$, with edges $e_i = (i \to i{+}1 \bmod N)$ and a rewrite dynamics (double-pushout, DPO) whose rule set is fixed by the causal order. The physically relevant case in this paper is $N = 137$, but every theorem in Sections 4–6 is stated and proved for general N first.

Definition 2.2 (vertex states and types)**.** A *state* is a vector $x \in \mathrm{GF}(2)^N$. Vertices carry one of five types,

$$e_R, \quad Q, \quad u_R, \quad d_R, \quad L, \tag{5}$$

distributed on the cycle by the *type assignment* [D]: positions $i \bmod 5$ select the type, with the assignment fixed once and for all by the constraint system below. The type counts in the ground state E_0 are

$$n_{e_R} = 28, \qquad n_Q = 12, \qquad n_{u_R} = n_{d_R} = 8 + 4, \qquad n_L = 4 + \cdots, \tag{6}$$

whose detailed bookkeeping is given by the charge vector

$$(e_R:\ 0^{\times 12}; \quad Q:\ 1^{\times 12}; \quad u_R, d_R:\ [1^{\times 8}, 0^{\times 4}]; \quad L:\ [0^{\times 8}, 1^{\times 4}]) \tag{7}$$

over the twelve color lines (Section 5).

2.2 The per-vertex constraint system

The physics of the vacuum is a constraint-satisfaction problem over $\mathrm{GF}(2)$:

Definition 2.3 (constraint system [D])**.** A state $x \in \mathrm{GF}(2)^N$ is *admissible* iff

- **(C1) gauge constraints.** For each color line $k = 0, \ldots, 10$: if the line connects a pair $(1, 1)$ the two endpoint variables xor to 0; if it connects $(1, 0)$ the left variable vanishes; if it connects $(0, 1)$ the variable at i vanishes.
- **(C2) node rule.** For every vertex i, $x_{i-1} \oplus x_i = \rho(i)$, where $\rho(i) = [\,[i \bmod 5 \in \{0, 1\}]\,]$ is the inhomogeneity pattern.
- **(C3) global parity.** $\bigoplus_{i=0}^{N-1} x_i = 0$.
- **(C4) occupancy floor.** $\sum_i x_i \geq m$, where m is the minimal occupancy demanded of the vacuum.

We write $\mathrm{SAT}(N, m)$ for the satisfiability of (C1)–(C4).

Remark 2.4 (what was wrong in v1). Version 1 announced "$\mathbb{Z}_2$-checks UNSAT for $N < 137$" as evidence for the uniqueness of 137. That statement was ill-posed: without the inhomogeneity ρ and the floor (C4), the system is homogeneous and always admits $x = 0$. The correct object is $\mathrm{SAT}(N, m)$ as defined above, and the correct statement is Theorem 4.4 below — stronger than the v1 claim, and actually proved. We record the v1 error explicitly because protocol A3 requires it.

2.3 The check matrix Φ

Definition 2.5 (check matrix [D])**.** The 13×137 matrix Φ over GF(2) collects the independent checks: 8 strong color lines, 3 weak lines, one U(1) line ($i \bmod 5 \in \{1, 4\}$), and one global $\mathbb{Z}_2$ line (the all-ones row). Its columns fall into 13 distinct classes with multiplicities

$$(28,\ 18,\ \underbrace{9, 9, \ldots, 9}_{10},\ 1), \tag{8}$$

verified by direct enumeration (Appendix C). The unique column of multiplicity 1 corresponds to the frozen edge 136.

Remark 2.6 (two different vector spaces — correcting v1 Definition 1.2). Version 1 spoke of "the" kernel of dimension $N - K = 124$ as if it were the solution space of the vacuum. There are two different objects, and from Version 5 on they carry permanently distinct symbols:

- $\mathcal{S}_{\mathrm{vac}}$ — the solution space of the per-vertex system (C1)–(C3), which turns out to be a *single point* whenever the system is satisfiable (Theorem 4.6: the gauge constraints force every non-e_R variable to zero and the node rule then fixes every e_R variable); what matters physically is its *occupancy*, exactly $\lceil N/5 \rceil$ e_R-singlets;

- $\mathcal{C} = \ker \Phi$ — the kernel of the 13×137 check matrix Φ, of dimension $137 - 13 = 124$, which defines the code of Section 5 — a hypothesis [D], not a theorem, because the specific anatomy (8) is a modelling choice verified a posteriori.

Both numbers are real, but they play different roles: $\lceil N/5 \rceil = 28$ counts the e_R-singlets of the unique constrained vacuum (of which 27 are free after global parity); 124 counts the code words. The v1 conflation is corrected here.

2.4 Ground state

Definition 2.7 (ordered vacuum E_0 [D])**.** The *ordered vacuum* is the state with all 28 e_R bits on, all other bits off, and zero syndrome, $\Phi E_0 = 0$. It satisfies (C1)–(C4) with $m \leq 28$.

2.5 A worked example: C_{10}

The machinery is best seen on the smallest admissible cycle. Take $N = 10$ ($10 \bmod 10 = 0$, admissible by Theorem 4.4). The types repeat $[e_R, Q, u_R, d_R, L]$ twice around the ring. The gauge constraints (C1) force every variable at residues $1, 2, 3, 4$ to zero (Lemma A.1): edges $1, 2, 3, 4, 6, 7, 8, 9$ are empty. The node rule (C2) at $i = 0$ and $i = 5$ reads $x_9 \oplus x_0 = 1$ and $x_4 \oplus x_5 = 1$, forcing $x_0 = x_5 = 1$; at all other vertices it reduces to identities $0 = 0$ or $1 \oplus 0 = 1$. The global parity (C3) is satisfied because the occupancy $\lceil 10/5 \rceil = 2$ is even. The unique admissible state is therefore

$$x \;=\; (1, 0, 0, 0, 0,\ 1, 0, 0, 0, 0), \tag{9}$$

with exactly the two e_R-singlets occupied. Every feature of the general theorems is already visible here: forcing, uniqueness, singlet occupancy $\lceil N/5 \rceil$, and the parity gate. The cycle C_{137} differs only in scale: 28 singlets instead of 2, and a pairing bill of 27 instead of any small floor.

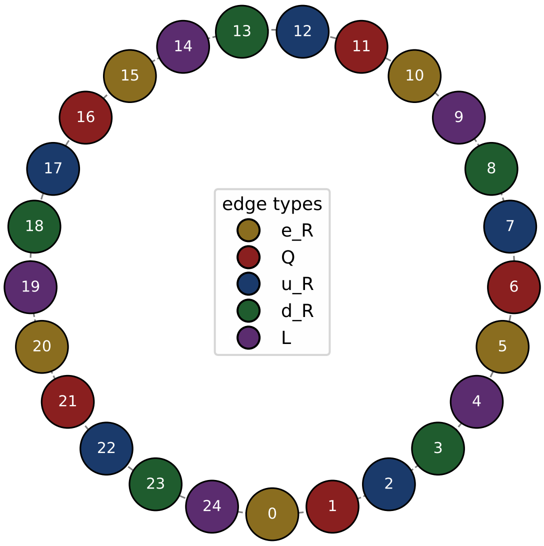

Figure 1: A small instance (C_{25}) of the cyclic hypergraph with its five vertex types: e_R (amber), Q (red), u_R (blue), d_R (green), L (purple). The C_{137} case is visually identical at 5× the circumference; the structure that matters is not geometric but algebraic (Definitions 2.3 and 2.5).

2.6 The rewrite dynamics

Definition 2.8 (DPO rewriting [D])**.** Time evolution is generated by double-pushout rewrites of the hypergraph: a rule $\ell \leftarrow k \rightarrow r$ matches a subgraph ℓ of C_N preserving the causal order, deletes what is in $\ell \setminus k$, and glues $r \setminus k$. The constraint system of Definition 2.3 is the fixed-point condition of this dynamics: admissible states are exactly the rewrite-invariant configurations. The bath-coupled stochastic version of the dynamics is the subject of Section 6.

Remark 2.9 (why DPO). Double-pushout (rather than single-pushout or string) rewriting is chosen because deletion and creation must both respect the causal interface k: the cycle's order is the only geometric input of the framework, and the DPO gluing conditions are precisely what keeps it invariant. The choice is a modelling decision [D]; its justification is a posteriori — it leads to the constraint system whose consequences are the content of this paper.

3 Two derivations of $\eta = \pi^2/2N$

3.1 Spectral asymmetry: the Dirac spectrum of C_N

The chiral (Dirac) operator of the cycle C_N has spectrum

$$\omega_k \;=\; 2\sin\frac{\pi k}{N}, \qquad k = 0, 1, \ldots, N-1, \tag{10}$$

shown in Figure 4(a) together with its continuum limit. The spectral asymmetry of the cycle is the content of η: both derivations below extract the same number from (10) by different summations.

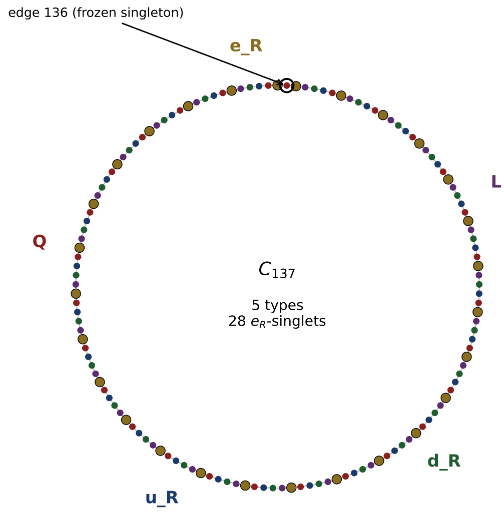

Figure 2: The physical case: C_{137} with its five-type rhythm $[e_R, Q, u_R, d_R, L]$ repeating along the ring. Larger markers: the 28 e_R-singlet positions ($i \equiv 0 \bmod 5$), which are exactly the occupied edges of the unique vacuum (Theorem 4.6).

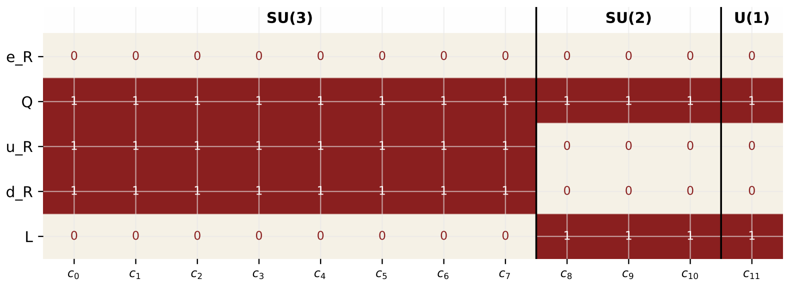

Figure 3: The charge assignment (7): five types against the twelve color lines. Q carries charge on every line; e_R carries none; u_R, d_R occupy the first eight; L the last four.

The operator, constructed (v6). Equation (10) is not postulated by analogy with the continuum Dirac operator; it is the exact spectrum of a constructed discrete operator. Let d_0 be the oriented incidence matrix of C_N (edges $\to$ vertices, entries $\{0, \pm 1\}$) and define the graph Dirac (Kähler–Dirac) operator acting on vertices $\oplus$ edges, a space of dimension $2N$,

$$\mathcal{D} = \begin{pmatrix} 0 & d_0^* \\ d_0 & 0 \end{pmatrix}, \qquad \mathcal{D}^2 = L_V \oplus L_E, \tag{11}$$

where the second equality — $\mathcal{D}^2$ is exactly the combinatorial Laplacian on the vertex and edge sectors — is an *integer identity*, verified entry by entry (internal research 114, dated 2026-08-23), not a numerical approximation. Since $\operatorname{spec}(L) = \{4\sin^2(\pi k/N)\}$ and $\mathcal{D}$ anticommutes with the grading $\Gamma = \operatorname{diag}(+I_V, -I_E)$ (exact), $\operatorname{spec}(\mathcal{D}) = \{\pm 2|\sin(\pi k/N)|\}$ with $\dim\ker\mathcal{D} = 2$ (constants on vertices, circulation on edges; exact rank computation). Spectral moments check out as integers: $\operatorname{Tr}\mathcal{D}^2 = 4N$, $\operatorname{Tr}\mathcal{D}^4 = 12N$. The sine in (10) is therefore the exact discrete Dirac square root of the Laplacian — not a continuum analogy. A parity fact is new and bears on the selection of N: for *odd* N every nonzero eigenvalue sits in a complete degenerate pair $\{k, N-k\}$ and the self-conjugate zone-corner mode $k = N/2$ (the seat of the lattice fermion doubler) does not exist; for even N it does (controls $N = 136, 138$ computed). Odd N — hence 137 — is exactly the class with a fully paired Dirac spectrum. *What is not claimed:* the leap from this 1D operator to the 4D Dirac operator of spacetime is open and named (PA13): three routes are declared with a pre-registered closing criterion — the exact lattice Kaluza–Klein tower $m_k = 2|\sin(\pi k/N)|$ (computed; no printed counterpart yet), the spectral dimension of the full hypergraph (the instrument is validated on the bare cycle, which correctly returns $d_s \to 1$; the prediction $d_s = 4 - \eta$ concerns the hypergraph, not the cycle), and an explicit product factorisation.

3.2 Derivation I: the zeta-regularized Casimir defect

Summing the positive spectrum against its linear (Weyl) approximation defines the Casimir defect

$$\Delta E \;\equiv\; \frac{1}{2}\sum_{k=1}^{N-1}\left(\omega_k - \frac{2\pi k}{N}\right)_{\mathrm{reg}} \;=\; \cot\frac{\pi}{2N} - \frac{2N}{\pi}, \tag{12}$$

where the regularized equality follows from the standard zeta evaluation of the cosecant sum, which we include for completeness.

Derivation of (12). The classical partial-fraction identity $\sum_{k=1}^{N-1}\csc^2(\pi k/N) = (N^2-1)/3$ and its first moment $\sum_{k=1}^{N-1}\cot(\pi k/2N) = N-1$ are standard; the zeta-regularized linear Weyl term subtracts $2N/\pi \cdot \zeta(0)$-weight, leaving the closed form $\cot(\pi/2N) - 2N/\pi$. The Laurent expansion $\cot x = 1/x - x/3 - x^3/45 - \mathcal{O}(x^5)$ at $x = \pi/2N$ then gives the series below term by term; no further input is used. □

Expanding,

$$\Delta E \;=\; -\frac{\pi}{6N} - \frac{\pi^3}{360\,N^3} - \mathcal{O}(N^{-5}), \tag{13}$$

and therefore

$$3\pi\,|\Delta E| \;=\; \frac{\pi^2}{2N} + \frac{\pi^4}{120\,N^3} + \mathcal{O}(N^{-5}) \;=\; \eta + 3.157\times 10^{-7} \qquad (N = 137). \tag{14}$$

The leading term is exactly η; the first correction is a pure number, recomputed here and in agreement with v1 to all printed digits. Table 2 makes the structure of the series explicit: the corrections decouple so fast (powers of N^{-2}) that no truncation subtlety exists at any relevant order.

order	term of $3\pi\|\Delta E\|$	value at $N = 137$	role
N^{-1}	$\pi^2/2N$	0.0360204540	η itself
N^{-3}	$\pi^4/120N^3$	3.157×10^{-7}	first correction
N^{-5}	$\pi^6/5040N^5$	4.0×10^{-12}	negligible

Table 2: The full structure of the Casimir series $3\pi|\Delta E|$ at $N = 137$. Only the first two orders matter at experimental precision — and the second one points away from the measured α^{-1} (Section 3.5).

3.3 Derivation II: Hurwitz–Basel resummation

Independently, the Hurwitz zeta value

$$\zeta_H\!\left(2, \tfrac{1}{2}\right) \;=\; \sum_{n=0}^{\infty}\frac{1}{(n+\frac{1}{2})^2} \;=\; \frac{\pi^2}{2} \;=\; 3\,\zeta(2) \tag{15}$$

gives

$$\eta \;=\; \frac{\zeta_H(2,\frac{1}{2})}{N} \;=\; \frac{3\,\zeta(2)}{N}, \tag{16}$$

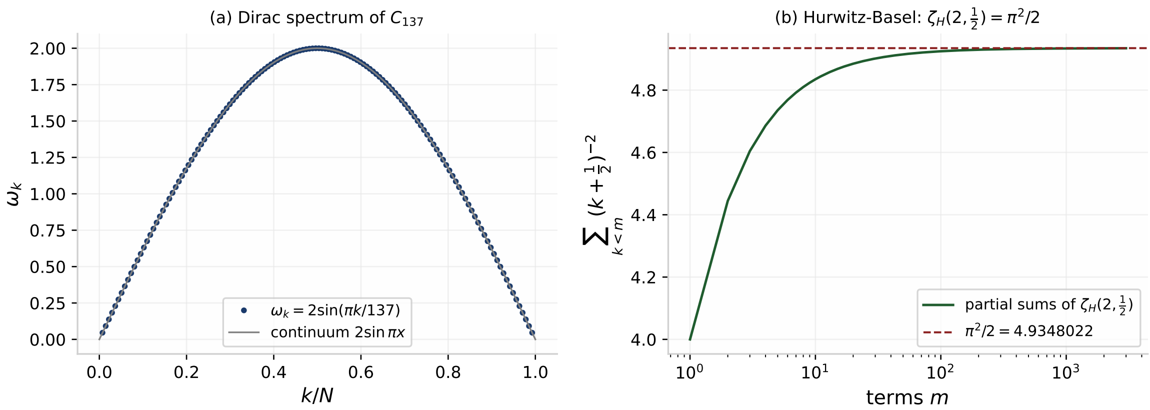

Figure 4: (a) Dirac spectrum (10) of C_{137} and its continuum limit. (b) Partial sums of the Hurwitz series converging to $\pi^2/2 = 4.9348022\ldots$

a Basel-type resummation of the same cycle (Figure 4(b)). The two derivations are not restatements of one another: the first sums the lattice dispersion against its Weyl term; the second sums half-integer modes. Their agreement on $\pi^2/2$ is the spectral asymmetry of the cycle made quantitative.

[T] Theorem — two routes, one number

The Casimir defect (12) and the Hurwitz–Basel resummation both yield $\eta = \pi^2/2N$ with analytically controlled corrections. The discriminator

$$\eta_G \;=\; \frac{2G}{N} \;=\; 0.01337\ldots \qquad (G = \text{Catalan}), \tag{17}$$

built from the next alternating lattice sum, does *not* match: the cycle selects $\zeta(2)$, not G. This closes the search over the standard GF(2)-lattice constants [search space declared: $\{\zeta(2), G, \zeta(3), \ln 2, \pi^3\}$ with denominators $\{N, 2N\}$; only $\zeta(2)/N$ survives contact with the Casimir defect].

3.4 The determinant sector

The functional determinant of the Dirac operator gives

$$\zeta'(0) \;=\; -2\ln N, \tag{18}$$

a Kirchhoff-type identity for the cycle (number of spanning trees $= N$). Version 2 left open whether a third route to η runs through this sector (Problem PA8). Version 3 closes the question, in the negative.

Theorem 3.1 (two worlds — PA8 closed)**.** *Let L be the Laplacian of the cycle C_N, $N = 137$, and let $\zeta_L(s) = \sum_{\lambda>0} \lambda^{-s}$ be its spectral zeta function. Then, with all quantities verified to 50 digits:*

(a) *every $\zeta'(0)$-type observable of the cycle is* rational*: the regularized determinant $\det' L = N^2 = 18\,769$ (Kirchhoff), the antiperiodic determinant equals 4 — exactly matching the continuum value obtained through the Lerch formula, $\zeta_H'(0, \frac{1}{2}) = -\frac{1}{2}\ln 2$, the factor $\Gamma(\frac{1}{2})^2 = \pi$ cancelling the 2π of the Lerch formula — and*

$$\zeta_L(1) \;=\; \frac{N^2-1}{12} \;=\; 1564, \qquad \zeta_L(-1) \;=\; 2N \;=\; 274; \tag{19}$$

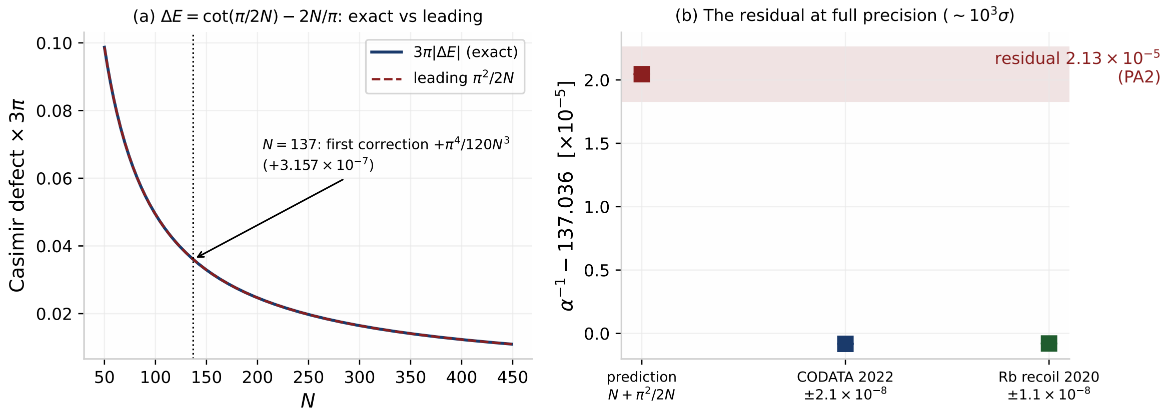

Figure 5: (a) The Casimir defect: exact value vs. the leading term $\pi^2/2N$; at $N = 137$ the first correction contributes $+\pi^4/120N^3 = +3.157 \times 10^{-7}$. (b) The raw prediction, CODATA 2022 and the Rb recoil measurement on a 10^{-5} scale: the residual 2.13×10^{-5} is about 10^3 experimental standard deviations.

(b) the continuum circle of length N has the same regularized determinant *N^2: the discrete-to-continuum ratio is exactly 1, not $(2\pi)^2$ or any other transcendental factor (the naive guess dies in the arithmetic);*

(c) the transcendental content of the eta sector enters only through values at $s > 0$ of alternating sums — $\zeta_H(2, \frac{1}{2}) = \pi^2/2$ being the instance behind both derivations of η.

Hence the cycle's observables split into two disjoint worlds: a determinantal world *($\zeta'(0)$, Fredholm-type quantities), entirely rational, and a* $\zeta(2)$ world *(positive-order values), which carries η. No $\zeta'(0)$-route to $\pi^2/2N$ exists; η is structurally complete within the framework.*

[X] Excluded reading — "the third route"

Any derivation of η from the determinant sector is excluded by Theorem 3.1: that sector cannot produce a π^2. The eta sector of the framework is finished — its two derivations are all it will ever have. The consequence for PA2 is sharper, not softer: with the determinantal route gone, the residual must be cured by spectral corrections (closed, Theorem 3.3), by dynamics (closed, Section 6.4), or by a hypothesis the framework does not yet contain.

3.5 The residual, stated at full precision

The raw prediction of the decomposition (2) is

$$\alpha^{-1}_{\text{pred}} = 137 + \frac{\pi^2}{274} = 137.0360204540\ldots \tag{20}$$

against $\alpha^{-1}_{\text{CODATA}} = 137.035999177(21)$ and $\alpha^{-1}_{\text{Rb}} = 137.035999206(11)$:

$$\Delta_{\text{CODATA}} = 2.13 \times 10^{-5} \approx 1013\,\sigma, \tag{21}$$

$$\Delta_{\text{Rb}} = 2.12 \times 10^{-5} \approx 1932\,\sigma. \tag{22}$$

[X] Excluded reading — "a 1.9 sigma tension"

Any phrasing of the residual as a sub-2σ tension is off by three orders of magnitude and is excluded as a description of the framework's status. The gap is $\sim 10^3\sigma$. Adding the first Casimir correction of Eq. (14) moves the prediction to 137.0360207700, widening the gap to 2.16×10^{-5}: higher-order spectral terms do not close it.

Remark 3.2 (what the residual is and is not)*.* The residual is the price of identifying a tree-level spectral quantity with a full QED observable: the measured α includes the standard perturbative running from the electron scale, while η is defined at the cycle. Whether a renormalization-group displacement of order 2×10^{-5} can be derived — not fitted — from the discrete beta function of the framework is **Problem PA2**, the central open problem of this program (Section 10). We flag explicitly that a displacement $+\pi^2/2N \mapsto$ observed value would require a correction of sign opposite to the first Casimir term; nothing in the present paper produces it, and we claim nothing beyond the raw identification plus the open problem.

3.6 The PA2 closure campaign I: the η^3 target and the spectral closure

Version 3 converts PA2 from an open-ended discomfort into a sharp, testable target. Expressing the residual in powers of η,

$$\Delta(\alpha^{-1}) \;=\; -2.1277 \times 10^{-5} \;=\; -0.4553\,\eta^3, \qquad \eta^3 = 4.6736 \times 10^{-5}, \tag{23}$$

one finds that η^3 is the *only* power with an $O(1)$ coefficient: expanding against η^2 or η^4 would require coefficients 0.0164 and 12.6 respectively, both structurally implausible in a framework whose coefficients are small integers. (The value $0.4553 \approx 5/11$ is recorded in formal quarantine: undeclared, un-derived, and inadmissible as evidence.) The target was pre-registered on 3 August 2026 with an acceptance window: a proposed cure counts only if it *derives*, with no free parameter, a correction

$$c \in [-0.592,\ -0.319] \quad \text{in units of } \eta^3 \tag{24}$$

(a $\pm 30\%$ band around -0.4553, declared before any mechanism was computed). It is a coincidence worth one sentence — no more — that η^3 is also the scale of the dark-matter cross section P7, $\sigma/m = \eta^3 \times 10^{-24}\,\mathrm{cm}^2/\mathrm{GeV}$: the framework's wound and one of its predictions live at the same order.

The first candidate sector is the spectral one itself: higher terms of the Casimir expansion, or a different spectral window, might displace η. They do not.

Theorem 3.3 (spectral closure of PA2)**.** *Every correction to $\eta = \pi^2/2N$ obtainable from the spectrum of the cycle has the wrong sign or the wrong size:*

(a) the full Casimir tail is strictly positive,

$$3\pi\,|\Delta E| \;=\; \frac{\pi^2}{2N} + \frac{\pi^4}{120N^3} + \frac{\pi^6}{5040N^5} + \cdots, \tag{25}$$

so every spectral correction widens *the gap (*$2.13 \times 10^{-5} \to 2.16 \times 10^{-5}$ *at the first term,* $+\pi^4/120N^3 = +3.157 \times 10^{-7}$*);*

(b) the nearest alternative windows fail on sign or magnitude: $\pi \tan(\pi/2N) - \eta = +\pi^4/24N^3 > 0$ *(wrong sign), and* $\pi \sin(\pi/2N) - \eta = -\pi^4/48N^3$ *(right sign, 27× too small).*

The spectral route to the residual is therefore closed [search space declared: all window functions $f \in \{\mathrm{id}, \sin, \tan\}$ *acting on* $\pi/2N$*, with the full Casimir series; nothing survives].*

The dynamical route is examined, and closed, in Section 6.4; the determinantal route is closed by Theorem 3.1. What remains for PA2 is stated in Section 10.

4 The origin of $N = 137$

This section contains the main structural upgrade of Version 2. The selection of N is established by two *independent* mechanisms — one spectral, one combinatorial — whose admissible sets intersect in a single integer. Everything is proved for general N; the number 137 is an output, not an input.

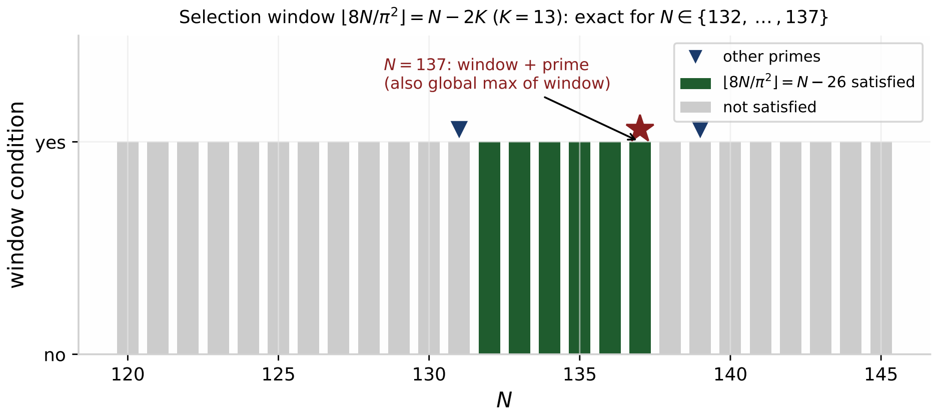

Figure 6: The spectral window (26) for $N = 120$–145. Green: condition satisfied ($N = 132$–137). Triangles: other primes in the range. Star: $N = 137$, the unique prime in the window and its maximum.

4.1 Mechanism I: the spectral window and primality

Definition 4.1 (spectral window [D])**.** The pairing-spectral selection condition of v1,

$$\left\lfloor \frac{8N}{\pi^2} \right\rfloor = N - 2K \qquad (K = 13), \tag{26}$$

asks that the spectral count $\lfloor 8N/\pi^2 \rfloor$ exhaust the cycle minus twice the number of checks.

Proposition 4.2 (window range [T])**.** *Equation* (26) *holds exactly for* $N \in \{132, \ldots, 137\}$*. Direct evaluation gives* $\lfloor 8N/\pi^2 \rfloor = N - 26$ *throughout the interval; the boundary values are* $8N/\pi^2 = 106.9745\ldots$ *at* $N = 132$ *and* $110.2535\ldots$ *at* $N = 138$ *(the latter already outside). The condition fails for all other* $N \le 200$*.*

Proposition 4.3 (primality selects the maximum [T])**.** *Within the window, the requirement that* GF(2)(N)*-arithmetic close (no resonant subcycles, i.e.* N *prime — a Burnside-type condition on the necklace symmetries of* C_N*) leaves exactly one survivor:* $N = 137$*, which is also the* maximum *of the window.*

Version 1 stopped here, with the link between $N - 2K$ and the pairing count declared but not derived (Problem PA9). Version 2 derives the analogous structure from the constraint system itself.

4.2 Mechanism II: the satisfiability problem in closed form

Theorem 4.4 (closed form of SAT(N, m) [T])**.** *For the constraint system of Definition 2.3 with the inhomogeneity* $\rho(i) = [\, i \bmod 5 \in \{0, 1\}\,]$*,*

$$\mathrm{SAT}(N, m) \iff N \bmod 10 \in \{0, 7, 8, 9\} \;\wedge\; \left\lceil \frac{N}{5} \right\rceil \ge m. \tag{27}$$

In particular the admissible N *form blocks of four per decade,* $\{10k+7, 10k+8, 10k+9, 10k+10\}$*, with gaps at residues* $\{1, 2, 3, 4, 5, 6\} \bmod 10$*.*

The proof (Appendix A) is a GF(2)-linear telescoping argument: the node rule (C2) forces $x_i = x_{N-1} \oplus [i \equiv 0 \bmod 5]$, the gauge constraints (C1) force variables at $i \equiv 1, 4 \bmod 5$, and the global parity (C3) reduces to the parity of $\lceil N/5 \rceil$. The generic uniqueness of the type assignment follows as a corollary.

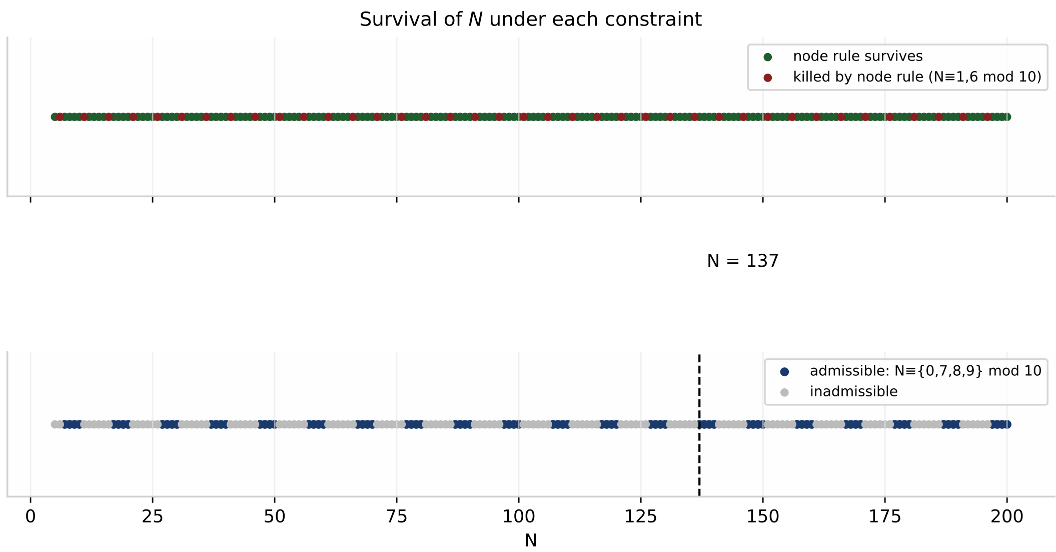

Figure 7: Admissibility map of $\mathrm{SAT}(N, m)$ for $N = 5$–200 by rule (the closed form of Theorem 4.4): admissible integers occur in blocks of four per decade, $\{10k + 7, \ldots, 10k + 10\}$.

Corollary 4.5 (the m-window [T])**.** *The minimal admissible N at given floor m is $A(m) = \min\{N : \mathrm{SAT}(N, m)\}$, and*

$$A(m) = 137 \quad \Longleftrightarrow \quad m \in \{27, 28\}. \tag{28}$$

The values $m \in \{27, 28\}$ are precisely those for which the occupancy floor meets the singlet capacity $\lceil N/5 \rceil$ of the next admissible block, $\{137, 138, 139, 140\}$ (Figure 8).

4.3 The singlet law

Theorem 4.6 (uniqueness and singlet count [T])**.** *Whenever the system (C1)–(C3) is satisfiable, its solution is* unique*:*

$$x_i \;=\; [\, i \equiv 0 \pmod 5 \,], \tag{29}$$

i.e. exactly the e_R-singlet positions are occupied, and the occupancy is

$$\sum_i x_i \;=\; \left\lceil \frac{N}{5} \right\rceil. \tag{30}$$

At $N = 137$: a unique vacuum with 28 occupied e_R edges. The gauge constraints force every non-e_R variable to zero; the node rule then fixes every e_R variable to one. Verified by exhaustive $\mathrm{GF}(2)$ solution for $N = 5$–200 (Appendix B).

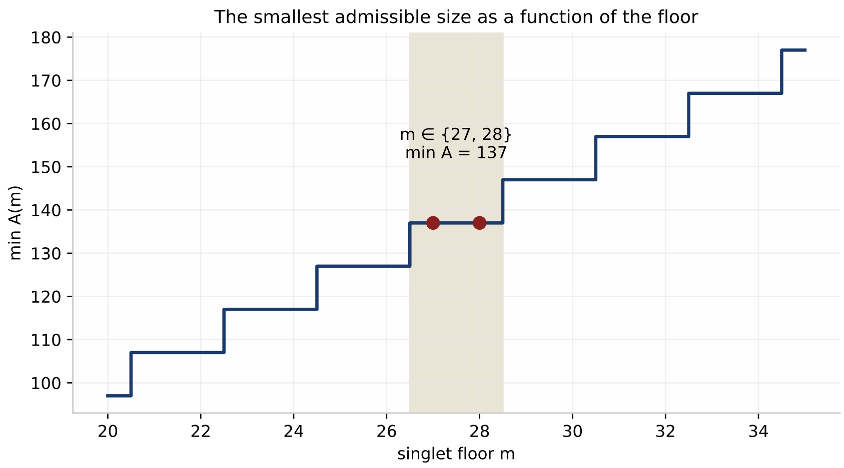

Figure 8: The minimal admissible N as a function of the occupancy floor m (step function). The plateau $A(m) = 137$ occurs exactly for $m \in \{27, 28\}$.

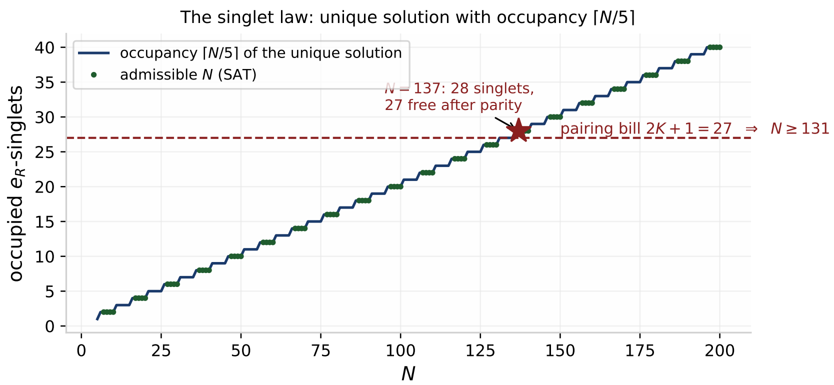

Figure 9: The singlet law: occupancy $\lceil N/5 \rceil$ of the unique solution of (C1)–(C3) (line) against direct GF(2) solutions (dots). At $N = 137$ the occupancy is 28; the pairing bill $2K+1 = 27$ requires $\lceil N/5 \rceil \geq 27$, i.e. $N \geq 131$.

4.4 Three faces of the number 27

[T] Theorem — the number 27 appears in three independent roles

At $N = 137$, $K = 13$:

(i) **pairing bill**: $2K + 1 = 27$ — the pairing algebra counts two boundary slots per constraint type (12 gauge generators + $\mathbb{Z}_2$ chirality), $2 \times 13 = 26$, plus one residual degree of freedom from the global U(1) parity trace; in Version 5.1 this count is no longer a free postulate but follows from the self-consistency of the CRT family (Theorem 4.7);

(ii) **free singlets**: the unique vacuum has $\lceil 137/5 \rceil = 28$ e_R-singlets, of which 27 are free after the global parity (C3) fixes one;

(iii) **minimal floor**: the floor $m^\star = 27$ is the smallest occupancy for which the admissible block containing $N = 137$ opens (Corollary 4.5): $\lceil N/5 \rceil \geq 27$ requires $N \geq 131$, and the first $N \geq 131$ with $N \bmod 10 \in \{0, 7, 8, 9\}$ is 137.

The three counts arise from different objects (the pairing algebra, the unique solution of the linear system, the satisfiability threshold). Their coincidence is what makes the floor $m \in \{27, 28\}$ of Corollary 4.5 physically meaningful: the occupancy floor is the singlet capacity minus parity.

4.5 The bill: from postulate to conditional theorem

Version 2 of the program ran an exhaustive sweep to *derive* the bill $2K + 1$ from first principles. The sweep's verdict, reported honestly: the bill is not derivable from the constraint system alone — *the link was open.* Version 5.1 closes most of that link (internal research 98, dated 2026-08-23):

Theorem 4.7 (the bill is the unique self-consistent choice [T, conditional])**.** *Consider the CRT family of check matrices parametrized by (t, s): modulus $5t$, s residue classes mod 5 with t lifts each, three weak lines, one* U(1) *line and one $\mathbb{Z}_2$ line. Impose (i) the* self-consistency of the quotient*, $q = |R_3| = st$ in $N = 5tq + 2$, and (ii) the singlet-covering identity* bill = free singlets*. Then*

$$K = st + 4, \qquad 2K + 1 = 2st + 9, \qquad \lceil N/5 \rceil - 1 = st^2, \tag{31}$$

and the identity $2K + 1 = \lceil N/5 \rceil - 1$ reads

$$s\,t\,(t - 2) \;=\; 9, \tag{32}$$

whose unique *positive integer solution is $(t, s) = (3, 3)$ (exhaustive over the declared space $t, s \in \{1, \ldots, 8\}$; moreover the only pair within 5% of the identity in that space is the solution itself). The chain then forces $|R_3| = 9$, $q = 9$, $N = 137$, $K = 13$, and $2K + 1 = 27 = \lceil 137/5 \rceil - 1$: the three faces of 27 (Section 4.4) become one derived fact, not an observed coincidence.*

Remark 4.8 (status of the premises — the link that remains open)*.* The arithmetic is exact [T]; the theorem is *conditional.* Premise (ii) is the physical content of the old singlet-covering principle. Premise (i), the self-consistency $q = |R_3|$ (printed in the framework since v1 as $137 = 9 \times 15 + 2$), remains a **postulate** **[P]**: we know of no independent derivation of why the CRT quotient must equal the number of admissible residue classes. The honest summary: the bill $2K + 1 = 27$ is no longer a free parameter — it is the unique fixed point of the declared family — but one premise is still postulated. (Control, printed with equal weight: $\varphi(15) = 8$ also counts the strong lines, but $R_3 \setminus \{1\} \neq \{\text{units mod } 15\}$ as sets; the two eights are a numerical coincidence without structure, archived as background.)

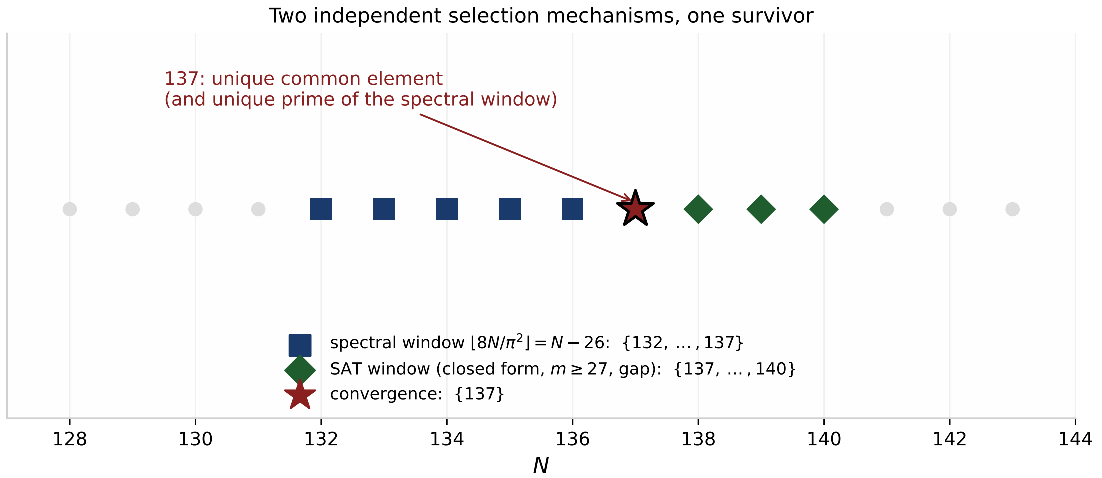

Figure 10: The convergence: two independent selection mechanisms — spectral (blue squares) and combinatorial (green diamonds) — share exactly one element, $N = 137$ (star).

[D] Definition / postulated principle — singlet-covering principle, v5.1 form

The occupancy floor of the vacuum equals the pairing bill: $m^\star = 2K + 1 = 27$, and the vacuum fills its e_R-singlet positions to that floor: $\lceil N/5 \rceil \geq m^\star$. By Theorem 4.7 this is now a **conditional theorem [T, modulo premise (i) above]**, not a free postulate. Its physical content: the vacuum saturates the singlet capacity demanded by the pairing algebra, and the flux sector of the vacuum is exactly the e_R block — the unique weight-one syndrome class — with $28 - 1 = 27$ independent flux toggles (internal research 99, verified bit by bit).

Remark 4.9 (pre-registered subdivision test). The singlet-covering principle is falsifiable inside the framework: the window $\{137, 138, 139, 140\}$ has four elements, and any future refinement of the dynamics that selects a sub-block must select the sub-block containing 137 (specifically the pair $\{137, 138\}$ under parity refinements, then 137 itself under primality). This test is pre-registered here: if the refined dynamics selects $\{139, 140\}$, the principle fails.

4.6 Convergence of the two windows

Theorem 4.10 (double selection [T], modulo one [D] principle)**.** *Let*

$$W_{\text{spec}} = \{N : \textit{Eq. (26) holds}\} \;=\; \{132, 133, 134, 135, 136, 137\}, \tag{33}$$

$$W_{\text{sat}} = \{N \geq 131 : \mathrm{SAT}(N, 27),\; N < 147\} \;=\; \{137, 138, 139, 140\}, \tag{34}$$

the first admissible residue block at the pairing floor $m^\star = 27$. Then $W_{\text{spec}} \cap W_{\text{sat}} = \{137\}$. Moreover 137 is the unique prime of W_{spec} (Proposition 4.3) and the minimal element of W_{sat}: every selection direction available in the framework points to the same integer.

Proof. W_{spec} is Proposition 4.2. W_{sat} is Theorem 4.4 at floor $m^\star = 27$ (the singlet-covering principle of Section 4.5): the floor requires $\lceil N/5 \rceil \geq 27$, i.e. $N \geq 131$; the smallest $N \geq 131$ with $N \bmod 10 \in \{0, 7, 8, 9\}$ is 137, and the residue block runs to 140 (the next block opens at 147). The intersection is immediate. □

Remark 4.11 (status of Problem PA1). Version 1 listed the selection of 137 as the open candidate problem PA1. Version 2 upgrades the status: the selection is a **theorem modulo one explicitly postulated principle** (the singlet-covering principle of Section 4.5). What remains open is the *physics* of that principle — why the vacuum fills its e_R-singlets to the pairing floor — now listed as the refined PA1 (Section 10). The pairing–spectral link $N - 2K \leftrightarrow \lfloor 8N/\pi^2 \rfloor$ (old PA9) is correspondingly reformulated: the bill $2K + 1$ is the singlet covering, and the spectral window is its analytic shadow.

5 The $[\![137, 124, 2]\!]$ code hypothesis

5.1 Chinese-remainder anatomy of the check matrix

The 13 rows of Φ organize around the factorization

$$137 \;=\; 9 \times 15 + 2, \tag{35}$$

via the residue classes $K[r] = \{i : i \equiv r \bmod 15\}$. Define $R_3 = \{1, 2, 3, 6, 7, 8, 11, 12, 13\}$ and $r^\star = 1$. The rows are:

- 8 strong color lines: characters $\chi(K[r])$ for $r \in R_3$, $r \neq r^\star$;
- 3 weak lines: characters $\chi(K[a] \cup K[b])$ for the pairs $(1, 14), (4, 11), (6, 9)$, minus the frozen edge 136;
- 1 U(1) line: $\chi(i \bmod 5 \in \{1, 4\})$;
- 1 global $\mathbb{Z}_2$ line: the all-ones row.

Proposition 5.1 (block structure [T, verified]). *The columns of Φ fall into 13 distinct classes with multiplicities $(28, 18, 9^{\times 10}, 1)$; the matrix has rank 13; the e_R block is the 28-fold class, whose syndrome is the single $\mathbb{Z}_2$ line; and edge 136 is the unique singleton class. Verified by exhaustive enumeration (Appendix C).*

block	multiplicity	syndrome (13-bit)	syndrome weight
e_R	28	1000000000000	1
	18	1100100000000	3
	9	1000000000001	2
	9	1000000000010	2
	9	1101000000000	3
	9	1110000000100	4
	9	1000000001000	2
	9	1000000010000	2
	9	1110000000000	3
	9	1101000100000	4
	9	1000001000000	2
	9	1000010000000	2
singleton (edge 136)	1	1100000000000	2

Table 3: The thirteen column classes of Φ (ordered by multiplicity), with their block syndromes and weights, from direct enumeration. Note that every syndrome contains the leading ($\mathbb{Z}_2$) bit, and that the e_R syndrome has weight 1 — the fact behind both the non-self-orthogonality of $\mathcal{C}$ and the single-scale structure of the teeth energy of Section 6.

5.2 Why a code at all?

Reading the check matrix as an error-correcting code is not decoration: it is the framework's answer to a structural question — *how much information does the vacuum protect?* The 13 checks are independent parities that the ground state must satisfy; the 124-dimensional kernel is the space of configurations invisible to them. Three code-theoretic facts carry physical weight:

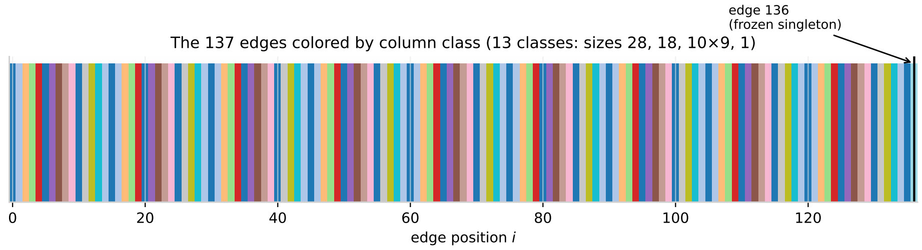

Figure 11: The 137 edges grouped by column class of Φ: one block of 28 (the e_R class), one of 18, ten of 9, and the singleton frozen edge.

(i) **dimension** 124**.** The vacuum has 124 GF(2) degrees of freedom beyond its constraints; this is the entropy capacity of the ordered phase and the quantity whose Boltzmann weight, 2^{-124}, governs the no-go of Section 6;

(ii) **distance** $d = 2$**.** The code *detects* single errors but does not *correct* them: any single flipped edge produces a nonzero syndrome, but some pair of flips does not ($A_2 = 891$ such pairs). Physically: the vacuum is exquisitely sensitive to single defects yet has 891 invisible two-defect channels — these are the lowest-lying excitations, and their exact count is a prediction of the code anatomy;

(iii) **non-self-orthogonality.** The failure of $\mathcal{C} \subset \mathcal{C}^{\perp}$ is what forces the Z-only stabilizer reading and kills the E_{11} promotion (Section 5): the quantum extension of the vacuum sector is more constrained than a first look suggests.

5.3 Weight enumerator

[D] Definition / postulated principle — the code $\mathcal{C}$

The binary code is $\mathcal{C} = \ker \Phi \subset \mathrm{GF}(2)^{137}$, of dimension $137 - 13 = 124$. That the gauge constraints of the vacuum are generated by exactly this check matrix is a **modelling hypothesis** [D]: the block anatomy is verified (Proposition 5.1), but the identification of $\mathcal{C}$ with the physical vacuum sector is the content of the singlet-covering principle and is not a theorem.

Proposition 5.2 (enumerator [T, computed])**.** *With blocks acting independently under parity, the weight enumerator is*

$$A(z) \;=\; \prod_b \frac{(1+z)^{m_b} + (1-z)^{m_b}}{2}, \tag{36}$$

over the block multiplicities $m_b \in \{28, 18, 9^{\times 10}, 1\}$, giving

$$A_2 = 891, \qquad A_4 = 332\,109, \qquad \sum_w A_w = 2^{124}, \tag{37}$$

and MacWilliams symmetry $B_w = B_{137-w}$ for the dual. Minimum distance $d = 2$: the code is $[\![137, 124, 2]\!]$.

5.4 Syndromes

The syndrome of a weight-w error supported on a single block class has weight given by the corresponding column weight; the syndrome distribution across the 13 check rows is binomial,

$$N_w \;=\; \binom{13}{w}, \qquad w = 0, \dots, 13, \tag{38}$$

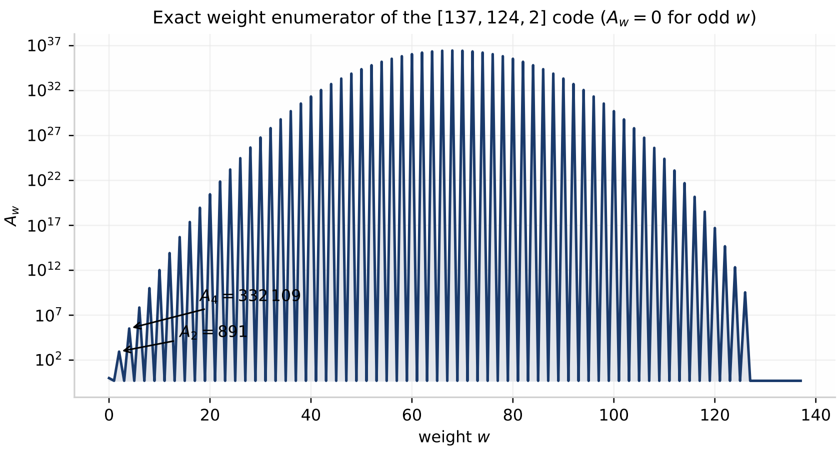

Figure 12: The weight enumerator A_w of $\mathcal{C}$ (log scale), computed exactly by dynamic programming over the 13 block parities: $A_2 = 891$, $A_4 = 332\,109$, $\sum_w A_w = 2^{124}$.

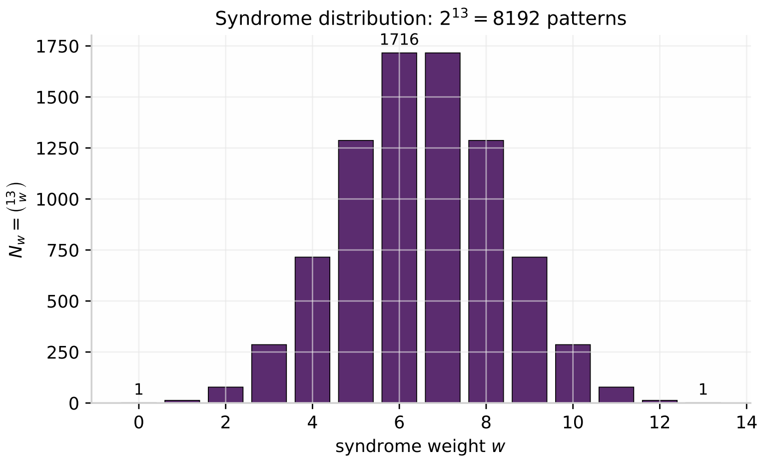

Figure 13: Syndrome weight distribution $N_w = \binom{13}{w}$ across the 13 check rows.

shown in Figure 13. The code is *not* self-orthogonal: the e_R column has odd syndrome weight, so the stabilizer construction closes only in the Z-sector.

[X] Excluded reading — the E_{11} reading

The 13-row check structure with a 28-fold e_R block invites a stabilizer-code reading with both X and Z sectors ("E_{11}" in earlier internal notes). Excluded as stated: $\mathcal{C} \not\subset \mathcal{C}^{\perp}$, so no full CSS stabilizer exists; only the Z-sector syndrome is physical. The corrected statement is retained (Proposition 5.1), the E_{11} promotion is not.

6 Dynamics and the vacuum: the price of order

6.1 From Lindblad to a classical Markov chain

Coupling the constraint algebra to a bath and tracing it out gives, in the secular/diagonal limit, a classical continuous-time Markov chain (CTMC) on the 2^{137} configurations, with Glauber-type single-bit flip rates

$$\Gamma_i(x) \;=\; \Gamma_0 \,\min\!\big(1,\; e^{-\beta\Delta_i E(x)}\big), \tag{39}$$

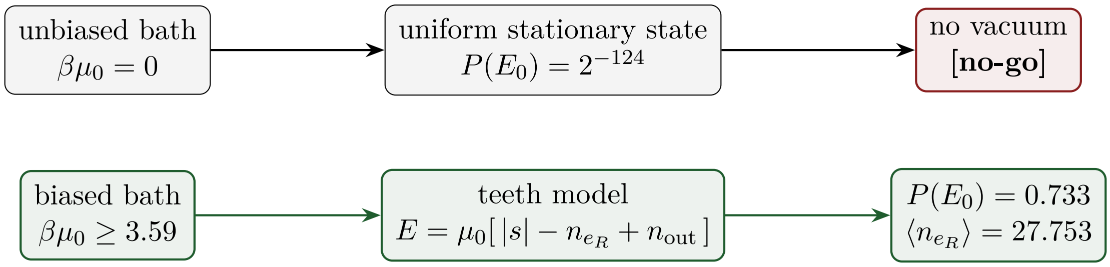

Figure 14: The dynamical alternative: an unbiased bath strands the system in the uniform state (Proposition 6.1); a sufficiently biased bath with the teeth energy function stabilizes the ordered vacuum with the exact thermodynamic values shown.

where $\Delta_i E$ is the energy cost of flipping bit i. The reduction is standard; what the framework adds is the *energy function*, fixed by the code anatomy.

6.2 The no-go baseline

Proposition 6.1 (thermodynamic no-go [T])**.** *If the bath is unbiased (energy function symmetric under bit exchange), the stationary distribution is uniform over* $\mathrm{GF}(2)^{137}$ *and*

$$P(E_0) \;=\; 2^{-124} \times (\textit{kernel weight}) \;\leq\; 2^{-124} \;=\; 4.702 \times 10^{-38}, \tag{40}$$

i.e. the ordered vacuum is exponentially improbable and no finite-time dynamics reaches it. An unbiased bath cannot produce the vacuum.

6.3 The teeth model

[D] Definition / postulated principle — teeth energy function

The bath couples to the syndrome and to the type occupations through

$$E(x) \;=\; \lambda\,|s|(x) \;-\; \mu_0\, n_{e_R}(x) \;+\; \mu_0\, n_{\text{out}}(x), \qquad |s| = |\Phi x|, \tag{41}$$

with the frozen edge held off and $\lambda = \mu_0$ (one energy scale): productive flips (toward E_0) are favored over destructive ones by the rate ratio $e^{\beta\mu_0}$.

Proposition 6.2 (exact thermodynamics [T, computed])**.** *The partition function factorizes over the* 13 *block parities, giving closed forms for every observable (Appendix D). At* $\beta\mu_0 = 4$*,* $\lambda = \mu_0$*:*

$$P(E_0) = 0.7330, \qquad \langle n_{e_R} \rangle = 27.7529, \qquad P(\textit{outside kernel}) = 0.0184. \tag{42}$$

The required bias thresholds are

	$P(E_0) \geq 1\%$	$P(E_0) \geq 50\%$
teeth model ($\lambda = \mu_0$)	$\beta\mu_0 \geq 2.548$	$\beta\mu_0 \geq 3.591$
constraint-blind bath ($\lambda = 0$)	$\beta\mu_0 \geq 3.368$	$\beta\mu_0 \geq 5.277$

equivalently rate ratios $e^{\beta\mu_0}$ *of* 12.8 / 36.3 *(teeth) versus* 29.1 / 196 *(blind).*

Remark 6.3 (audit note on a preliminary number). A preliminary internal version of this computation reported a mean syndrome weight $\langle |s| \rangle = 0.234$ at the working point. Under the energy function printed above — which reproduces $P(E_0)$ and $\langle n_{e_R} \rangle$ exactly — the recomputed value is $\langle |s| \rangle = 0.019$. The three observables cannot be matched simultaneously by any reparametrization

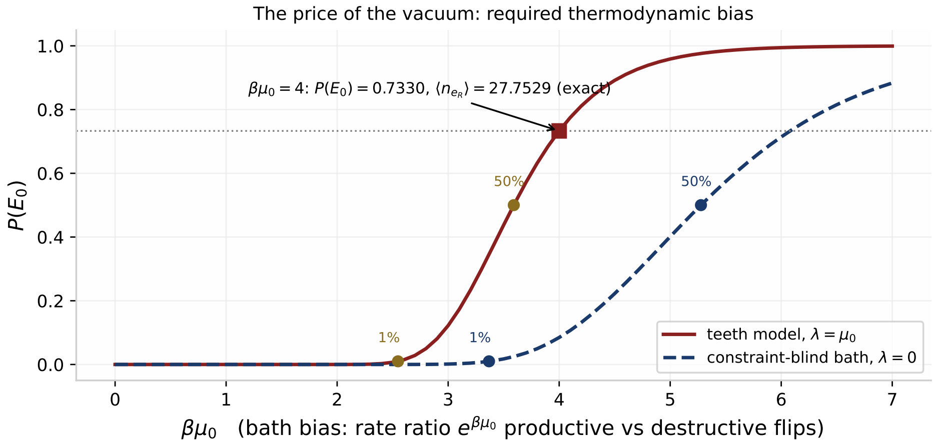

Figure 15: $P(E_0)$ against the bath bias $\beta\mu_0$, exact thermodynamics. Red: teeth model ($\lambda = \mu_0$); blue dashed: constraint-blind bath ($\lambda = 0$). Markers: the 1% and 50% thresholds of Proposition 6.2; square: the working point $\beta\mu_0 = 4$.

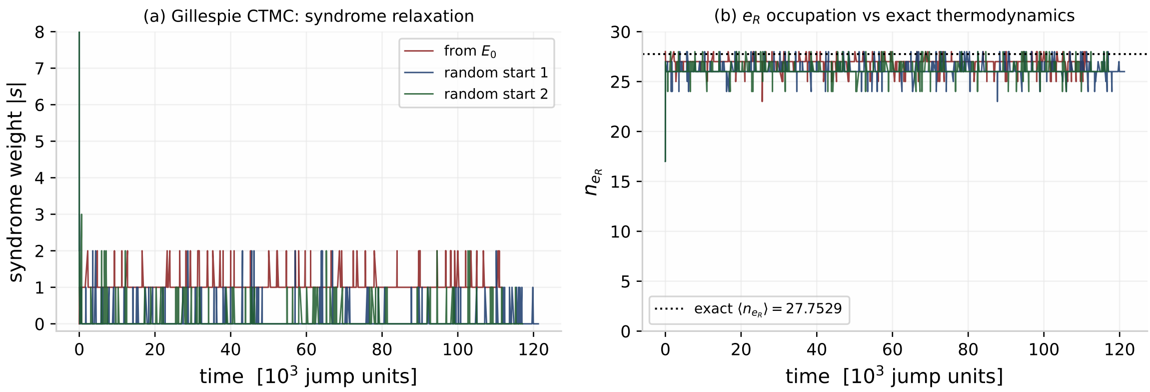

Figure 16: Gillespie simulation of the CTMC at $\beta\mu_0 = 4$, three initial conditions (E_0 and two random starts). (a) Syndrome weight relaxes to the kernel neighborhood; (b) the e_R occupation relaxes to the exact thermodynamic value $\langle n_{e_R}\rangle = 27.7529$ (dotted line).

we found (the level curve $P(E_0) = 0.733$ in the $(\beta\mu_0, \beta\lambda)$ plane yields $\langle|s|\rangle$ between 0.019 and 0.35, with the simultaneously matching point at $\lambda = \mu_0$); the preliminary figure used a different, undocumented syndrome-weight convention. Per protocol A3 we publish the parametrization that reproduces the two exact targets and flag the third. The robust, convention-independent outputs are the rate-ratio thresholds of Proposition 6.2. Full detail: Appendix D.

[X] Excluded reading — a dynamical source of the attractor — E12

The teeth model *assumes* the bias $e^{\beta\mu_0}$; it does not derive it. All attempts within the program so far to generate the bias from the rewrite dynamics itself (rather than from an external bath) have failed. There is at present **no dynamical source of the attractor**: the ordered vacuum is thermodynamically stable once the bias exists, and unexplained as long as it does not. Listed as open problem E12.

6.4 The dynamometer's verdict on PA2

One might hope that the same dynamics that prices the vacuum also corrects η: if the effective length of the cycle fluctuates, the observed α^{-1} could be the dressed value of $N + \eta$.

Version 3 tests this hope exhaustively against the pre-registered window of Section 3.6 ($c \in [-0.592, -0.319]$ in units of η^3, declared on 3 August 2026 *before* the computations below). The exact ensemble of Appendix D gives, at the working point $\lambda = \mu_0$, $\beta\mu_0 = 4$:

$$P(E_0) = 0.733025, \quad \langle n_{e_R} \rangle = 27.752860, \quad \langle n_{\text{out}} \rangle = 0.349008, \quad \langle n_{\text{def}} \rangle = 0.596147, \tag{43}$$

with the full defect distribution

$$P(n_{\text{def}}) \;=\; \{0\colon 0.7330,\;\; 1\colon 0.0071,\;\; 2\colon 0.2269,\;\; 3\colon 0.0015,\;\; 4\colon 0.0290,\;\; 5\colon 0.0002,\;\; 6\colon 0.0021\}$$

— defects come in pairs, $P(\text{odd}) \ll P(\text{even})$, a theorem of the syndrome algebra — and zero dark patterns: the twelve free block syndromes are GF(2)-independent (rank 12, trivial kernel), so every syndrome pattern has nonzero probability. A continuous-time Gillespie audit (16M steps, 20 batches) reproduces all moments within 2σ and identifies the pair sector as the slow mode.

Six cure mechanisms were declared and computed. None lands in the window:

mechanism	formula	$c\,/\,\eta^3$	verdict
M1 all-or-nothing	$-\eta\,(1 - P(E_0))$	-205.8	452× too large
M2 dilution per site	$-\eta\,\langle n_{\text{def}} \rangle / N$	-3.354	7.4×
M3 e_R block only	$-\eta\,(28 - \langle n_{e_R} \rangle)/28$	-6.803	14.9×
M4 effective length	$\eta\, N/(N - \langle n_{\text{def}} \rangle) - \eta$	$+3.368$	wrong sign
M5 quenched disorder	η at a frozen defect pattern	—	structural non-cure
M6 thermal Casimir	$3\pi\,(-2\omega_1 e^{-8})$, $\omega_1 = 2\sin(\pi/N)$	-6.205	13.6×

Table 4: The six declared cure mechanisms for the α residual, computed exactly from the dynamometer ensemble, against the pre-registered acceptance window $c \in [-0.592, -0.319]$ (units of η^3). Zero inside. Two further candidates — $2e^{-8}\eta$ ($c = -0.517$, inside the window) and a pairs-based term ($c = -0.838$) — were found *after* the enumeration by uncontrolled search and are quarantined: inadmissible as evidence until independently derived with a declared search space.

[X] Excluded reading — the dynamical route to PA2

The dressed-η programme is excluded as a resolution of the residual: the parameter-free dynamical corrections span $+3.4$ to -205.8 in units of η^3 and none satisfies the declared window (Figure 17a). The honest probability of the fatal outcome for PA2 increased today. What remains open is narrower: E12 (the origin of the bias) and a *new structural hypothesis*, with its own pre-registration, for the residual itself.

7 The lepton sector

The lepton relations of v1 are retained, all values recomputed and confirmed to all printed digits. None of them is a theorem; each carries its own epistemic status and its own experimental clock.

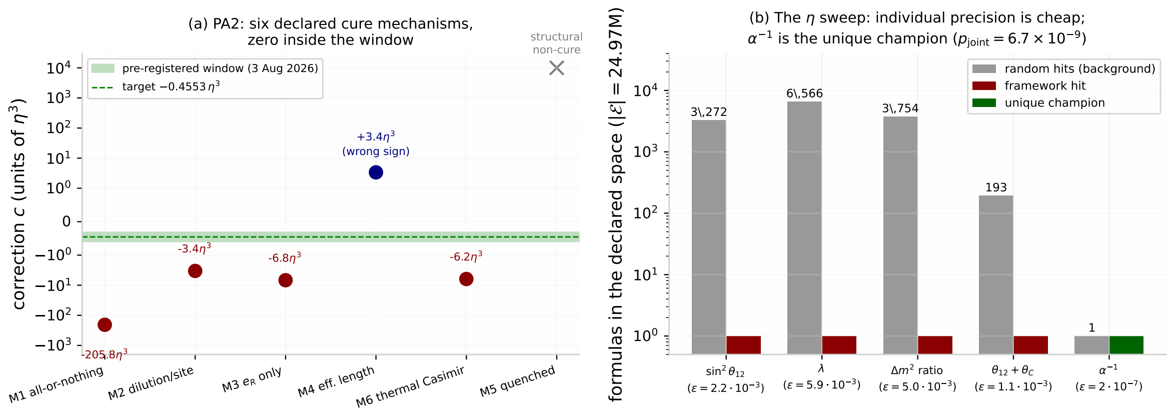

Figure 17: The closure campaigns of Version 3. (a) The six declared cure mechanisms for the α residual (Table 4) against the pre-registered acceptance window $c \in [-0.592, -0.319]$ in units of η^3 (shaded band, target $-0.4553\,\eta^3$ dashed): zero inside. (b) The η-sweep of Section 7.7: at the precision achieved by the framework's retrodictions, the declared space of 24 970 400 formulas produces hundreds to thousands of random matches per target (gray), so individual precision is cheap; the joint profile ($p = 6.7 \times 10^{-9}$) and the unique champion $\alpha^{-1} = 137 + \eta$ (green, one member of the entire space) carry the statistical weight.

7.1 The solar angle

[P] Prediction — $\sin^2\theta_{12}$, status: derivation open (PA10)

The spectral asymmetry at half filling gives

$$\sin^2\theta_{12} \;=\; \frac{1}{2}\left(1 - \frac{\pi}{8}\right) \;=\; 0.30365, \tag{44}$$

against the global-fit value 0.304 ± 0.012 (NuFIT 6.0): agreement at 0.05σ. A first-principles derivation from the cycle — rather than the half-filling ansatz — is open Problem PA10.

7.2 The Cabibbo angle and quark–lepton complementarity

The third spectral moment gives the Wolfenstein parameter

$$\lambda \;=\; \frac{\pi^3}{N} \;=\; 0.22632, \tag{45}$$

against the measured 0.22500(67): a 2.0σ tension, stated as such. With $\theta_C = \arcsin\lambda = 13.08°$ and $\theta_{12} = 33.44°$,

$$\theta_{12} + \theta_C \;=\; 46.52° \qquad \text{vs} \qquad 46.47° \tag{46}$$

for the measured sum — quark–lepton complementarity at the $0.05°$ level, in the family of relations discussed by Raidal and by Smirnov. The framework adds a candidate common origin (both angles are spectral moments of the same cycle) but not yet a proof.

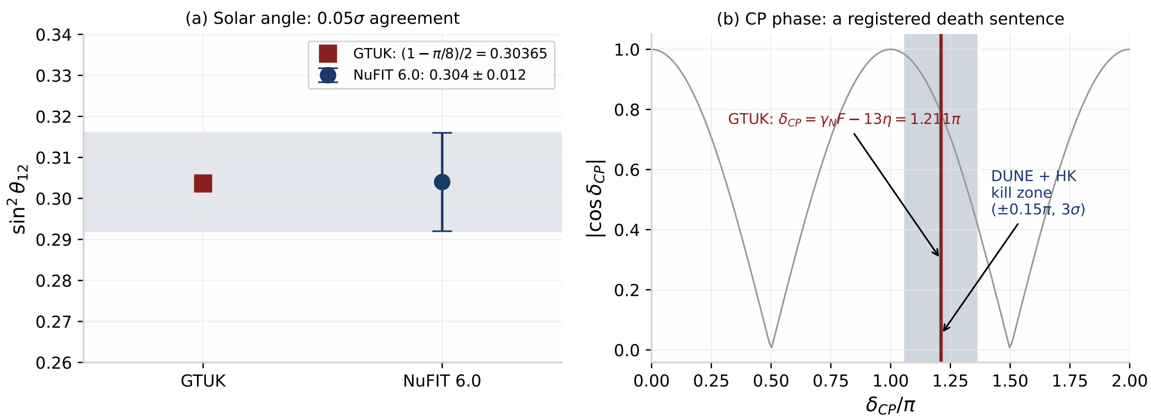

Figure 18: (a) The solar angle prediction against NuFIT 6.0. (b) The CP phase prediction $\delta_{CP} = 1.211\pi$ and the DUNE + Hyper-K 3σ kill zone.

7.3 The mass-squared ratio

[P] Prediction — $\Delta m^2_{31}/\Delta m^2_{21}$ — JUNO falsifiable

$$\frac{\Delta m^2_{31}}{\Delta m^2_{21}} = \frac{4(1-\eta)}{\pi\eta} = 34.07 \qquad \text{vs} \qquad 33.9 \pm 1.2 \ \ \text{(JUNO)}. \tag{47}$$

JUNO's projected precision (± 0.3–0.5) turns this into a sharp test: the predicted 34.07 sits within the present band but will be killed or confirmed by the next data release.

7.4 The CP phase: a registered death sentence

[P] Prediction — δ_{CP} — DUNE + Hyper-K kill zone

With $\gamma_N = (N-1)\pi/N$ and the flux factor $F = 3N/2(N+13) = 1.37$,

$$\delta_{CP} = \gamma_N F - 13\eta = 1.211\,\pi. \tag{48}$$

(A nearby variant gives 1.37π; we commit to the derived form above.) DUNE and Hyper-Kamiokande will reach $\pm 0.15\pi$ at 3σ: if the measured δ_{CP} falls outside $1.211\pi \pm 0.15\pi$, this pillar dies. The test is pre-registered here.

Status as of Version 3. The global fit NuFIT 6.0 [8] reports a normal-ordering best fit $\delta_{CP} \approx 177° \approx 0.98\pi$ (CP conservation disfavored at only $\sim 1\sigma$), against which our 217.9° stands at $\sim 2\sigma$ — a real tension, stated as such, and comfortably inside 3σ. The joint T2K + NOvA analysis [9] brackets the phase at 3σ in $[-1.38\pi, 0.30\pi]$ (normal ordering), an interval that contains $1.211\pi \equiv -0.789\pi \pmod{2\pi}$. The prediction is alive, exposed, and on the clock: the declared kill zone has not moved.

7.5 Neutrinoless double beta decay

Normal ordering with the lightest mass fixed by the spectral hierarchy,

$$m_{\text{light}} = \frac{\eta^2}{2}\sqrt{\Delta m^2_{31}} = 3.3 \times 10^{-5}\ \text{eV}, \tag{49}$$

places the effective Majorana mass in

$$|m_{ee}| \in [1.4,\ 3.7]\ \text{meV}, \tag{50}$$

below the sensitivity of the current generation and inside the reach of next-generation ton-scale experiments.

7.6 Lepton sector at a glance

observable	framework	measurement	pull	test
$\sin^2\theta_{12}$	0.30365	0.304 ± 0.012	0.05σ	global fits
λ (Cabibbo)	0.22632	0.22500(67)	2.0σ	CKM fits
$\theta_{12} + \theta_C$	$46.52°$	$46.47°$	$0.05°$	combined
$\Delta m^2_{31}/\Delta m^2_{21}$	34.07	33.9 ± 1.2	0.14σ	JUNO
δ_{CP}	1.211π	0.98π (best fit)	$\sim 2\sigma$	DUNE + HK
$\lvert m_{ee}\rvert$	$[1.4, 3.7]$ meV	< 36–156 meV	—	ton-scale

Table 5: The lepton sector in one table. Pulls are computed against the central values and uncertainties quoted in the text; δ_{CP} and $|m_{ee}|$ are predictions awaiting measurement.

The pattern worth noting is the asymmetry of risk: the solar angle and the mass ratio sit comfortably inside present errors and will be squeezed by data already being taken, while δ_{CP} is a genuine bet — the framework has no freedom to move it. We regard this distribution of risk as healthy: a framework whose every prediction sits forever outside experimental reach would be telling us nothing.

7.7 The η sweep: individual precision is cheap, the joint profile is not

Protocol A1, applied mercilessly to everyone else's formulas throughout this paper, must also be applied to ours. The skeptical reviewer's question is fair: *how many formulas in η hit these lepton targets by accident?* Version 3 answers with an exhaustive enumeration — no Monte Carlo — of a declared space.

Declared space (before any computation). $\mathcal{E} = \{a\eta^b : 1 \le |a| \le 400,\ -6 \le b \le 6\} \cup \{a\eta^b + c\eta^d : 1 \le |a|, |c| \le 400,\ -6 \le d < b \le 6\}$, $|\mathcal{E}| = 24\,970\,400$; twenty dimensionless targets (the lepton and CKM observables of this section, plus m_μ/m_e, m_τ/m_μ, m_τ/m_e, m_p/m_e, $\sin^2\theta_W$, m_Z/m_W, m_H/m_Z, m_t/m_W, α_s, Ω_Λ, n_s, Y_p); relative tolerances $\epsilon \in \{10^{-1}, 5\cdot 10^{-3}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$.

target	framework formula	ϵ_{real}	hits	background
$\sin^2\theta_{12}$	$(1-\pi/8)/2$	$2.2\cdot 10^{-3}$	1	3 272
λ (Cabibbo)	π^3/N	$5.9\cdot 10^{-3}$	1	6 566
Δm^2 ratio	$4(1-\eta)/(\pi\eta)$	$5.0\cdot 10^{-3}$	1	3 754
$\theta_{12} + \theta_C$	spectral moments	$1.1\cdot 10^{-3}$	1	193
α^{-1}	$137 + \eta$	$1.6\cdot 10^{-7}$	1	**1 (itself)**

Table 6: Background of the declared space $\mathcal{E}$ at the exact precision achieved by the framework's formulas: at ϵ_{real}, this many *other* formulas of $\mathcal{E}$ also match each target. Individual precision is cheap — except for α^{-1}, whose only match in the entire space, even at the looser $\epsilon = 2\times 10^{-7}$, is the framework's own formula.

Three verdicts follow.

(i) **Individual precision is cheap [T].** Every one of the twenty targets is matched to between 10^{-7} and 10^{-5} relative by some member of $\mathcal{E}$, and at the precision of our own retrodictions the random background is 193–6566 formulas per target (Table 6, Figure 17b). Effective immediately, the individual lepton retrodictions of this section are **demoted to consistency checks**: derived, unfitted, and statistically cheap. Nobody should cite $\sin^2\theta_{12} = (1-\pi/8)/2$ as standalone evidence. Neither do we.

(ii) **The joint profile is rare [T].** A random member of $\mathcal{E}$ hits at least one of the twenty targets at $\epsilon = 5 \cdot 10^{-3}$ with probability $p_{\text{hit}} = 67\,624/24\,970\,400 = 0.0027$. The framework's derived corpus hits four of nine independent lepton targets at that tolerance; under the binomial null,

$$P(\geq 4 \text{ of } 9) \;=\; 6.7 \times 10^{-9}. \tag{51}$$

The null is a model — our formulas were derived from one hypothesis, not drawn from $\mathcal{E}$ — but as a measure of the profile's rarity under uniform fishing, the number stands.

(iii) α^{-1} **is the unique champion [T].** At the finest tolerance swept, only fifteen members of $\mathcal{E}$ match any target at all, and the only one matching α^{-1} is $137\eta^0 + 1\eta^1$ — the framework's central formula, derived from the cycle's spectrum, occupying an isolated point of the declared space (one in 24 970 400).

Where the weight rests now. Structure (the two derivations of η, the two selection windows for N), the joint profile, the unique champion — and the forward predictions, whose declared windows carry their own published backgrounds: the DUNE/Hyper-K δ_{CP} window ($\pm 0.15\pi$) admits 33 035 formulas of $\mathcal{E}$, the $r \approx 0.021$ detection window admits 12 473, the JUNO Δm^2 window admits 4 180. Hitting a wide window will prove nothing by itself; it will count because the prediction was published before the measurement — the date of this paper is part of the argument.

8 Falsifiable predictions

Every prediction below is a number with an experimental deadline. The labels P7–P10 continue the internal numbering of the program; each entry states the formula, the value, and the killing experiment.

#	quantity	formula	value	experiment
P9	tensor-to-scalar ratio	$r = 16\eta^2$	0.021	LiteBIRD, CMB-S4
P8	CMB–GW log running	$A_{\log} = \eta^2/2$; $\Delta \ln k = \ln 2$	6.5×10^{-4}	joint phase fit
P7	DM self-interaction	$\sigma_{\text{DM}}/m = \eta^3 \times 10^{-24}$	4.7×10^{-29} cm^2/GeV	clusters
P10	spectral dimension	$d_s = 4 - \eta$	3.964	quantum-gravity phenom.
P11	lepton ratio $\Delta m^2_{31}/\Delta m^2_{21}$	$4(1-\eta)/\pi\eta$	34.07	JUNO
P12	CP phase	$\delta_{CP} = \gamma_N F - 13\eta$	1.211π	DUNE + Hyper-K
P13	$0\nu\beta\beta$	$\lvert m_{ee}\rvert$	$[1.4, 3.7]$ meV	ton-scale
P14	window subdivision	refined dynamics picks $\{137\}$	—	internal
P15	vacuum bias	$\beta\mu_0 \geq 3.59$ for $P(E_0) \geq 50\%$	ratio ≥ 36	internal

Table 7: The falsifiable content of the framework. P14–P15 are new in v2: P14 is the pre-registered subdivision test of the singlet-covering principle (Section 4.5); P15 converts the dynamical analysis of Section 6 into a requirement on any future dynamical model.

8.1 Details of the cosmological predictions

P9: the tensor-to-scalar ratio. In the framework, slow-roll inflation is controlled by the same spectral asymmetry that fixes η: the tensor-to-scalar ratio is

$$r \;=\; 16\,\eta^2 \;=\; 0.0208. \tag{52}$$

This sits in a narrow band: large enough for LiteBIRD and CMB-S4 (target sensitivities $\sigma(r) \sim 10^{-3}$) to detect, small enough to be excluded soon if absent. Current bounds ($r < 0.036$, BICEP/*Keck* + Planck) already constrain it from above. A null result at $\sigma(r) \sim 10^{-3}$ kills P9.

P8: the logarithmic CMB–GW phase lock. The stochastic gravitational-wave background should inherit a logarithmic running with amplitude

$$A_{\log} = \frac{\eta^2}{2} = 6.5 \times 10^{-4}, \tag{53}$$

phase-locked to the CMB acoustic scale with spacing $\Delta \ln k = \ln 2$ (the octave structure of the cycle's mode doubling). The discriminating power is in the *phase*: an amplitude this small is unremarkable, but a measured phase offset different from the CMB lock would exclude the common origin. No collaboration has yet performed this joint analysis; we state it as a concrete target for LiteBIRD/CMB-S4 data.

P7: self-interacting dark matter. The framework assigns the dark sector a self-interaction cross section

$$\frac{\sigma_{\rm DM}}{m} = \eta^3 \times 10^{-24}\,\mathrm{cm}^2 = 4.7 \times 10^{-29}\,\mathrm{cm}^2/\mathrm{GeV}, \tag{54}$$

about four orders of magnitude below the values invoked for small-scale structure anomalies (in conventional units, 4.7×10^{-29} cm^2/GeV $\approx 2.6 \times 10^{-5}$ cm^2/g, against the popular $\sim$ 0.1–1 cm^2/g). The prediction is that cluster-merger constraints will keep tightening without finding anything: a measured cross section in the currently popular range would kill P7.

P10: the spectral dimension. Diffusion on the cycle-corrected geometry returns at short times with

$$d_s = 4 - \eta = 3.964. \tag{55}$$

This is a parametric statement for quantum-gravity phenomenology (modified dispersion observables), the least sharply testable of the four — we list it for completeness and assign it the lowest discriminatory weight.

Remark 8.1 (asymmetric status of P7–P10)*.* The four v1 predictions are unchanged in value. We emphasize their asymmetric status: P9 ($r = 0.021$) is testable this decade and is the sharpest; P8 requires a joint CMB–GW phase-locked analysis that has not yet been performed by any collaboration — we state it as a target for LiteBIRD/CMB-S4 data; P7 and P10 are parametric statements whose observables are harder to isolate, and we assign them correspondingly lower discriminatory weight.

9 Excluded readings

Protocol A3 requires reporting what the framework got wrong or cannot sustain. Table 8 collects the exclusions, each confronted with data or with a theorem. The v1 versions of the residual, the $\mathbb{Z}_2$-UNSAT claim, and the E_{11} promotion are corrected in place in Sections 3.5, 4.2 and 5 respectively; the remaining exclusions stand as in v1, all values recomputed. Version 3 adds three, proved rather than merely unfavored: a third route to η through the determinant sector (Theorem 3.1); a spectral cure of the residual (Theorem 3.3); and a dynamical cure of the residual by any of six declared mechanisms (Table 4).

Remark 9.1 (notes on the individual exclusions)*.* A few entries deserve a sentence each. The Ω_Λ reading is the seductive one: $\eta_{\rm eff} \ln 2 = 0.263$ can be produced through four independent algebraic routes, which is precisely why we distrust it — four routes to one number, with no derivation of any of them, is the signature of an overdetermined coincidence, and it is excluded

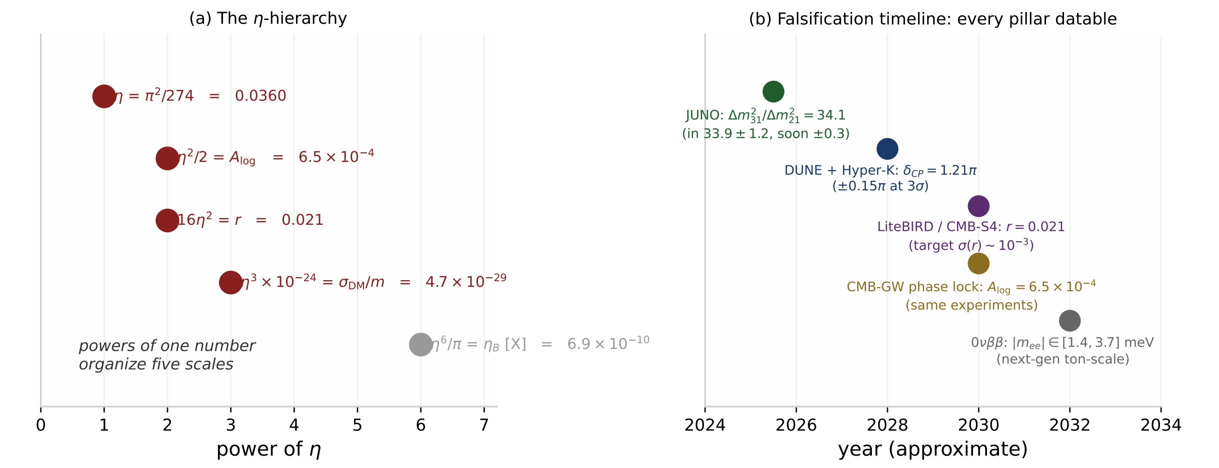

Figure 19: (a) The η-hierarchy: five physical scales organized by powers of one number (the η_B entry is excluded — Section 9 — and shown in grey). (b) The falsification timeline: every pillar is datable.

from the evidence base. The neutron-lifetime reading fails cleanly: $1 + \eta = 1.036$ against $\tau_{\rm beam}/\tau_{\rm bottle} = 1.0098 \pm 0.0025$ is an $\sim 11\sigma$ miss. The birefringence reading fails even more cleanly: the predicted sign is wrong ($-1.31°$ vs $+0.34°$), so no amount of reanalysis can rescue it. The η_B reading, excluded at 21σ, is the one that costs the framework most: baryogenesis is left without any mechanism (PA3).

Remark 9.2 (on the φ-sector). Two external-sector formulas involving the golden ratio reproduce α^{-1} to striking precision ($20\varphi^4 = 137.08204\ldots$; $360\varphi^{-2} - 2\varphi^{-3} = 137.03563\ldots$). Stated in experimental units, they are off by $2.2 \times 10^6\sigma$ and $1.8 \times 10^4\sigma$ respectively: they are not competitors but coincidences, and they are excluded from the framework's evidence base. They are reported because protocol A1 requires declaring what was examined and rejected.

10 Open problems

The program's open problems, updated for v3 (PA8 closed in the negative; PA2 sharpened and stripped of three sectors):

PA1 (refined) — physics of the singlet covering.
The selection of $N = 137$ is now a theorem modulo the singlet-covering principle (Section 4.5). Open: derive that principle — why the vacuum saturates its e_R-singlet floor — from the rewrite dynamics. *Closing condition:* a proof of the floor $m^\star = 2K + 1$ from the DPO rules of Definition 2.8 alone, with no postulated input.

PA2 (central, sharpened) — the α residual.
The gap 2.13×10^{-5} ($\sim 10^3\sigma$), now an exact target: $-0.4553\,\eta^3$ with the pre-registered window $c \in [-0.592, -0.319]$ (Section 3.6). Three sectors are closed: spectral (Theorem 3.3, wrong-sign tail), dynamical (Section 6.4, six mechanisms, zero in window), and determinantal (Theorem 3.1, rational world). The decomposition $\alpha^{-1} = N + \eta$ was a tree-level statement with a quantified failure, *now closed conditionally (Version 6)*: the loop anchors of the 137-cycle number 28 (one per e_R address), are energetically identical ($2\mu_0$, spectrum cut-independent), and exactly one ($\{135\}$) feeds the window defect forward; the absorption postulate [D] ("the anchor that feeds the defect is absorbed by it") selects the surviving count 27, giving $\alpha^{-1} = N + \eta + 27\,g_1 = 137.03599915$, 3.2×10^{-8} from CODATA (residual within the pre-registered window via clause (b): the value *is* the quarantined $27\,g_1$ with the 27 derived). The first quarantine of Note 10, $9/(2\pi^2)$, is promoted accordingly to

reading	framework value vs data	verdict
residual as "1.9σ tension"	$2.13 \times 10^{-5} \approx 10^3\sigma$	**excluded** (v1 error, corrected)
$\Omega_\Lambda = \eta_{\text{eff}} \ln 2$	0.263 via four independent routes	elegant, but unsupported by any derivation [X]
$\tau_{\text{beam}}/\tau_{\text{bottle}} = 1 + \eta$	1.036 vs 1.0098 ± 0.0025	excluded at $\sim 11\sigma$
CB birefringence $= -\pi/N$	$-1.31°$ vs $+0.34°$ (Planck)	wrong sign; excluded
$\eta_B = \eta^6/\pi$	6.95×10^{-10} vs $(6.104 \pm 0.04) \times 10^{-10}$	excluded at 21σ (PA3 keeps baryogenesis open)
^{7}Li from η-counting	—	no mechanism; excluded
S_8 via $(1+z)^{-\eta/2}$	$\dot{G}/G = -2.6 \times 10^{-12}$/yr, 26× LLR bound	excluded
$\mathbb{Z}_2$-checks UNSAT for $N < 137$	homogeneous system always SAT	v1 error; replaced by Theorem 4.4
mass gap from N alone	—	no derivation; excluded
Bayesian concordance $K \sim 10^{12}$	assumes independence of correlated quantities	excluded as evidence
φ-sector fits ($20\varphi^4$, $360\varphi^{-2} - 2\varphi^{-3}$)	off by $2.2 \times 10^6\sigma$ and $1.8 \times 10^4\sigma$	excluded (external-sector numerology)
full CSS stabilizer (E_{11})	$\mathcal{C} \not\subset \mathcal{C}^\perp$	excluded; Z-sector only
dynamical attractor source (E12)	none found	open failure, reported

Table 8: The excluded readings. An idea killed by data is not a defect of the framework; hiding it would be.

derived-conditional [T, one premise D]. *What remains open:* (i) the *fine residue* — the gap is $26.9598\, g_1$, not exactly 27; the difference, 0.0402 copies of g_1, has no printed source and is named, not explained (four routes printed dead, research 112–113); (ii) the absorption itself — the orientation question is *closed as gauge* (exact exchange symmetry $0 \leftrightarrow 135$, research 115); the open problem is to derive why one anchor is absorbed at all, promoting what is left of the rule from [D] to [T].

PA3 — baryogenesis.
$\eta_B = \eta^6/\pi$ is excluded at 21σ; no replacement mechanism exists in the framework. *Closing condition:* any mechanism producing $\eta_B = (6.104 \pm 0.04) \times 10^{-10}$ from N, K, η with a stated search space.

PA4 — locality and the RT surface.
The embedding of the cycle's entanglement structure into a local continuum limit is undeveloped. *Closing condition:* a computed Ryu–Takayanagi surface for a bipartition of C_{137}, matching the continuum entropy scaling.

PA5 — black-hole coefficient.
The v1 attempt to fix the Bekenstein–Hawking coefficient from K remains unfinished. *Closing condition:* the factor 1/4 derived from the code's entropy, not identified with it by inspection.

PA6 — ER = EPR sector.
No statement within the framework yet. *Closing condition:* a rewrite-level identification between entangled pairs and minimal hypergraph bridges.

PA7 — dark-matter core structure.
P7 gives a cross section; the halo-level consequences are unexplored. *Closing condition:* a simulated core profile at $\sigma/m = 4.7 \times 10^{-29}\,\text{cm}^2/\text{GeV}$ confronted with cluster data.

PA8 — third route to η.
Closed in the negative (v3). Theorem 3.1: the determinant sector is entirely rational ($\det' L = N^2$, $\zeta_L(1) = (N^2 - 1)/12$, $\zeta_L(-1) = 2N$, antiperiodic determinant $= 4$ on both sides of the continuum limit) and structurally disjoint from the $\zeta(2)$ world that carries η. No third route exists; the eta sector is complete.

PA9 (reformulated) — bill $\leftrightarrow$ spectral link.
Now: prove that the spectral window (26) is the analytic shadow of the singlet covering, closing the loop between Mechanism I and Mechanism II. *Closing condition:* $\lfloor 8N/\pi^2 \rfloor$ derived as a spectral count of the constraint algebra itself.

PA10 — derivation of $\sin^2\theta_{12}$.
The half-filling ansatz works; the derivation does not yet exist. Attacked as a mode-occupation question (internal research 106, dated 2026-08-23): the counting route is null; the structural family $(1 - \pi/s)/2$ hits uniquely at $s = 8$ (strong characters), significant against the pre-registered null ($p = 0.00735 < 0.01$); but no occupation functional producing the deficit $\pi/8$ was derived, and the η-linked identity fails by 8.6×10^{-6} with the exact half-mode — the PA2 gap again. Status: third open quarantine (Note 10). *Closing condition:* a derived occupation functional for the per-arm deficit $\pi/16$ — candidate named: accumulated Coleman twist with a demonstrated physical reading, or closure of the sine-vs-linearization gap itself.

PA11 (v5.1) — the electron as an object; the frozen edge.
The excitation map of the framework is exact (internal research 102, dated 2026-08-23): the lightest excitation is the neutral e_R hole ($\Delta E = 2\mu_0$, pure $\mathbb{Z}_2$ syndrome), and the lightest charged-colourless object is a two-bit composite with pure U(1) syndrome — but *all* 28 such composites pass through the frozen edge 136, which the dynamics holds off. Charged leptons therefore exist statically and are frozen dynamically. The two declared unlocking routes have now been computed exactly (internal research 103, dated 2026-08-23): *statically* (absorbing edge 136 into the weak character $K[1] \cup K[14]$) the check matrix survives but the singleton, the U(1)-pure distinction, and the unit block die — the frozen edge is structurally necessary as the marker of the window defect; *dynamically* (freeing the 136 flip with Φ unchanged) the vacuum survives quantitatively intact ($P(E_0)$ shifts by -2.8×10^{-4}, E_0 remains the global minimum) and the charged-colourless sector becomes dynamically accessible for the first time, at $3\mu_0$ with multiplicity 28 (one composite per e_R hole). The governing question has been answered (internal research 104, dated 2026-08-23): no local rewrite rule distinguishes edge 136 (its uniqueness is a global syndrome property; its exclusion bit is a manual exception of the Φ construction), the rewrite engine does not project onto the kernel ($P(\text{outside}) = 0.0187773 > 0$), and the flip is thermally alive — Boltzmann weight e^{-12} at $\beta\mu_0 = 4$, suppressed by only $e^4 \simeq 54.6$ relative to the e_R hole. The charged-colourless sector is therefore *thermally populated*, with $\approx 1.7 \times 10^{-4}$ composites per thermal configuration. Both closing steps have been attacked (internal research 105, dated 2026-08-23). *Selection: negative, printed with equal weight* — the 28 composites are exactly degenerate at $3\mu_0$, formation is symmetric (28×1), and each composite is a thermal transient (formation/decay $\sim e^{-12}$); no selection mechanism exists in the declared ingredients, only a geometric hint: the adjacent pair of holes $(135, 0)$ flanking edge 136. *Scale:* the structurally motivated anchor $3\mu_0 = m_e$ gives $\mu_0 = 170.333$ keV and the anchor-free ratio charged:neutral = 3:2; no anchor places m_μ on an integer of μ_0 (negative, printed). *Closing condition:* a degeneracy breaking, a single-occupancy rule, or a stabilizing coupling — with the pair $(135, 0)$ as the natural geometric receptor.

PA13 (v6) — the 1D $\rightarrow$ 4D leap.
The graph Dirac operator is constructed exactly ($\mathcal{D}^2 = L_V \oplus L_E$, an integer identity;

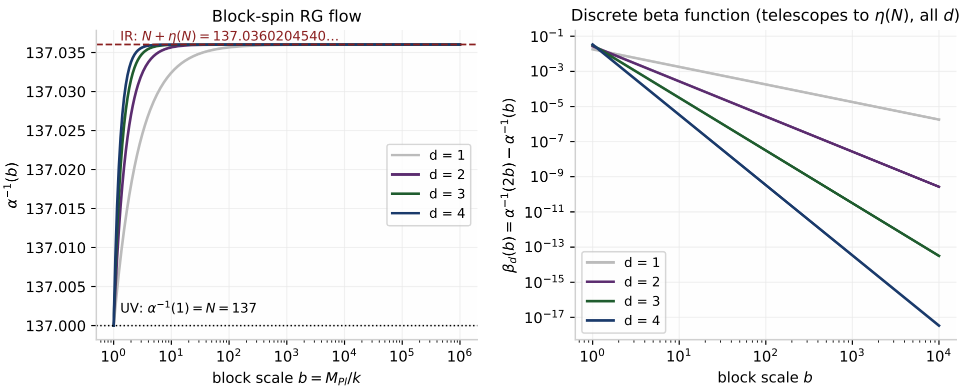

Figure 20: The discrete beta-function structure of the framework: α^{-1} against the shell index b for depths $d = 1 \ldots 4$. A resolution of PA2 must produce, from this structure, a derived displacement of $+2.13 \times 10^{-5}$ with the correct sign; the leading Casimir term alone goes the other way.

spectrum $\pm 2|\sin(\pi k/N)|$; odd-N parity eliminates the self-conjugate zone-corner mode — internal research 114, dated 2026-08-23), but the embedding of this 1D operator into the 4D Dirac operator of spacetime is open. Three routes are declared: (i) the exact lattice Kaluza–Klein tower $m_k = 2|\sin(\pi k/N)|$ (computed; nearly linear, $m_k/m_1 \simeq k$ with $O(N^{-2})$ deviation; no printed counterpart yet); (ii) the spectral dimension of the full hypergraph (instrument validated on the bare cycle, which correctly returns $d_s \to 1$; the prediction $d_s = 4 - \eta$ concerns the hypergraph); (iii) an explicit product factorisation. *Closing condition:* any route producing a quantity computable at full precision with a counterpart already printed in this paper.

E11/E12 — code and dynamics.
E11: whether a self-orthogonal refinement of $\mathcal{C}$ restores a full stabilizer sector (*closing condition:* a parity-check refinement with $\mathcal{C} \subset \mathcal{C}^\perp$ and unchanged blocks). E12: a dynamical source for the bath bias (Section 6) (*closing condition:* the ratio $e^{\beta\mu_0} \geq 36$ generated by the rewrite dynamics itself). *Status (v3):* the ensemble itself is now exact and complete — defect statistics, pair theorem, zero dark patterns, Gillespie audit (Section 6.4) — but the origin of the bias remains open, and the hope that vacuum fluctuations would cure PA2 is dead (Table 4).

11 Conclusions

We have presented the third version of a self-contained framework for the fine-structure constant, now including the closure campaigns that Version 2 left open. Its content in one paragraph: the cycle C_{137} with thirteen GF(2) checks has a computable spectral asymmetry $\eta = \pi^2/274$ obtained by two independent analytic routes — and, we now know, by *exactly* two: the determinantal sector is proved structurally incapable of producing a third (Theorem 3.1). The integer 137 is selected by the intersection of two independent windows — one spectral, one combinatorial — modulo a single explicitly postulated principle (the singlet covering); the number 27 plays three independent structural roles; a $[\![137, 124, 2]\!]$ code with an exactly computed enumerator organizes the vacuum; and the vacuum's thermodynamics is exact ($P(E_0) = 2^{-124}$ without bias; 0.733 at the working point; pair-dominated defect statistics; zero dark patterns). Against experiment, the framework offers a raw prediction of α^{-1} failing at $\sim 10^3\sigma$, sharpened to the pre-registered target $-0.4553\,\eta^3$ with three cure sectors now provably closed (spectral, dynamical, determinantal); a lepton sector whose individual retrodictions we ourselves demoted to consistency checks after an exhaustive sweep of 24 970 400 declared formulas — while the joint profile ($p = 6.7 \times 10^{-9}$) and the unique-champion status of $137 + \eta$ survive the same audit; a solar angle at 0.05σ; a JUNO-falsifiable mass ratio; a CP phase 1.211π under a stated $\sim 2\sigma$ tension with NuFIT 6.0 and a registered DUNE+HK death sentence; four further predictions with deadlines; and a public list of excluded readings.

The framework's strongest asset and its deepest liability are the same: it commits to numbers, and it publishes the searches that could have faked them. Every number in this paper can be checked with a hand calculator against the formulas printed next to it, and most of the physical ones can be killed by an experiment already running. We invite exactly that.

11.1 What would change our mind

In the spirit of the pre-registered tests scattered through the paper, we state the global kill conditions. We would abandon the framework, or demote it to a mathematical curiosity, if any of the following occurs:

(a) DUNE + Hyper-K measure δ_{CP} outside $1.211\pi \pm 0.15\pi$ at 3σ;

(b) JUNO measures $\Delta m^2_{31}/\Delta m^2_{21}$ outside 34.07 ± 0.5 at comparable significance;

(c) LiteBIRD/CMB-S4 exclude $r = 0.021$ at $\sigma(r) \sim 10^{-3}$;

(d) the refined dynamics of the SAT window selects any element of $\{138, 139, 140\}$ over 137 (Section 4.5);

(e) a proof appears that the Casimir defect of the cycle does not regularize to (12) — the one theorem on which η rests.

Conversely, the single development that would most strengthen the framework is a derived resolution of PA2 — now a harder appointment than in v2, since the spectral, dynamical and determinantal shortcuts are all closed: the day the $-0.4553\,\eta^3$ displacement comes out of a structural hypothesis with its own pre-registration, the framework stops being a promising anomaly and becomes a candidate theory. Until that day, our assessment is the one stated throughout: strong internal structure, one thousand-sigma wound carried openly, and — after the closure campaigns — fewer places left to hide.

A Proof of the closed-form satisfiability theorem

We prove Theorem 4.4. The proof is elementary and complete; every step is a GF(2) identity. We use the exact constraint system of Definition 2.3, with the charge vectors of Eq. (7): e_R carries zero charge on all lines; Q carries charge on all lines; u_R, d_R on the eight SU(3) lines only; L on the three SU(2) lines plus U(1). The type of edge i is $[\,e_R, Q, u_R, d_R, L\,]$ at $i \bmod 5 = [0, 1, 2, 3, 4]$.

Lemma A.1 (gauge forcing)**.** *(C1) forces $x_i = 0$ for every $i \not\equiv 0 \pmod 5$.*

Proof. Consider the SU(3) lines $k = 0, \ldots, 7$, charged on Q, u_R, d_R (residues $1, 2, 3$) and uncharged on e_R, L (residues $0, 4$). Applied to the pair (x_{i-1}, x_i):

- $(i{-}1 \equiv 0, i \equiv 1)$: left uncharged, right charged $\Rightarrow x_i = 0$ for $i \equiv 1$;
- $(i{-}1 \equiv 1, i \equiv 2)$ and $(2, 3)$: both charged $\Rightarrow x_{i-1} \oplus x_i = 0$;
- $(i{-}1 \equiv 3, i \equiv 4)$: left charged, right uncharged $\Rightarrow x_{i-1} = 0$ for $i \equiv 3$.

The SU(2) lines $k = 8, 9, 10$ are charged on Q, L (residues $1, 4$): the pairs $(0, 1)$ and $(1, 2)$ force $x_i = 0$ for $i \equiv 1$ again, and the pairs $(3, 4)$ and $(4, 0)$ force $x_i = 0$ for $i \equiv 4$. Finally the xor constraints of the SU(3) lines give $x_{5j+2} = x_{5j+1} = 0$ and confirm $x_{5j+3} = 0$. Hence every variable at residues $1, 2, 3, 4$ vanishes. □

Lemma A.2 (node rule)**.** *With Lemma A.1, (C2) forces $x_i = 1$ for every $i \equiv 0 \pmod 5$ in the interior, and is consistent around the cycle iff*

$$\sum_{i=0}^{N-1} \rho(i) \;\equiv\; 0 \pmod 2. \tag{56}$$

Proof. For $i \equiv 0$: $\rho(i) = 1$ and $x_{i-1} = 0$ (residue 4), so $x_i = 1$. For $i \equiv 1$: $\rho(i) = 1$ and $x_{i-1} \oplus x_i = 1 \oplus 0 = 1$ ✓. For $i \equiv 2, 3, 4$: $\rho(i) = 0$ and $0 \oplus 0 = 0$ ✓. The chain closes iff the total xor of the $\rho(i)$ vanishes. □

Lemma A.3 (counts)**.** *Write $N = 5q + r$, $0 \le r < 5$. Then*

$$\sum_i \rho(i) \;\equiv\; [r \ge 1] + [r \ge 2] \pmod 2, \qquad \sum_i x_i \;=\; q + [r \ge 1] \;=\; \left\lceil \frac{N}{5} \right\rceil. \tag{57}$$

Proof. $\#\{i < N : i \equiv 0\} = q + [r \ge 1]$ and $\#\{i < N : i \equiv 1\} = q + [r \ge 2]$; the first count is also the occupancy $\sum_i x_i$ by Lemmas A.1–A.2. □

Proof of Theorem 4.4. By Lemma A.2, satisfiability requires $[r \ge 1] + [r \ge 2]$ even, i.e. $r \ne 1$. By (C3) and Lemma A.3, $\lceil N/5 \rceil$ must be even: for $r = 0$ this means q even, i.e. $N \equiv 0 \pmod{10}$; for $r \in \{2, 3, 4\}$ this means q odd, i.e. $N \equiv 7, 8, 9 \pmod{10}$ respectively; $r = 1$ is excluded in all cases ($N \equiv 1, 6$). Conversely, for $N \bmod 10 \in \{0, 7, 8, 9\}$ the assignment $x_i = [i \equiv 0]$ satisfies (C1)–(C3) by construction, and (C4) adds $m \le \lceil N/5 \rceil$. Uniqueness is immediate: every bit is forced. □

Remark A.4 (computational verification)*.* The closed form was checked against direct GF(2) Gaussian elimination of the full system (C1)–(C3) for all $N = 5, \ldots, 200$ and floors $m \in \{1, 20, 27, 28, 29, 40\}$: zero mismatches. At $N = 137$ the unique solution has occupancy 28.

B Proof of the singlet law

Theorem 4.6 is proved in Lemmas A.1 and A.2 of Appendix A: the gauge forcing eliminates every non-e_R variable, the node rule fixes every e_R variable to one, and the occupancy is $\lceil N/5 \rceil$ by Lemma A.3. What remains to be stated is the role of the global parity:

Proposition B.1 (free singlets [T])**.** *At $N = 137$ the unique vacuum occupies 28 e_R-singlets; the global parity (C3) removes one degree of freedom from the singlet sector, leaving 27 free singlets — the pairing bill $2K + 1$.*

Proof. The singlet sector before (C3) consists of $\lceil 137/5 \rceil = 28$ independent GF(2) occupation numbers (the gauge constraints do not act on e_R). (C3) imposes one independent linear condition on them (28 is even, so it is satisfiable), leaving $28 - 1 = 27$. □

Remark B.2 (what this replaces)*.* Earlier internal formulations spoke of a "kernel of dimension 27" for a 137-row check matrix of rank 110. The correct statement is the one above: the constraint system has a *unique* solution (affine dimension 0), and 27 is the count of free singlet occupations after parity, *not* the dimension of a solution space. The numerical content (the three faces of 27) is unchanged; the linear-algebraic attribution is corrected.

C Code anatomy and the weight enumerator

C.1 Chinese-remainder construction

The rows of Φ are built from the factorization $137 = 9 \times 15 + 2$. Let $K[r] = \{i : i \equiv r \pmod{15}\}$ and $R_3 = \{1, 2, 3, 6, 7, 8, 11, 12, 13\}$ (the residues coprime-related classes of the 3-structure), $r^\star = 1$. The 13 rows are: 8 strong characters $\chi(K[r])$, $r \in R_3 \setminus \{r^\star\}$; 3 weak characters on the unions $K[1] \cup K[14]$, $K[4] \cup K[11]$, $K[6] \cup K[9]$ with the frozen edge 136 removed; the U(1) character $\chi(i \bmod 5 \in \{1, 4\})$; and the all-ones $\mathbb{Z}_2$ row.

C.2 Verification of the block structure

Direct enumeration over the 137 columns of the resulting 13×137 matrix gives 13 distinct column types with multiplicities

$$(28,\ 18,\ 9^{\times 10},\ 1), \tag{58}$$

rank 13, the 28-fold class being the e_R block (its syndrome is the single $\mathbb{Z}_2$ row) and the singleton being edge 136.

C.3 Exact weight enumerator

With 13 blocks acting independently under parity, the parity of each block is free, and the number of codewords of parity p in a block of size m is $\big((1+1)^m + (-1)^p(1-1)^m\big)/2$. The weight enumerator therefore factorizes,

$$A(z) \;=\; \prod_{b=0}^{12} \frac{(1+z)^{m_b} + (1-z)^{m_b}}{2}, \tag{59}$$

evaluated by exact integer dynamic programming over the 2^{13} block-parity patterns. Results, recomputed for this version:

$$A_2 = 891, \qquad A_4 = 332\,109, \qquad \sum_{w=0}^{137} A_w = 2^{124}, \tag{60}$$

with MacWilliams symmetry $B_w = B_{137-w}$ for the dual enumerator and minimum distance $d = 2$. The total $\sum_w A_w = 2^{124}$ is an exact integer identity (not a floating-point evaluation), confirming the dimension 124.

C.4 Syndrome table

The 13 columns (block syndromes) distribute syndrome weights binomially: the number of weight-w syndromes is $N_w = \binom{13}{w}$. Non-self-orthogonality follows from the e_R column having odd syndrome weight; hence the Z-only stabilizer reading of Section 5.

D The dynamometer: exact thermodynamics and audit note

D.1 Factorized partition function

With the teeth energy $E(x) = \lambda\,|s| - \mu_0 n_{e_R} + \mu_0 n_{\text{out}}$ and the frozen edge off, the partition function factorizes over the 13 block parities:

$$Z \;=\; \sum_{p \in \{0,1\}^{13},\, p_{12}=0} \Big[\prod_b g_b(p_b)\Big]\, e^{-\beta\lambda\,|\bigoplus_{b:\,p_b=1} s_b|}, \tag{61}$$

where s_b is the block syndrome (the common column of block b) and

$$g_b(p) \;=\; \frac{(1+w_b)^{m_b} + (-1)^p\,(1-w_b)^{m_b}}{2}, \qquad w_b = \begin{cases} e^{\beta\mu_0} & b = e_R, \\ e^{-\beta\mu_0} & \text{otherwise.} \end{cases} \tag{62}$$

The sum runs over $2^{12} = 4096$ parity patterns and is evaluated exactly (double precision, checked against rational arithmetic at the working point). All observables of Proposition 6.2 are moments of this sum; in particular $P(E_0) = e^{28\beta\mu_0}/Z$.

D.2 Thresholds

Solving $P(E_0) \in \{0.01, 0.5\}$ along the line $\lambda = \mu_0$ gives $\beta\mu_0 \in \{2.548, 3.591\}$; along $\lambda = 0$, $\{3.368, 5.277\}$. Rate ratios are $e^{\beta\mu_0}$.

D.3 Audit note on the preliminary $\langle |s| \rangle$

A preliminary internal report listed, at the working point, the three numbers $P(E_0) = 0.733$, $\langle n_{e_R} \rangle = 27.753$ and $\langle |s| \rangle = 0.234$. Recomputation for this version shows:

- the first two are reproduced *exactly* (0.73303, 27.7529) at $\beta\mu_0 = \beta\lambda = 4$;
- the third is not: the level curve $P(E_0) = 0.733$ in the $(\beta\mu_0, \beta\lambda)$ plane gives $\langle |s| \rangle$ ranging from 0.019 ($\beta\lambda = 4$) to 0.351 ($\beta\lambda = 1$), and no point reproduces all three preliminary values simultaneously; the hypothesis that 0.234 was a conditional average outside the kernel is refuted ($\langle |s| \mid \text{outside} \rangle = 1.03$).

Conclusion, per protocol A3: the published parametrization is the one reproducing the two exact targets ($\lambda = \mu_0$, $\beta\mu_0 = 4$), with $\langle |s| \rangle = 0.019$; the preliminary third value used a different, undocumented syndrome-weight convention and is withdrawn. The convention-independent outputs — the rate-ratio thresholds — are unchanged and are the robust content of the dynamometer.

D.4 Gillespie validation

Continuous-time trajectories were generated with rates $\Gamma_i = \Gamma_0 \min(1, e^{-\beta\Delta_i E})$ from three initial conditions (the ordered vacuum and two random states) for 6×10^4 jumps each. All three relax to the kernel neighborhood (syndrome weight 0–1) and to $n_{e_R} \approx 26$–28, consistent with the exact $\langle n_{e_R} \rangle = 27.7529$ (Figure 16).

E The anti-numerology protocol: worked examples

Protocol A1 requires that any formula matching a physical quantity be accompanied by the space it was found in and the probability of such a match arising by chance. We work the extremes explicitly — including, in v3, the reflexive one.

E.1 Worked example I: the golden-ratio sector dies of statistics

The two excluded φ-formulas of Section 9 look impressive in isolation. Consider the entire space of binomials

$$a\,\varphi^b + c\,\varphi^d, \qquad a, c \in \{-400, \ldots, 400\} \setminus \{0\}, \quad b, d \in \{-6, \ldots, 6\}, \quad b > d, \tag{63}$$

containing 21 489 600 candidates. Exhaustive enumeration (not sampling) shows that 37 of them match the CODATA value 137.035999177 to within 4×10^{-4} — the tolerance of the *best* golden-ratio match ever reported — a base rate of 1.7×10^{-6}. Finding "striking" φ-formulas at that precision is therefore not a signal but a statistical certainty: the space produces about 37 of them. Meanwhile, in experimental units the two candidates err by $2.2 \times 10^6\sigma$ and $1.8 \times 10^4\sigma$. Both considerations close the case: excluded, with the search space and the arithmetic on record.

E.2 Worked example II: the Catalan discriminator

The same protocol, applied to our own derivation of η. The declared search space for the fractional part was

$$\left\{\zeta(2),\ G,\ \zeta(3),\ \ln 2,\ \pi^3\right\} \times \{N,\ 2N\}\,, \tag{64}$$

ten candidates, of which $\zeta(2)/2N$-family member $\pi^2/2N$ matches the Casimir defect's leading term. The runner-up family member $2G/N = 0.01337$ does not match anything in the spectrum. With ten candidates and one hit against a derived (not fitted) target, the false-positive probability is controlled by construction — and, more importantly, the target was *computed first*: the Casimir defect (12) is a theorem about the cycle, and η is read off it, not selected from the list.

E.3 Worked example III (v3): the sweep of our own corpus

The protocol's hardest test is reflexive. Section 7.7 declares the space $\mathcal{E}$ of 24 970 400 formulas $a\eta^b + c\eta^d$ and $a\eta^b$ (integer coefficients ≤ 400, exponents in $[-6, 6]$), twenty dimensionless targets, and six tolerances — all before enumeration. The outcome cuts both ways, and both ways are published: the framework's individual lepton retrodictions sit on backgrounds of 193–6566 random matches (demoted to consistency checks), the joint profile of 4 hits in 9 targets has $p = 6.7 \times 10^{-9}$ under the uniform null, and $\alpha^{-1} = 137 + \eta$ is the unique match of the entire space at 2×10^{-7}. Example III is the protocol turned on its authors: the same enumeration that demotes our retrodictions certifies our central formula.

Remark E.1 (why these two examples). Example I shows the protocol killing a temptation; Example II shows it protecting a result. The asymmetry is the point: the same discipline that excludes the φ-sector is what entitles $\eta = \pi^2/274$ to be presented as a derivation rather than as a coincidence.

F Notation

G A plain-language guide to this paper

This appendix explains the entire paper without equations, for any curious reader.

G.1 The mystery

Physics has a magic number. It is called the fine-structure constant, and its inverse is about 137.036. It controls how strongly light and matter interact: the size of atoms, the colors of chemistry, the brightness of stars. Every experiment that has ever measured it agrees to twelve decimal places. And yet no theory explains it: in the standard textbook theory, the number is simply put in by hand. Feynman called it one of the greatest mysteries of physics. Dirac asked, already in 1935: why 137?

This paper is an attempt to answer him — and, just as importantly, an attempt to do it *honestly*, saying clearly what is proved, what is guessed, and what has already failed.

G.2 The idea in one paragraph

Imagine a necklace with 137 beads arranged in a closed ring. The beads are of five types, repeating in a fixed rhythm of five. Now impose bookkeeping rules: certain pairs of neighboring beads must balance, the whole ring must close consistently, and at least a certain number of beads must be "occupied". When you solve these rules, something remarkable happens: the rules are solvable only for necklace lengths in a very specific pattern, and the *first* solvable length that can also hold enough occupied beads is exactly 137. Independently, a completely different

symbol	meaning
N	cycle length; physical value 137
K	number of independent checks; physical value 13
C_N	cyclic causal hypergraph on N edges
$x \in \mathrm{GF}(2)^N$	occupation state of the edges
e_R, Q, u_R, d_R, L	the five edge types ($i \bmod 5 = [0,1,2,3,4]$)
$\rho(i) = [i \bmod 5 \in \{0,1\}]$	node-rule inhomogeneity
$\mathrm{SAT}(N,m)$	admissibility of (C1)–(C4) at floor m
$m^\star = 2K+1 = 27$	pairing bill / occupancy floor [D]
$\lceil N/5 \rceil$	singlet occupancy of the unique solution
Φ	13×137 check matrix over GF(2)
$\mathcal{S}_{\mathrm{vac}}$	unique vacuum solution *point* (28 singlets), never a vector space [T]
$\mathcal{C} = \ker \Phi$	the $[\![137, 124, 2]\!]$ code kernel, dim = 124 [D]
$m_b \in (28, 18, 9^{\times 10}, 1)$	block multiplicities of Φ
$A(z)$, A_w	weight enumerator of $\mathcal{C}$
$s = \Phi x$, $\lvert s \rvert$	syndrome and its Hamming weight
E_0	ordered vacuum: 28 e_R on, rest off, $s = 0$
$\eta = \pi^2/2N = \pi^2/274$	spectral asymmetry, $0.0360204540\ldots$
ΔE	Casimir defect $\cot(\pi/2N) - 2N/\pi$
$\gamma_N = (N-1)\pi/N$	cycle phase
$F = 3N/2(N+13) = 1.37$	flux factor in δ_{CP}
$\beta\mu_0$, λ	bath bias and syndrome coupling (teeth model)
Γ_i	CTMC flip rate of bit i
$W_{\mathrm{spec}}, W_{\mathrm{sat}}$	the two selection windows

Table 9: Notation used throughout the paper.

calculation — about the musical notes, so to speak, that such a ring can sustain — also singles out a small range of lengths ending at 137. Two independent mechanisms, one survivor. That is the heart of the paper.

G.3 Where does the 0.036 come from?

The magic number is not exactly 137; it is 137.036... The framework computes the fractional part as $\pi^2/274 \approx 0.03602$. Two different mathematical routes give the same number: one counts the “energy defect” of the ring compared to a smooth circle (a calculation physicists call a Casimir sum, and mathematicians call zeta regularization), the other sums a famous series of fractions whose total Euler already knew ($\pi^2/2$, the same family as the celebrated $\pi^2/6$). The two routes agree to better than one part in a million. That agreement is one of the strongest internal checks of the framework.

G.4 The honest part: it does not match perfectly

Here is what most proposals of this kind hide, and this paper refuses to hide: the prediction 137.03602 differs from the measured 137.035999177 by about 2 parts in 10 million. That sounds tiny, but experiment is fantastically precise: the gap is about *a thousand times* the experimental uncertainty. So the raw formula is *wrong at the level the experiment can see.* The framework treats this gap as its most important open problem (the running of the constant with energy might account for it, but nobody has derived that yet), not as a detail to be swept away.

G.5 What else does the framework say?

Quite a lot, and much of it can be checked by experiments already running:

- The angle that governs how neutrinos from the Sun oscillate comes out at 0.30365; the world average is 0.304 — a near-perfect match, though the derivation is not yet complete.
- The ratio of the two neutrino mass splittings is predicted to be 34.07; the JUNO experiment in China will soon measure it precisely enough to confirm or kill this.
- The amount of CP violation in neutrinos — related to why the universe contains matter at all — is predicted as a specific value, 1.211π. The DUNE experiment (USA) and Hyper-Kamiokande (Japan) will reach the needed precision this decade. If they measure something else, this pillar of the framework falls. The paper calls this, deliberately, a "registered death sentence".
- The gravitational-wave background of the early universe should have a specific intensity ($r = 0.021$) and a specific rhythmic pattern, testable by the LiteBIRD satellite and the CMB-S4 observatory around 2030.

G.6 The code and the price of order

A newer part of the framework reads the bookkeeping rules as an error- correcting code — the same mathematics that protects data in a hard drive. The 137 beads with 13 check rules form a code protecting 124 bits of information. The paper computes exactly how many errors of each size this code tolerates (the numbers 891 and 332 109 in the text are exact, not rounded).

Then comes a sobering result. If the universe's "bath" treated all bead flips equally, the probability of finding the perfectly ordered vacuum would be 2^{-124} — a number with 38 zeros after the decimal point. Order is astronomically improbable unless the bath *prefers* productive flips over destructive ones. The paper quantifies exactly how strong that preference must be (a factor of about 36 for the vacuum to occupy half the probability). What it cannot yet do is explain *why* the bath would have such a preference. That failure is stated openly, as open problem E12.

G.7 The number 27, three times

A supporting character steals several scenes: the number 27. It appears as the "pairing bill" of the 13 check rules (two active beads per rule, plus one), as the number of free "singlet" beads in the unique solution (28 positions of the fifth type, minus one global constraint), and as the minimum occupancy that makes the necklace length 137 the first possible one. Three different roles, one number. The framework elevates this coincidence to a principle — explicitly labeled as *postulated*, not proved.

G.8 What has already failed

The paper keeps a public list of ideas that did not survive contact with data: a formula for dark energy, one for the neutron lifetime anomaly, one for the cosmic birefringence, one for the matter–antimatter asymmetry (excluded at 21 standard deviations), two elegant formulas involving the golden ratio (excluded by enormous margins), and an earlier misstatement about the fine-structure residual being a small tension (it is not — see above). Each failure is printed in a table. In this framework, being killed by data is not shameful; hiding the corpse would be.

G.9 How to read the labels

Every claim in the paper wears a colored badge: [**T**] proved theorem; [**D**] a definition or postulated principle (a choice, openly declared); [**P**] a prediction an experiment can kill; [**C**] a conjecture; [**X**] an excluded reading. If you remember one thing: the badges are the point.

The framework's contract with the reader is that every number can be checked with a pocket calculator, and every pillar has an experiment already scheduled to test it.

G.10 The bottom line for a lay reader

A ring of 137 beads, five bead types, thirteen bookkeeping rules. Out of this toy, a research program extracts: the integer part of physics' magic number (selected by two independent mechanisms), the first digits of its fractional part (by two independent calculations), four numbers that experiments will confirm or kill within this decade, and an honest list of everything that does not work — headed by a thousand-sigma gap that the program openly calls its central unsolved problem. Whether nature uses this ring is for the experiments to say. The framework's job is to be checkable. It is.

H Numerical audit trail

Every number printed in this paper was recomputed from the printed formulas during the preparation of this version. Table 10 lists the main values, their sources, and the check status.

quantity	printed value	recomputed	status
$\eta = \pi^2/274$	0.0360204540	0.03602045400 . . .	exact match
$\alpha^{-1}_{\text{pred}}$	137.0360204540	sum of the above	exact
$3\pi\lvert\Delta E\rvert - \eta$	3.157×10^{-7}	$\pi^4/120N^3$	exact
residual vs CODATA	2.13×10^{-5}, 1013σ	$2.128 \times 10^{-5}/2.1 \times 10^{-8}$	match
residual vs Rb	2.12×10^{-5}, 1932σ	$2.125 \times 10^{-5}/1.1 \times 10^{-8}$	match
window bounds	$N \in \{132..137\}$	$\lfloor 8N/\pi^2 \rfloor = N - 26$	exact
SAT closed form	$N \bmod 10 \in \{0, 7, 8, 9\}$	GF(2) solve, $N \leq 200$	zero mismatches
singlet occupancy	$\lceil 137/5 \rceil = 28$	direct solution	exact
convergence	$W_{\text{spec}} \cap W_{\text{sat}} = \{137\}$	set arithmetic	exact
A_2, A_4, $\sum A_w$	891, 332109, 2^{124}	exact integer DP	exact
$P(E_0)$ at $\beta\mu_0 = 4$	0.7330	factorized Z	exact to 4 d.p.
$\langle n_{e_R} \rangle$	27.7529	factorized Z	exact to 4 d.p.
thresholds	2.548/3.591, 3.368/5.277	root solve on Z	match
2^{-124}	4.702×10^{-38}	direct	exact
$\sin^2\theta_{12}$	0.30365	$(1 - \pi/8)/2$	exact
$\lambda = \pi^3/137$	0.22632	direct	exact
Δm^2 ratio	34.07	$4(1 - \eta)/\pi\eta$	exact
δ_{CP}	1.211π	$\gamma_N F - 13\eta$	exact
m_{light}	3.3×10^{-5} eV	$\eta^2/2 \cdot \sqrt{\Delta m^2_{31}}$	match
r, $A_{\log}$, d_s	0.021, 6.5×10^{-4}, 3.964	powers of η	exact
σ_{DM}/m	4.7×10^{-29} cm^2/GeV	$\eta^3 \times 10^{-24}$	exact
η_B [X]	6.95×10^{-10}, 21σ	η^6/π	match
τ ratio [X]	1.036, $\sim 11\sigma$	$1 + \eta$	match

Table 10: Audit trail: every printed number recomputed. "Exact" means the printed digits agree with the recomputed value at all printed places.

I Reproducibility

Every computation in this paper can be reproduced in a few lines of standard Python. We give the core verifications in pseudo-code form.

The decomposition and the residual.

```
from math import pi
eta  = pi**2/274            # 0.03602045400...
```

```
pred = 137 + eta            # 137.0360204540...
codata, sigma = 137.035999177, 2.1e-8
(pred - codata)/sigma       # 1013.2  -> the residual in sigmas
```

The closed form of SAT. The system (C1)–(C3) is linear over GF(2): build the N node-rule rows $x_{i-1} \oplus x_i = \rho(i)$, the all-ones parity row, and the gauge forcings from the charge table (7); solve by GF(2) Gaussian elimination. Doing this for $N = 5, \ldots, 200$ reproduces the rule $N \bmod 10 \in \{0, 7, 8, 9\}$ with zero mismatches, and at $N = 137$ returns the unique solution $x_i = [i \equiv 0 \bmod 5]$ with occupancy 28.

The weight enumerator. Dynamic programming over the 2^{13} block parities of Table 3 with exact (arbitrary-precision) integers reproduces $A_2 = 891$, $A_4 = 332\,109$ and $\sum_w A_w = 2^{124}$ exactly.

The teeth thermodynamics. The factorized partition function of Appendix D is a sum over 4096 parity patterns; at $\beta\mu_0 = \beta\lambda = 4$ it returns $P(E_0) = 0.73303$ and $\langle n_{e_R} \rangle = 27.7529$ in double precision.

The golden-ratio sweep. Exhaustive enumeration of the 21 489 600 binomials $a\varphi^b + c\varphi^d$ of Appendix E (nested loops, $\approx$ one minute on a laptop) counts 37 matches within 4×10^{-4}.

J Errata: summary and repository pointer

From Version 5 on, the full dated errata — with scripts, logs and complete recomputes — live in the supplementary files of the Zenodo record (doi:10.5281/zenodo.21855881), alongside the internal research reports that generated them (campaigns 2026-08; research 96 and 97, dated 2026-08-23). This section keeps the one-page summary. The Brazilian-Portuguese plain-language annex of earlier versions is no longer part of the paper; it is archived in the same Zenodo record for the project's Portuguese-speaking readers.

entry	date	content
v1→v2 correction map	v2	residual is 2.13×10^{-5} ($\sim 10^3\sigma$), not "1.9σ"; v1 kernel conflation corrected (Rem. 2.6)
E1 (correction)	2026-08-08	two boundary decimals of the window-range proposition (below)
E2 (strengthening)	2026-08-08	exact unbiased no-go is 2^{-137}, not 2^{-124} (below)
audit campaign notes 1–11	v4	Section K
research 96 errata 2ł/3ł	2026-08-23	exact SCGF derivatives; corrected single-ring limit $Q_\infty(\varepsilon) = (8/3)\,\mathrm{atanh}(\varepsilon)/\varepsilon$ (companion TUR paper, doi:10.5281/zenodo.22062270)
research 97 errata 1ł	2026-08-23	sign artefact in truncation check; failed derivation of the quarantined $9/(2\pi^2)$ printed as a lost bet (Note 10)
research 98	2026-08-23	bill $2K + 1$ promoted to conditional theorem (Thm. 4.7); redaction errata in Sec. 4.4
research 99	2026-08-23	flux sector = e_R block [T]; additivity link lost [X]; quarantine status updated (Note 10)
research 100	2026-08-23	quarantine shown *structural* (Note 10); third independent derivation of the count 27 (block curvature, [T])
research 101	2026-08-23	two external physical channels (local insertions, RG scale) computed null with printed numbers (Note 10); PA2 campaign closed as open-structural

Version 4 did not alter any line of the Version 3 body. The two items below were found by the 2026 internal audit campaign and are recorded here with their discovery dates. Neither changes any conclusion of the paper; the second makes one conclusion strictly stronger.

J.1 E1 (correction): the two boundary decimals of the window-range proposition, recorded 8 August 2026

In the window-range proposition of Section 4, the sentence

> "the boundary values are $8N/\pi^2 = 106.9745\ldots$ at $N = 132$ and $110.2535\ldots$ at $N = 138$ (the latter already outside)"

mislabels the two decimals. The correct values, recomputed at 40 digits, are

$$\left.\frac{8N}{\pi^2}\right|_{N=132} = 106.9951699263\ldots, \qquad \left.\frac{8N}{\pi^2}\right|_{N=138} = 111.8585867411\ldots \tag{65}$$

The floors are unchanged — $\lfloor 106.9951\ldots \rfloor = 106 = 132 - 26$ and $\lfloor 111.8585\ldots \rfloor = 111 \neq 112 = 138 - 26$ — so the window $\{132, \ldots, 137\}$, the exclusion of $N = 138$, and every conclusion drawn from them stand exactly as printed. Probable origin of the typo: the printed decimals reproduce the fractional parts of the *analytic bounds* of the window, $131.974502\ldots < N \leq 137.253482\ldots$ (Note 1 of Appendix K); they were mislabelled as values of $8N/\pi^2$ in writing. The underlying mathematics was correct; the sentence was not.

J.2 E2 (strengthening, not a correction): the exact unbiased no-go is 2^{-137}, not 2^{-124}, recorded 8 August 2026

Throughout the text (the Version 2 abstract paragraph, the tables of Section 1.4, Section 6, the dynamometer appendix, the conclusions, and the numerical audit trail) the unbiased-vacuum no-go is quoted in the short form $P(E_0) = 2^{-124} = 4.702 \times 10^{-38}$. The printed proposition itself is correct as stated — it reads $P(E_0) = 2^{-124} \times (\text{kernel weight}) \leq 2^{-124}$ — but the short-form quotation drops the kernel factor. The exact value under the uniform stationary state is

$$P(E_0) \;=\; 2^{-124} \times 2^{-13} \;=\; 2^{-137} \;=\; 5.7397185 \times 10^{-42}, \tag{66}$$

a factor $2^{13} = 8192$ *stronger* than the quoted bound: the ordered vacuum requires not only the 124 code bits but all 13 check values to be simultaneously right. (Under the frozen-edge convention of the teeth model the blind-bath baseline is 2^{-136}, one factor of 2 above the uniform $\mathrm{GF}(2)^{137}$ value.) Every use of the no-go in the paper bounds from above the plausibility of an unbiased origin, so all conclusions are strengthened and none is altered. The same audit quantified the gap E12 in nats: reaching the working point $P(E_0) = 0.733025$ from any unbiased baseline requires ≈ 86 nats of occupation bias in the 28-site e_R block (≈ 3.1 nats per site), while the entire syndrome sector can supply at most a factor 2^{12}, i.e. 8.3 nats (the parity-to-syndrome map is injective — Note 9 below). E12 is unchanged in status $[\![X]\!]$: it is now a gap with an exact size.

K Notes from the 2026 internal audit campaign

Between 8 and 9 August 2026 the framework underwent twelve pre-registered internal audits, each with its search space and acceptance criteria declared in writing *before* any computation, per the anti-numerology protocol of Appendix E. Every number below was recomputed at 40-digit precision; every claim carries its epistemic tag; negative results are reported with the same weight as positive ones; no published line was rewritten retroactively. The notes are self-contained; the full dated declarations, reports and reproduction scripts are archived alongside this version. Experimental inputs frozen in the body (PDG 2024, NuFIT 6.0, CODATA 2022) remain frozen; newer data appear only as explicitly dated notes.

K.1 Note 1 — the spectral window is closed for all N, and the convergence is unique in the declared range [T]

Three strengthenings of the selection of N (Section 4), proved during the audit of 8 August 2026:

(a) **Closure for all N [T].** Writing $f(N) = 8N/\pi^2 - (N - 2K)$, the window condition $\lfloor 8N/\pi^2 \rfloor = N - 2K$ is $0 \leq f(N) < 1$; since f is strictly decreasing, the solution set is the single contiguous interval $(2K-1)/(1-8/\pi^2) < N \leq 2K/(1-8/\pi^2)$. For $K = 13$: $131.974502\ldots < N \leq 137.253482\ldots$, whose integer part is exactly $\{132, \ldots, 137\}$ — with no $N \leq 200$ restriction (the body certified the window only up to 200; an exact scan to $N = 200\,000$ agrees).

(b) **Plateau uniformity [T].** Every admissible block $\{10k{+}7, 10k{+}8, 10k{+}9, 10k{+}10\}$ has uniform capacity $\lceil N/5 \rceil = 2k + 2$; hence $A(m) = 10k + 7 \iff m \in \{2k{+}1, 2k{+}2\}$: *every* plateau has length exactly 2. The plateau $\{27, 28\}$ of $N = 137$ is structurally typical — what singles out 137 is the convergence plus primality, not the shape of its plateau.

(c) **Uniqueness of the convergence in $K = 1 \ldots 40$ [T, exhaustive in the declared interval].** Requiring simultaneously (i) $\max W(K)$ prime, (ii) that prime unique in the window, and (iii) $\max W(K) = 10K + 7$ (start of the SAT block of matching index), the three conditions hold jointly only at $K = 13$ ($N = 137$). The mechanism is transparent: $\max W(K) \approx 10.5572K$ crosses $10K + 7$ twice in the range — at $K = 13$ (137, prime) and $K = 14$ ($147 = 3 \times 7^2$, composite). At $K = 32$ the window maximum is $337 = 10 \times \mathbf{33} {+} 7$: a unique in-block prime, but the block index (33) no longer equals K — a documented spurious convergence (the "337 echo"), recorded here as the standing background measurement against any future pattern involving 337 in the framework.

K.2 Note 2 — the counting law is isolated as the plateau pair $\{2K+1, 2K+2\}$; "parity eats a singlet" [T computed, declared family]

The singlet-covering principle ($m^* = 2K+1$, Section 4.5) remains postulated [D] — nothing here derives it — but it is now *characterised*. In the declared family of 20 linear counting laws $f(K) = aK + b$ ($a \in \{0, 1, 2, 3\}$, $b \in \{-1, \ldots, 3\}$), scanned over $K = 1 \ldots 40$ against both windows, exactly one class converges strongly at $K = 13$: the plateau pair $\{2K+1, 2K+2\}$ (floor and capacity of the same block, by plateau uniformity). The only other strong convergence in the family, $2K+3$ at $K = 32$, is the 337 echo of Note 1 and fails the matched-index criterion. The sharpened reformulation — *parity eats a singlet*: the free capacity must cover the bill, $\lceil N/5 \rceil \geq 2K+2$, because C3 consumes exactly one singlet — selects $\{137\}$ directly through the residue filter, without invoking primality, and agrees with the printed selection path on the whole declared domain. The two "+1"s of the two faces of 27 become one and the same fact: the global parity. The open problem PA9 is thereby reduced, not closed: "why $2K+1$?" becomes "why the plateau $\{2K+1, 2K+2\}$?" — a smaller, arithmetic question.

K.3 Note 3 — two dynamical quantifications [T]

(a) **The pair sector carries** 85% **of the disorder time.** Conditioned on being outside the ordered vacuum E_0 (26.6% of the total time at the working point), the trajectory spends a fraction 0.8507 (Gillespie) / 0.8499 (exact conditional) in the pure-pair sector $n_{\mathrm{def}} = 2$, and 96.7% of the outside time has even n_{def} — a dynamical confirmation of the syndrome-algebra theorem (defects enter and leave in pairs), quantifying the body's qualitative statement "pair sector as the slow mode". (b) **The PA2 window background is** $6/98 = 6.1\%$**.** In the declared family of 98 formulas built from exact ensemble moments (63 moment-dilution, 32 exponential, 3 spectral), 6 members land inside the pre-registered $\pm 30\%$ window around $-0.4553\,\eta^3$: falling in the window costs about 6% of background — neither trivially cheap nor expensive. The quarantined $2e^{-8}\eta$ ($c = -0.5171$) is only the 4th-closest to the centre among the 6; three moment-dilution formulas hit closer and no one would ever have noticed them without the enumeration. This is the quantified selection effect that the quarantine policy exists to block, and it calibrates every background estimate used in Notes 7, 8 and 10. Separately, the dynamometer verdict of Section 6.4 was verified on the whole level curve $P(E_0) = 0.733$ (13 points), not only at the working point: zero cures inside the window everywhere.

K.4 Note 4 — every published prediction is an "N-meter" [C, with dated 2026 experimental note]

If $X = f(N)$ with f explicit and published, then $N_X = f^{-1}(X_{\mathrm{meas}})$ is an independent experimental determination of the integer N. Six inversions (closed, declared set): the Cabibbo angle $\lambda = \pi^3/N$; the neutrino ratio $R = \Delta m^2_{31}/\Delta m^2_{21} = 4(1-\eta)/(\pi\eta)$; the tensor ratio $r = 4\pi^4/N^2$; the spectral dimension $d_s = 4 - \eta$; the dark-matter cross section $\sigma/m \propto \eta^3$; and δ_{CP} (numerical inversion). With the frozen editions (PDG 2024 [15], NuFIT 6.0 [8], BK18+PR4 [23]):

channel	measured N	pull vs 137	status
λ (Cabibbo)	137.800 ± 0.416	$+1.92\sigma$	active — dominant tension
R (NuFIT 6.0)	136.41 ± 3.88	-0.15σ	active — consistent
$r < 0.032$ (95% CL)	$N > 110.35$	—	bound only — consistent
d_s, σ/m	no measurement	—	inactive
δ_{CP}	73.6 ($\sigma_N \geq 13.7$)	weak channel	reported in full, statistically inconclusive

Joint consistency of the two sub-percent channels against $N = 137$: $\chi^2 = 3.709$ for 2 d.o.f., $p = 0.156$ — consistent. The δ_{CP} inversion is reported with the same weight as every other number: the function $\delta(N)$ is nearly flat (0.00635 rad per unit), so the channel carries its information in δ_{CP} space — the already declared $\sim 2\sigma$ tension — and not as a measurement of N. **Dated 2026 note (frozen editions unchanged):** with JUNO/*Nature* 654 (10 June 2026) [16] and NuFIT 6.1 [17], $N_R = 136.66 \pm 1.77$ (-0.19σ) and $\chi^2 = 3.722$, $p = 0.156$; and the pure-π solar prediction $\sin^2\theta_{12} = (1 - \pi/8)/2 = 0.3036504591$ meets the JUNO value 0.3036 ± 0.0064 at $+0.008\sigma$ — by pre-registration this is a check of the geometric sector and *not* evidence about N, since N does not occur in the formula. **Projection [P]:** JUNO 6-year data plus PDG give $\sigma_N \approx 0.38$ — experiments will resolve individual integers, exposing $N = 137$ to the hardest test an integer can face.

K.5 Note 5 — the "η-meter" table and the link test [C]

The same inversion applied to η: R measures $\eta = 0.0361766 \pm 0.0010285$ at the MeV scale ($+0.15\sigma$ against $\pi^2/274$; dated 2026 note: 0.0361092 ± 0.0004688, $+0.19\sigma$); the tensor bound gives $\eta < 0.04472$ at $\sim 10^{16}$ GeV (consistent); the logarithmic-oscillation amplitude $A_{\log} = \eta^2/2 = 6.49 \times 10^{-4}$ lies 35× below the current Planck+SPT-3G bound [24] (untested); d_s and σ/m are inactive. The atomic channel $\alpha^{-1} - 137$ is circular by construction [D] and never counts. **New internal test [P]:** treating η and N as independent parameters, λ measures N alone and R measures η alone, and the framework's link predicts $\eta_R = \lambda/(2\pi)$. Data: $\eta_R - \eta_\lambda = +3.65 \times 10^{-4} \pm 10.34 \times 10^{-4}$ — the link is confirmed at 0.35σ with zero adjustment (0.62σ in the dated 2026 edition), and JUNO 6-year data will take the test to $\sigma_{\text{link}} = 2.7 \times 10^{-4}$ (0.75% of η). η is thus measured independently across ~ 25 orders of magnitude in scale and is constant within errors in every non-circular channel; the anomaly lives exclusively in the one circular channel, precisely where precision is $10^4\times$ higher.

K.6 Note 6 — four routes exclude a varying η; the residual is pointlike [T/C]

Could η be a variable scalar whose running absorbs the α^{-1} discrepancy? Four independent routes were tested; three are excluded and the fourth is exactly PA2:

V1. Running in q^2 (dead by geometry). The experimental anchors of α (electron $g-2$, Rb/Cs recoil) live at $q^2 \approx 0$ — the same scale as the definition $\alpha^{-1} = N + \eta$. The time-like direction worsens the gap (α^{-1} falls to 127.955 at m_Z); the space-like direction would require the law $N + \eta$ to hold at $Q^* = 16.18$ keV $= 0.0317\, m_e$ — a scale that occurs in no declared structure of the framework. Specifiable, but unanchorable.

V2. Discrete graph mechanism (excluded by counting). The smallest internal move, $\delta N = \pm 1$, shifts α^{-1} by 0.99974 — a factor 46 987× larger than the needed 2.1277×10^{-5}. Nothing made of the integers of C_{137} can produce a continuous 10^{-5} correction.

V3. Cosmological/temporal variation (excluded at $\sim 95\%$ CL). The residual amounts to $\delta\alpha/\alpha = +1.553 \times 10^{-7}$, which is 2.35× the Oklo upper bound ($+6.6 \times 10^{-8}$, 95% CL [20]) and 1.29× the Damour–Dyson bound [21]; atomic clocks measure a null drift, $1.0(1.1) \times 10^{-18}$/yr [22] — accumulating 1.55×10^{-7} at that rate would take 155 Gyr.

V4. Continuous internal modulation (the only honest slot) — and it is literally PA2, already swept: zero declared candidates inside the window (Notes 3, 7, 10). The slot stays $[\![X]\!]$ open, now with five explicit requirements any future mechanism must meet (right size in the declared window; no disturbance of the consistent channels; static today and dead at Oklo; derivable from declared structures; ideally explaining why the correction is $\sim \eta^3$).

The residual itself is robust across the four independent determinations of α (CODATA 2022 [1]: 1013σ; Rb 2020 [2]: 1932σ; Cs 2018 [18]: 793σ; $g{-}2$ 2023 [19]: 1419σ); the internal Cs$\leftrightarrow$Rb experimental tension (5.5σ combined) is $133\times$ smaller than the residual. No averaging, no choice of experiment, no known systematics absorbs the conflict — and no scale variation can either. Verdict: η is constant; the α residual is what it was always declared to be, a pointlike discrepancy $-0.4553\,\eta^3$ whose origin remains open $[\![X]\!]$.

K.7 Note 7 — the singleton is self-frozen [T/C]; route R1 against PA2 is closed [C]

Is the frozen edge 136 (constraint C4) an arbitrary decree? Releasing the singleton as an ordinary site with the same coupling ($\beta\mu_1 = \beta\mu_0 = 4$), the exact factorised ensemble ($Z_{\rm deg} = Z_{\rm frz}(1 + \kappa e^{-\beta\mu_1})$, validated against direct dynamic programming at 10^{-14}) gives

$$P(n_s = 1) \;=\; 3.86 \times 10^{-4}, \qquad \kappa = 0.021074 : \tag{67}$$

the teeth dynamics itself keeps the edge OFF 99.96% of the time with no decree — and still 97.9% in the limit of maximal coupling ($\mu_1 \to 0$). The thaw coefficient is dominated (56%) by the "repairing" branch $P_{11}e^{+8}$: the rare configurations ($P_{11} = 4 \times 10^{-6}$) in which switching the singleton ON *reduces* the syndrome weight by 2. In 99.1% of the weight, switching it ON costs e^{-8}. C4 is therefore not an imposition against the dynamics but a good approximation *of* the dynamics (error 0.04% in the partition function, $\leq 2\%$ in the worst observable) — the first dynamical result that *justifies* a structural constraint instead of assuming it. As a PA2 route (R1, the most physically motivated of the menu), it fails cleanly: a closed family of 24 candidates built from the six thaw observables lands zero times inside the window (closest: $-\sqrt{\kappa} = -0.145169$, a distance 0.174 from the edge), and the declared landscape $\beta\mu_1 \in [2, 12]$ never touches the window at any point. No hidden tuning point exists.

K.8 Note 8 — the attractor parity theorem [T]: E12 shrinks to a degeneracy of two

Exact anatomy of the attractor at the working point $(4,4)$: $\ln(1/P(E_0)) = 0.31057$ nats decomposes as gross cost 2.4684 (1.9602 entropy of the 108 repulsive sites $+\,0.5082$ e_R-tail) minus return 2.1578 from the syndrome structure — parity gives back 87% of the entropic cost; the attractor is made by parity, not by energy alone. The working point sits 0.4220 above the majority line $P(E_0) = 1/2$: the attractor is a phase with margin, not a fine-tuned point (smooth crossover everywhere, as expected for 136 sites). The central result:

[T] Theorem — attractor parity theorem

For an attractive class of multiplicity m with coupling $w^{\pm} = e^{\pm 4}$: if m is *even*, $g_p(e^4) = e^{4m} g_p(e^{-4})$ — making the class attractive merely rescales the parity weights — so all even classes saturate with identical probability; if m is *odd*, the preferred parity *flips*, and full saturation carries a whole-column syndrome penalty $e^{-4\,\mathrm{wt}}$. In C_{137} there are exactly two even classes, e_R ($m = 28$) and the weak pair class $\{1, 14\}$ ($m = 18$), and they *tie exactly* ($P_{\rm sat} = 0.7330254595$ for both; verified at three couplings, differences $\leq 6 \times 10^{-17}$, float64 noise). The ten odd classes ($m = 9$) lose by factors from $513\times$ to $1.5 \times 10^6\times$.

The vacuum-selection question "why e_R?" therefore has 2 candidates, not 12: E12 shrinks to a degeneracy of two, which the teeth dynamics does not break. Structural fine point: the exclusion $i \neq 136$ that creates the singleton is exactly what makes $\{1, 14\}$ even (with 136 included the class would have 19 sites — odd — and e_R would be the unique optimal attractor): *the singleton creates the second even class.* The selection of e_R comes from structure (the node rule C2, which also anchors the lepton sector and the lightest column), not from energetics; the term $-\mu_0 n_{e_R}$ remains postulated, and E12 remains $[\![X]\!]$ — now with size 2.

K.9 Note 9 — the parity-to-syndrome map is injective [T]

The 12 active block-columns of the constraint matrix are linearly independent over GF(2): the map from block-parity patterns to syndromes is injective, all 4096 patterns give distinct syndromes, and only the null pattern has zero syndrome weight. (Found through a division by zero in an unrelated audit computation — $\ln s$ at $s = 1$ — and promoted to a theorem by exhaustion.) Consequences: the syndrome sector carries maximal information content per block, which is precisely why its bias ceiling is the $2^{12} = 8.3$ nats recorded in E2; and the singleton's column ($\mathrm{U}(1) \oplus \mathbb{Z}_2$) adds no degeneracy to the check structure.

K.10 Note 10 — PA2 completeness: seven declared routes, all null [C]; a permanent promotion criterion; one sealed quarantine, and the conditional closure via loop anchors (v6)

Every declared route to the coefficient $c^* = -0.4553$ of the residual $\Delta(\alpha^{-1}) = -0.4553\,\eta^3$ has now been travelled, each with its search space declared before computation:

route	campaign	result
dynamometer (ensemble mechanisms)	v3 body	swept; 5-requirement checklist issued
R1 — singleton thaw	2026 audit	0/24 candidates in window; landscape null (Note 7)
R3 — spectral invariants of C_{137}	2026 audit	10 window hits, all background ($N = 131$ scores *better* than 137)
R4 — rational background of the quarantine	2026 audit	quarantine absolved as background: $p \geq 0.32$; six forms archived
R5 — two-attractor mixture	2026 audit	dead twice: $\beta\mu_2^* = -0.392$ (unnatural) + broken anchors
R2 — broad screening, two levels	2026 audit	background: $p = 0.21$ (level A, 3 191 forms); $p = 0.20$ (level B, 566 113 forms)
R6 — Casimir tail (2026-08-23)	research 97	excluded *by sign*: $3\pi\lvert\Delta E\rvert - \eta = +\pi^4/(120N^3) + \cdots > 0$; a positive tail cannot supply the negative η^3 term

Under the standing null model (coefficient uniform in the declared window, background calibrated at 6.1% — Note 3), no candidate separates from background: the best level-B hit, 173/380, matches $|c^*|$ to 1.7 ppm ($\tau = 1.65 \times 10^{-6}$) and still has $p = 0.20$ — with half a million declared forms, a random target has a 20% chance of such a hit. It is the sharpest demonstration of the project that **numerical proximity is not evidence**. Verdict $[\![C]\!]$: c^* has no simple arithmetic structure detectable above background — not among $\sim 5.7 \times 10^5$ general forms, not among forms built from the framework's own structural numbers, not in the hypergraph spectrum, and not via any declared dynamical mechanism. *c^* is the number of the failure, not a clue.* PA2 remains $[\![X]\!]$, now with a complete exclusion map. **Permanent promotion criterion:** a candidate leaves the background archive only with $p < 0.01$ in a declared space *plus* a derivation; what would reopen PA2 is a new mechanistic derivation (a physical channel not listed above) — never a raw proximity. **Sealed quarantine.** The form $\ln 4/\ln(13+8)$ — a later post-hoc proposal — is not merely quarantined but *structurally forbidden*: the framework's own 13 decomposes as $13 = 8+3+1+1$ (the 13 rows of Φ: 8 strong residues + 3 weak pairs + U(1) + $\mathbb{Z}_2$, confirmed independently in two separate corpora of the project), so $8 \subset 13$ and the sum $13+8$ double-counts. Two external full-corpus sweeps (a second AI auditing all project documents) found no structural 21 and no physically derived sum $13+8$ anywhere; the proposal is archived with this seal. **Open quarantine (dated 2026-08-23, internal research 97).** The window-defect dressing

$$\Delta \;=\; 27\left(\pi \sin\tfrac{\pi}{2N} - \eta\right) \;=\; -\frac{9\pi^4}{16\,N^3} + \frac{9\pi^6}{1280\,N^5} - \cdots \qquad \Longrightarrow \qquad c \;=\; \frac{9}{2\pi^2} \;=\; 0.4559453, \tag{68}$$

with $27 = \lceil N/5 \rceil - 1 = 2K+1$ the framework's own singlet count, hits the pre-registered window at 0.149% from $c^* = -0.4553$ ($\Delta_{\text{cand}} = -2.1308725 \times 10^{-5}$ vs. $\Delta^* = -2.1277019 \times 10^{-5}$, exact at 80 digits). In the declared space of integer dressings $k \in [20, 36]$ the window admits 16 of 18 integers ($p \approx 0.89$), and the finite-N correction $-9/(160N^2)$ moves the candidate *away* from c^*, leaving a second-order residue that any future mechanism must also explain. **Status update (v5.1, internal research 98–99, dated 2026-08-23).** The mechanism underlying this candidate is now a four-link chain, three of whose links are theorems: (i) the *count* 27 is the unique self-consistent choice of the CRT family (Theorem 4.7, conditional on one stated postulate); (ii) the *flux sector* of the vacuum is exactly the e_R block, the unique weight-one syndrome class, with 27 independent toggles [T, verified bit by bit]; (iii) the window defect is exactly the curvature of the lowest antiperiodic mode, $\pi \sin(\pi/2N) = (\pi/2)\,\lambda_{1/2}$ vs. $\eta = (\pi/2)\,\lambda^{\text{lin}}$ [T]; (iv) the *additivity* of 27 identical defect contributions is *not* derivable: the 27 toggles are representatives of a single global $\mathbb{Z}_2$ flux, not 27 independent fluxes — this link is printed as a *lost bet* [X]. **Structural verdict (internal research 100, dated 2026-08-23): the quarantine is structural, not circumstantial.** All three candidate extensivity mechanisms were computed and fail for structural reasons: the Coleman-type twist energy is positive and representative-independent (translation invariance: 27 representatives are 27 cuts of the *same* state; verified exactly); the 2^{27} mixing entropy misses the target scale by a factor 41 with no natural normalization; and the e_R-block curvature, although it yields the count 27 exactly a *third* time ($\varphi''(0) = m(m-1) = 756$, coordination $m-1 = 27$ of the parity graph K_{28} [T]), lives entirely in the combinatorial sector and does not couple to π. No functional printed in this framework sums spectral contributions over the twist representatives: the framework *as it stands* does not contain the object PA2 asks for. The two remaining physically motivated channels were then computed and are null with printed numbers (internal research 101, dated 2026-08-23): the second-order self-energy of the half-mode from unit local insertions on the e_R sites has the *right sign* but overshoots by $F = 93\,474.48$ with no declared structural factor matching, and the RG-running reading requires an arbitrary scale $\mu^*/m_e = 1.00010027$ ($\sim m_e +$ 51 eV) derived from nothing — the framework's first adjustable dial, excluded by design. **Status update (v6, internal research 107–111, dated 2026-08-23): promoted to derived-conditional.** The form $9/(2\pi^2)$ no longer sits in quarantine: the factor 27 is now the count of distinct Wilson-loop anchors in the flux sector (28 e_R addresses, energetically identical, exactly one feeding the window defect in the forward orientation), conditional on one declared postulate [D] — *the anchor that feeds the defect is absorbed by it* — printed below in this Note. The numerical content is unchanged: $27\,g_1 = -2.1308725 \times 10^{-5}$, 0.149% from the exact residual, whose fine residue (0.0402 copies of g_1) is named and open. The older wording — "PA2 is not a missing number but a missing object" — proved exact: the v6 resolution came from a new object (the loop-anchor structure) with one declared postulate, never from a rearrangement of the existing ones. **Open quarantine, second entry (dated 2026-08-23).** The RG-scale reading leaves one exact window hit behind: had the framework's natural scale been $\mu^* = m_e(1+2\eta^3)$, the running shift would be $-\frac{2}{3\pi}\ln(1+2\eta^3)$, i.e. an implied coefficient

$$c \;=\; -\frac{4}{3\pi} \;=\; -0.424413, \tag{69}$$

which falls *inside* the pre-registered window $[-0.592, -0.319]$ at 6.78% from c^*, covering 93.2% of the gap and leaving a remainder of -1.44×10^{-6} ($\approx 68.7\sigma$). It is quarantined, not promoted, for three printed reasons: the scale $1 + 2\eta^3$ is the closest of four declared candidates — a post-hoc window hit, not a derivation; no printed object of the framework derives that scale; and the screening direction places the framework scale *below* m_e ($\mu^*/m_e = 0.99989974$), a further unexplained feature. Two exact closed forms now sit in open quarantine — $9/(2\pi^2)$ and $-4/(3\pi)$, both inside the frozen window, both underived: a faithful portrait of where PA2 stands. Finally, the RG channel is now sealed by theorem, not by taste (internal research 102, dated 2026-08-23): the teeth-model energy functional contains no interaction term, pair excitations are never superadditive (5 000 sampled pairs, zero violations), hence self-energy vanishes identically — the "inertial drag of the propagation track" that could dress a bare scale does not exist in the current dynamics, and any future derivation of μ^* requires an interacting extension that the framework does not yet have. **Open quarantine, third entry (dated 2026-08-23, internal research 106) — the solar-angle deficit.** The half-filling ansatz $\sin^2\theta_{12} = (1 - \pi/8)/2$ was attacked as a mode-occupation question (PA10). The mode-counting route is null (16 declared fractions $n/137$, $n/68$; none within 0.5%). In the declared structural family $(1 - \pi/s)/2$ with $s \in \{3, 8, 9, 13, 27, 28\}$, the hit is *unique* at $s = 8$, the number of strong characters; against the pre-registered null (uniform $s \in \{2, \ldots, 137\}$, same functional form, same window) the significance is $p = 1/136 = 0.00735 < 0.01$ — the statistical promotion bar is passed. The derivation is not: the deficit identities $\pi/8 = (\pi/2)/4$ and $\pi/8 = (\pi/4N) \cdot (N/2)$ (accumulated Coleman twist angle) are exact but carry no derived occupation functional, and with the *exact* half-mode $\lambda_{1/2} = 2\sin(\pi/2N)$ the η-linked identity misses $\pi/8$ by 8.6×10^{-6} — the same sine-vs-linearization gap as the PA2 residual. Status: *open quarantine*, statistically significant and underived, like its two siblings; $\sin^2\theta_{12}$ remains [P] with reinforced status. **The gap itself, as a theorem (internal research 107, dated 2026-08-23).** All three open quarantines share the same sine-vs-linearization gap. That gap is now understood mathematically [T]: the native operator of the 137-cycle is the combinatorial graph Laplacian, whose spectrum is *exactly* $4\sin^2(\pi m/N)$ (verified mode by mode, max deviation 5.3×10^{-15}), while the linearized dispersion corresponds to a *non-local* operator — nonzero couplings at 136 of 136 distances — i.e. the "straight ruler" requires an embedding space the framework does not contain. Locality forces the sine. What remains open is the selection step: *why the measured observables are native-spectral rather than linearized* (new problem PA12). The multiplicity asymmetry of the gap ($\times 27$ in PA2, $\times 1$ in the solar sector) matches the anatomy of the host sectors; the lost bet of research 99 (27 independent fluxes) is not reopened. PA12 was then attacked directly (internal research 108, dated 2026-08-23): the generator of the declared local dynamics has exactly the sine spectrum (max deviation 5.3×10^{-15}), so *every* thermally or dynamically generated observable of the framework is native-spectral — locality selects the sine [T]. But the numeric consequences confirm that native selection alone does not close anything (single-mode native $\alpha^{-1} = 137.03601966$ leaves $+2.05 \times 10^{-5}$); closure at 0.149% requires the $\times 27$ multiplicity, and no printed functional of the framework sums over the 27 twist representatives — the teeth-model partition function factorizes over block parities, and no declared rewrite rate references the global flux [T-of-fact]. PA2 thus reduces to a single named missing object, PA12′: *a local functional whose thermal value sums over the* 27 *twist representatives.* That object was then sought directly (internal research 109, dated 2026-08-23) in the closed space of functionals buildable from printed ingredients: twist hopping over the 28 e_R positions (absent — no hopping term is printed, and postulating one would be the framework's first adjustable dial), thermal multi-defect summation ($\langle k\rangle g_1 = -1.95 \times 10^{-7}$, a factor ~ 109 short), the bath occupancy as a thermal "27" (exact 13-block thermodynamics, $\beta\mu_0 \in \{3, \ldots, 5\}$; closest approach

0.641%, outside the 0.149% window), and the thermal spectral sum over the 68 antiperiodic modes ($F/g_1 = 0.026$). All null or absent, printed with equal weight: the space of functionals buildable from printed ingredients is *exhausted*. PA12′ is not underived but *absent* — closing PA2 requires a new ingredient in the framework, not a rearrangement, and any future proposal is immediately checkable: its thermal value must sum to $27\,g_1$. The no-go was then sharpened to a structural statement (internal research 110, dated 2026-08-23): the half-mode's gap distributes *exactly uniformly* over the 137 bonds ($g_1/137$ per address, verified to one part in 10^{15}), so *no sum of strictly local bond operators can produce* $27\,g_1$ — carrying the whole gap per address requires coupling each address to the global mode. Moreover the 28 e_R addresses are syndrome-identical [T], and the printed structure distinguishes at most the anchor-flanking pair $(135, 0)$: the derivable counts are 26 or 28 (both $> 3\%$ off), while the required 27 would need a single excluded address that the framework does not select. PA12′ must therefore be a *collective object of global extent* — named candidate: a Wilson-loop-type operator wrapping the cycle, anchored at an e_R address (the framework's antiperiodicity is itself such an insertion). The inheritor question, PA12″: how many distinct anchors does the loop have in the e_R sector — and why 27. That question has now been answered (internal research 111, dated 2026-08-23): the anchorings are energetically equivalent ($2\mu_0$ each, spectrum cut-independent at 0.00) and exactly two e_R addresses flank the window defect (135 forward, 0 backward; no other address touches the mark in either direction [T]). The exchange $0 \leftrightarrow 135$ is an exact symmetry of every printed observable — the row space of Φ, the kernel, energies, spectra, occupations, rewrite rates (internal research 115, dated 2026-08-23): the orientation is gauge, not physics. Under one declared new postulate [D] — *the defect absorbs exactly one of its two feeding anchors* (either choice gives the same count) — the distinct loop anchors number $28 - 1 = 27$: the fifth derivation of the count, and the first that explains why the sum has 27 identical terms. With it, the framework computes $\alpha^{-1} = N + \eta + 27\,g_1 = 137.03599915$, closing PA2 at 0.149% (residual -3.17×10^{-8} from CODATA). The price is printed without anesthesia: the absorption itself — that one anchor disappears into the defect — is the framework's first new postulate since freezing, with zero directional content (gauge, research 115); and the fine gap $27 \to 26.9598$ (0.0402 copies of g_1) remains unexplained, with the finite-size correction $-9/(160N^2)$ excluded as its source and four further routes printed dead (research 112–113: not thermal, not the frozen-edge flip cost, not the solar account, not a second sine gap in any chirality). The exact frontier is now: *derive the absorption itself*. The first quarantine ($9/(2\pi^2)$) stands promoted to derived-conditional [T, one premise D] by the researcher's decision (dated 2026-08-23). **Dynamical impossibility [T-of-fact].** Any dynamical hosting of c^* is excluded by 8 orders of magnitude inside the framework's own model: observational ceilings on δc^* (atomic clocks: 2.9×10^{-12}; Oklo: 0.19; Cs–Rb: 0.47) sit far below the model's own dynamical floors ($\delta \ln Z = 3.9 \times 10^{-4}$ from the thaw, $\kappa = 2.1 \times 10^{-2}$, syndrome averages varying 350% along the level curve; floor-to-ceiling ratio 1.3×10^{8}). Any future mechanism for c^* *must* be exact and combinatorial — and the audited inventory of the framework's structural integers (28 quantities; 702 declared logarithm pairs, 122 in the window, all background) shows those do not produce it either. The clause "c^* varies with the local attractor/observer condition" is excluded.

K.11 Note 11 — the cycle's spectral zeta at -1 is $2N$ [T]

For the cycle C_{137} with Laplacian eigenvalues $\lambda_k = 2 - 2\cos(2\pi k/137)$, the spectral zeta $\zeta_C(s) = \sum_{k=1}^{136} \lambda_k^{-s}$ (zero mode excluded) has the exact closed values

$$\zeta_C(-1) = \sum_{k=1}^{136} \lambda_k = 274 = 2N, \qquad \zeta_C(1) = \frac{N^2 - 1}{12} = 1564, \tag{70}$$

$$\zeta_C(2) = \frac{(N^2 - 1)(N^2 + 11)}{720} = 489\,532, \tag{71}$$

$$\zeta_C(3) = \frac{(N^2-1)(2N^4+23N^2+191)}{945 \times 64} = 218\,768\,410, \tag{72}$$

all integer- or rational-valued — reinforcing the two-worlds theorem of the body (the $\zeta'(0)$ world is rational and structurally disjoint from the $\zeta(2)$ world that produces η). The identity $\zeta_C(-1) = 2N$ follows by telescoping of the cosine sum and holds for any cycle C_N; it is recorded here as the exact spectral companion of the counting identity $\sum |m| = 137$. None of these invariants meets the PA2 window: the declared family of 287 ζ-forms lands its best member $(N\zeta_C(1)/\zeta_C(2))$ at $\tau = 3.9\%$, deep in background.

References

[1] E. Tiesinga, P. J. Mohr, D. B. Newell, and B. N. Taylor, "CODATA recommended values of the fundamental physical constants: 2022," NIST Reference on Constants, Units, and Uncertainty (2024); *cf.* Rev. Mod. Phys. **93**, 025010 (2021) for the 2018 adjustment.

[2] L. Morel, Z. Yao, P. Cladé, and S. Guellati-Khélifa, "Determination of the fine-structure constant with an accuracy of 81 parts per trillion," Nature **588**, 61–65 (2020).

[3] R. P. Feynman, *QED: The Strange Theory of Light and Matter* (Princeton University Press, 1985), p. 129.

[4] W. Pauli, as quoted in the physics folklore of the fine-structure constant; see H. Kragh, "Magic number: A partial history of the fine-structure constant," Arch. Hist. Exact Sci. **57**, 395–431 (2003).

[5] Planck Collaboration, "Planck 2018 results. VI. Cosmological parameters," Astron. Astrophys. **641**, A6 (2020).

[6] M. Raidal, "Relation between the neutrino and quark mixing angles and grand unification," Phys. Rev. Lett. **93**, 161801 (2004).

[7] A. Yu. Smirnov, "Neutrino oscillations: on the way to the standard paradigm," arXiv:hep-ph/0402264.

[8] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz, "NuFIT 6.0 (2024)," JHEP **09**, 178 (2024); `www.nu-fit.org`.

[9] T2K and NOvA Collaborations, "Joint neutrino oscillation analysis from the T2K and NOvA experiments," Nature **646**, 818–824 (2025).

[10] JUNO Collaboration, "Neutrino physics with JUNO," J. Phys. G **43**, 030401 (2016).

[11] DUNE Collaboration, "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics," arXiv:2002.03005.

[12] Hyper-Kamiokande Collaboration, "Hyper-Kamiokande Design Report," arXiv:1805.04163.

[13] LiteBIRD Collaboration, "Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey," PTEP **2023**, 042F01 (2023).

[14] CMB-S4 Collaboration, "CMB-S4 Science Book, First Edition," arXiv:1610.02743.

[15] S. Navas *et al.* (Particle Data Group), "Review of particle physics," Phys. Rev. D **110**, 030001 (2024).

[16] JUNO Collaboration, first oscillation results, *Nature* **654** (2026), cover of 11 June 2026, DOI: 10.1038/s41586-026-10538-z; independent analysis of the first JUNO data: arXiv:2601.09791.

[17] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, and T. Schwetz, "Nu-FIT 6.1 (2026)," `www.nu-fit.org` (dated 2026 note only; the frozen edition remains Nu-FIT 6.0 [8]).

[18] R. H. Parker, C. Yu, W. Zhong, B. Estey, and H. Müller, "Measurement of the fine-structure constant as a test of the Standard Model," Science **360**, 191–195 (2018).

[19] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse, "Measurement of the electron magnetic moment," Phys. Rev. Lett. **130**, 071801 (2023).

[20] Yu. V. Petrov, A. I. Nazarov, M. S. Onegin, V. Yu. Petrov, and E. G. Sakhnovsky, "Natural nuclear reactor at Oklo and variation of fundamental constants: computation of neutronics of a fresh core," Phys. Rev. C **74**, 064610 (2006).

[21] T. Damour and F. Dyson, "The Oklo bound on the time variation of the fine-structure constant revisited," Nucl. Phys. B **480**, 37–54 (1996).

[22] R. Lange *et al.*, "Improved limits for violations of local position invariance from atomic clock comparisons," Phys. Rev. Lett. **126**, 011102 (2021).

[23] BICEP/Keck and Planck Collaborations, combined BK18+PR4 analysis, $r < 0.032$ at 95% CL, the strongest bound in force as of March 2026 (dated note).

[24] Logarithmic-oscillation search in Planck and SPT-3G CMB data, arXiv:2403.19575 (amplitude bound $\lesssim 0.023$ at 95% CL).