UCLA/Arisaka-GG-C4-Leptogenesis
v23.3 — September 11, 2026 — *Preprint*

The Origin of Matter

Holonomic *CP* Violation, Majorana Neutrinos, and Leptogenesis at the Lock Scale

Katsushi Arisaka

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

Contact: arisaka@physics.ucla.edu

Abstract

Every atom in every reader of this page is a leftover. For each billion quark–antiquark pairs the early universe made and unmade, one quark survived; that residue is all the matter there is. This study asks where the residue came from, and answers with a chain that has no adjustable number in it.

The answer is geometric. In GG-Theory the sixth dimension is a compact phase circle and the fifth an open scale interval. Loops close only on the circle, so every physical phase — the whole content of *CP* violation — is a record of going around it, and the circle admits only a five-fold set of phases. *CP* violation exists because the flavor alphabet built on that set is chiral, and no choice was available to make it otherwise. The same argument says why the strong and electromagnetic forces do not violate *CP*: their circuit closes with nothing left over.

The quark sector cannot pay for us — it falls ten to twelve orders of magnitude short — so what remains is the neutrino route, where the two Majorana phases turn out to be fixed parities rather than dials. The heavy neutrinos sit at 3.85 TeV, set by a ledger integer chosen for reasons that have nothing to do with baryogenesis.

The chain lands at $\eta_B = 7.34 \times 10^{-10}$ against a measured $(6.12 \pm 0.04) \times 10^{-10}$: a +20% central deviation with nothing fitted, at a declared band of $\times 3^{\pm 1}$. The full flavored Boltzmann integration, which assumes least, lands at $2.8^{+2.0}_{-1.2} \times 10^{-10}$ — consistent at 1.3σ, with a factor-2.2 deficit charged to standard transport effects. That deficit is the number to watch. The same rule that permits this history forbids proton decay outright, and both halves are killable: by DUNE and Hyper-Kamiokande on the leptonic phase, by tonne-scale double beta decay, and by a proton decay that must never be seen.

Executive Summary

The question. The universe we live in is made of matter and not of antimatter, and the imbalance is small: about one surviving baryon per 1.6 billion photons. Sakharov showed in 1967 that producing such an excess needs three things at once — baryon-number violation, C and CP violation, and departure from thermal equilibrium. The Standard Model has all three in principle and none of them in sufficient measure: its electroweak transition is a crossover rather than a first-order phase transition, and its single CP phase delivers a dimensionless asymmetry measure of order 10^{-20} where 10^{-10} is needed. This study asks what GG-Theory says instead, and takes the answer all the way to a number.

Where a phase can live. The six-dimensional bulk is $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$, with $I_r = [0, K]$ the open scale interval — the renormalization-group coordinate $r = \ln(E_{\text{GUT}}/\mu)$ made geometry — and S^1_θ the compact phase circle. That asymmetry decides everything that follows. A holonomy is a property of a *closed* loop, and an interval has none: the dr component of the connection can always be gauged to zero, and there is no closed circuit in r on which to evaluate anything. Every physical phase in the theory is therefore a circle-side object. *The circle remembers; the interval forgets.*

Why *CP* violation is not optional. Regard a Yukawa texture as a U(1)-valued connection on the flavor graph. Rephasing the fields is exactly per-family re-zeroing of the dials; in a spanning-tree gauge every tree edge is set to zero and the residual physical data are exactly the loop holonomies. On the GG-6 bulk — one circle, block integer $L = 5$, and decisively no second flavor torus on which an angle could float — those holonomies are valued in the pentagonal lattice $(\pi/5)\mathbb{Z}$. The Higgs doublet retains the fraction $2/L$ of the unified block, which forces the branch phase $\pi(1 - 2/L) = 3\pi/5$; and CP is conserved if and only if every loop angle lies in $\{0, \pi\}$. Since $3\pi/5$ does not, CP violation must exist. One unphysical lattice step leaves the physical phase $2\pi/5 = 72^\circ$, completed by the quantum-defect step $-\pi/K$ to $13\pi/35 = 66.9^\circ$.

The same picture says why the strong and electromagnetic sectors are different. Color writes a *triangle* on the circle — three steps of $2\pi/3$ — and the triangle's circuit is trivial: the loop closes with nothing left over, the holonomy is 1, and CP is exact. *Closure without residue confines; closure with residue violates.*

A picture we tested and had to withdraw. One might hope to see the weak interaction's phase as *swept out* along a helix: the scale coordinate r increasing down a decay cascade while the gauge phase turns. Part IV builds that helix quantitatively for both sectors and then tests it. If the phase accumulated at a fixed rate along r, the two sectors would share a pitch. They do not: the quark cascade turns 8.8° per unit r and the lepton cascade 24.1°, a factor of 2.73 apart over nearly the same interval. A universal pitch is excluded. The helix is where the phase *lives*, not what *generates* it — which is precisely the framework's own reading, now supported by a test that could have failed the other way.

Why the quark sector cannot pay for us. For $N = 3$, unitarity ties every plaquette of the flavor graph to a single Jarlskog invariant — one area for the whole graph, twice the area of the unitarity triangle. The lepton area is $|J_{\rm PMNS}| \simeq 9.4 \times 10^{-3}$ against the quark area $J_{\rm CKM} \simeq 3.1 \times 10^{-5}$: a factor of about three hundred. And even that understates it, because what enters baryogenesis is the area *times* the product of quark mass-squared splittings over the twelfth power of the temperature. At $T = 100$ GeV that measure is 1.2×10^{-19}; the T^{-12} makes it steeply temperature dependent, and across the plausible range the shortfall against 6.1×10^{-10} is ten to twelve orders of magnitude. The quark sector is a bystander. That is the sentence that forces the lepton route.

The Majorana door. A fermion multiplet carrying weight n under the phase circle admits a Lorentz- and gauge-invariant Majorana self-pairing — a $\Delta L = 2$ bilinear — if and only if $n = 0$; on every $n \neq 0$ sector the only invariant symmetric self-pairing is the charge-conserving Dirac pairing. Lepton number is winding number, exact on every dial-carrying species. So the neutral sector is not a model preference but the unique door, and the trilinear LHN is the unique fiber-neutral portal. The light masses follow from the seam coefficient without a new scale, $m_\star = \pi v^2/E_{\rm GUT} \simeq 8.35$ meV, giving $(m_1, m_2, m_3) \simeq (0.065,\ 8.59,\ 50.5)$ meV in *normal ordering as an output*, $\Sigma m_\nu \simeq 59$ meV. The two Majorana phases are fixed parity data of the rank-two block, $(\alpha_{21}, \alpha_{31}) = (0, \pi)$, and $\alpha_{31} = \pi$ is the letter that carries the excess. It does so in a way worth stating precisely, because the obvious reading of it is wrong: a Majorana parity of π is CP-conserving in the flavor-summed asymmetry — that sum vanishes here exactly — and CP-violating in the flavor *differences*, which do not. Differential washout turns a set of asymmetries that adds to nothing into a net excess.

Leptogenesis at four TeV. The heavy Majorana states sit at $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-K} \simeq 3.85$ TeV, with splittings of order $1/K$. That is not a scale chosen for accessibility: it is the scale the ledger integer $K = 35$ fixed for its own reasons, the same K that sets the Kaxion's frequency window and the vacuum's e^{-2K}. Standard thermal leptogenesis with hierarchical heavy states is capped by the Davidson–Ibarra bound at $M_N \gtrsim 10^9$ GeV, forever beyond every collider. The known escape is the resonant regime; the literature buys it by *engineering* a degeneracy. Here the degeneracy is not engineered but *faced in public*: the on-resonance condition is the named hypothesis H-$\phi_{\rm geo}$-1, and the honest bracket is stated on both sides.

The number. The chain's throughput reconstructs the ledger's exponential, $\epsilon^{\rm res} \kappa \simeq (2.5 \times 10^{-6})(1.0 \times 10^{-2}) \simeq 2.5 \times 10^{-8} = e^{-K/2}$, and standard sphaleron transfer and entropy dilution finish the job: $\eta_B = 7.04\,(28/79)\,(1.25/g_*)\,e^{-K/2} = 7.34 \times 10^{-10}$ against a measured $(6.12 \pm 0.04) \times 10^{-10}$ — the triplet's equilibrium yield $1.25/g_*$ now displayed rather than absorbed. The excess is one part in a billion because half a lock depth *is* $10^{-7.6}$, and the universe paid the toll exactly once. What must not be read from the $+20\%$ central (headline $\approx 7 \times 10^{-10}$) is precision: the ledger-level band is a factor of three either way. This study therefore goes further than its parents: it performs the full flavored Boltzmann integration those papers declared as awaited, with every input pinned to a derived GG quantity — the flavored washouts $K_\alpha = (3.5,\ 29.3,\ 22.5)$ computed from the derived masses and mixing matrix, the kernels validated against the standard analytic efficiency to four per cent — and scans only the freedoms the cohort itself declares. The error bar contracts by a factor of three and a half, and the sign of the excess emerges as an *output*: the flavor-summed asymmetry is exactly zero, and matter wins because the electron flavor, decoupled from the heavy third mass state by the derived $\sin^2\theta_{13} = 3/140$, is erased ten times more slowly than the others. Ninety-six per cent of it leaves through the electron door. The best-estimate prediction is $\eta_B = 2.8^{+2.0}_{-1.2} \times 10^{-10}$ against the observed 6.12×10^{-10} — consistent at 1.3σ on the log scale, with the central factor-2.2 deficit assigned to named, standard, one-sided transport physics, and the Boltzmann-level deficit now *measured*, $E = 7.1$ [4.4, 11.4], where the parents could only bracket 2–20. All of it conditional on C4, which no Boltzmann-level work can discharge.

Absolutely stable protons, absolutely stable electrons. The same exact baryon number that makes the ledger work forbids proton decay. Baryon number can change only through the $\mathrm{SU}(2)_L$ anomaly, in units of $N_{\mathrm{gen}} = 3$; a proton decay is $\Delta B = -1$ and $-1 \notin 3\mathbb{Z}$, so no anomalous process, however violent, can remove a single nucleon. Neutron–antineutron oscillation ($\Delta B = 2$) is forbidden by the same rule. The electron is stable for a different reason: it is the ground state of its stratum and possesses no exit gate — stability as an *absence of destinations.* Where grand unification predicts proton decay and is confirmed by finding it, this framework predicts $\tau_p = \infty$ and is refuted by a single event. That a universe that made an excess of baryons and a proton that can never decay are compatible is not a loophole but the selection rule doing both jobs at once.

How to kill it. Six dated tests. A measured δ_{PMNS} inconsistent with 196.4° at DUNE and Hyper-Kamiokande retires the leptonic lattice. A confirmed $0\nu\beta\beta$ line anywhere in the tonne-scale reach of 9–20 meV is as fatal as an absence at 1.8 meV would be unprovable. An established inverted ordering retires the paper with no fallback spectrum. A cosmological Σm_ν robustly below ~ 58 meV under this framework's own $w = -1$ prior falsifies the inherited block — and this is the tightest live tension of all six, with current data already at the boundary. Absence of the lock cluster near 3.85 TeV at HL-LHC retires the chain regardless of every leptonic success. And one proton decay, in any channel, ends the whole baryonic logic.

Closing reflection. Two things here are, so far as we can establish, without precedent. Leptogenesis by Majorana neutrinos at a scale of a few TeV — reachable by a machine that exists — is not the standard picture, which places the mechanism nine orders of magnitude higher. And a baryon-to-photon ratio obtained from a single integer, rather than accommodated by adjustable heavy masses and phases, is a different kind of claim from the one the leptogenesis literature has been able to make. Both may be wrong. Neither can be quietly rescued: the first is exposed at HL-LHC, the second is exposed at the joint where its one postulated exponent sits. We would rather be refuted in public than be safe.

Keywords: GG-Theory; baryogenesis; resonant leptogenesis; Majorana neutrinos; *CP* violation; flavor holonomy; pentagonal lattice $(\pi/5)\mathbb{Z}$; Jarlskog invariant; $\delta_{\mathrm{CKM}} = 13\pi/35$; $\delta_{\mathrm{PMNS}} = \pi + 2\arcsin(1/7)$; Majorana phases $(0, \pi)$; the dial-free door; geometric seesaw; lock scale $M_{\mathrm{lock}} = 3.85$ TeV; Davidson–Ibarra evasion; sphaleron transfer 28/79; baryon-to-photon ratio η_B; $\Delta B \in 3\mathbb{Z}$; absolute proton stability; neutrinoless double beta decay; DUNE; Hyper-Kamiokande; falsifiability.

It Could Have Come Out at 10^{-6} or 10^{-14} — and There Would Have Been Nothing to Do About It

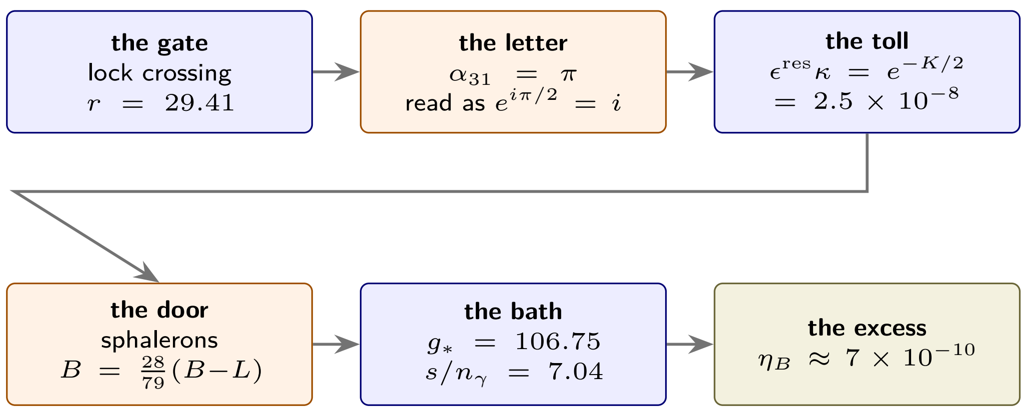

Blue is the interval's contribution — scales, amounts, the departure from equilibrium. Orange is the circle's — orientations, the letter, the direction of the bias. Green is what is left in the sky. The asymmetry is the product of exactly six factors, and every one of them is either a ledger integer, a theorem of the framework, or textbook equilibrium bookkeeping. Nothing was fitted to 6.12×10^{-10}.

Where the matter came from, in six boxes. The universe keeps about one baryon for every 1.6 billion photons. Sakharov's three conditions say what producing such an excess requires — baryon-number violation, *CP* violation, departure from equilibrium — and the Standard Model has all three in principle and none in sufficient measure: its single *CP* phase delivers a dimensionless asymmetry of order 10^{-20} where 10^{-10} is needed. This chain supplies the missing ten orders from structure fixed elsewhere. The color key is inside the picture, and the point of the page is what the picture does not say twice: *there is no dial in it*: K is 35 because $LJ = 35$ and $M^2 + N^2 + L^2 = 35$, g_* is the Standard Model's, 28/79 is Harvey–Turner, 7.04 is thermodynamics, and nothing was fitted to the measured $(6.12 \pm 0.04) \times 10^{-10}$. **And the honesty is the point of the page.** The ledger level lands at $\eta_B \approx 7 \times 10^{-10}$; the full flavored Boltzmann integration, which assumes least, lands at $2.8^{+2.0}_{-1.2} \times 10^{-10}$ — consistent at 1.3σ on the logarithmic scale, with a central factor-2.2 deficit assigned to named, standard, one-sided transport effects. That deficit is now measured rather than bracketed, and it is the number to watch: if the completed transport cannot supply it, the chain dies by throughput. In full: Figure 16, p. 83, with the measurement in Figure 15, p. 81.

Overview of Contents

Part I. The Question, and the Shape of the Answer 14
Part II. The Arena: One Fiber that Remembers, One that Forgets 19
Part III. The Holonomic Origin of *CP* Violation 27
Part IV. The Helix: A Picture Built, Tested, and Corrected 35
Part V. The Quark Sector: One Area for the Whole Graph 40
Part VI. Why the Quark Sector Cannot Pay for Us 44
Part VII. The Majorana Door: The Only Way In 48
Part VIII. Leptogenesis at Four TeV 55
Part IX. From Lepton Number to Baryon Number: The Predicted η_B 72
Part X. Absolutely Stable Protons, Absolutely Stable Electrons 84
Part XI. What Is New Here, and What Is Not 89
Part XII. The Falsification Program, 2026–2040 95
Part XIII. Discussion 103
Part XIV. Conclusions 106
Part XV. Two Talks: Why Are We Here? 109
Part XVI. For the Curious Reader 125
Part XVII. Technical Appendices 129
Part XVIII. References 134

List of Figures

1 The arena: the compact fiber and the open fiber. 20
2 Holonomy, in three rungs. 22
3 The chiral alphabet on the circle. 29
4 The pentagon's two angles. 30
5 Two figures on one circle, two fates. 32
6 The three orientation reversals. 33
7 The quark cascade as a helix. 36
8 The lepton cascade as a helix. 36
9 The pitch test, and what its failure buys. 37
10 The two unitarity triangles at true relative area. 43
11 Why the quark sector cannot pay. 46
12 Where on the ladder the excess is paid. 57
13 The regulator landscape, drawn. 64
14 The birth of the asymmetry, integrated. 70
15 The predicted η_B: three levels and the measurement. 81
16 The birth of the excess, end to end. 83
17 The $\Delta B \in 3\mathbb{Z}$ selection rule. 86
18 The decision calendar, 2026–2040. 101

List of Tables

1 Provenance of the load-bearing elements. 17
2 The ledger and where each integer acts in this study. 26
3 The five operations, located. 32
4 Why each symmetry has the fate it has. 34
5 The pitch test: the two sectors' measured pitches. 37
6 The quark sector: predicted against measured. 42
7 The Standard Model's *CP* measure, and its temperature sensitivity. 45
8 The lepton sector: what is fixed, and by which integer. 54
9 The nine-step audit of the chain. 59
10 The regulator landscape at the lock crossing. 63
11 The six conditions of the candidate regulator. 65
12 The flavored washouts, computed from derived inputs. 67
13 Validation of the integration. 69
14 The chain arithmetic, factor by factor. 78
15 What could move, and what could not. 79
16 The uncertainty budget, before and after the integration. 80
17 Two absolute stabilities, two different mechanisms. 88
18 The honest ledger of the whole chain. 91
19 Cosmological bounds on the neutrino mass sum, August 2026. 98
20 The kill table. 102
21 What each letter is worth. 115
22 Dictionary of symbols. 130
23 Numerical audit of every quantity in this study. 132

Contents

Executive Summary . . . 2
List of Figures . . . 6
List of Tables . . . 7
Preface . . . 12
Notation . . . 12

Part I. The Question, and the Shape of the Answer 14
1 One survivor in a billion . . . 14
1.1 Sakharov's three conditions, and what the Standard Model does with them 14
2 What a geometric theory is obliged to say . . . 15
2.1 The chain in one paragraph . . . 15
2.2 What is inherited, and from where . . . 16
2.3 What GG-P5-CPT and GG-W2-Holonomy added after GG-A9-SM and GG-A10-Cosmos . . . 16
3 How to read this study . . . 18

Part II. The Arena: One Fiber that Remembers, One that Forgets 19
4 The bulk, and the asymmetry inside it . . . 19
4.1 Physical picture: a grooved record and a one-way stylus . . . 19
5 A one-page primer: holonomy, from the bench to the bulk . . . 20
5.1 Rung one: the Aharonov–Bohm phase . . . 21
5.2 Rung two: the Wilson loop . . . 21
5.3 Rung three: two algebras, and only one of them a group . . . 21
6 The decisive computation: the connection, fiber by fiber . . . 22
7 Holonomy on the bench: three worked examples . . . 23
7.1 A vector walked around a cone of latitude . . . 23
7.2 The Aharonov–Bohm phase, with laboratory numbers . . . 23
7.3 The two transport algebras, with proofs the size of a postcard . . . 24
8 The Kaluza–Klein tower, done honestly . . . 24
9 From the bench to the bulk: what is the same, what is new . . . 25
10 The ledger, and why five integers . . . 25
11 What this Part has bought . . . 26

Part III. The Holonomic Origin of *CP* Violation 27
12 Flavor as a connection on a graph . . . 27
13 The alphabet: why the letters are multiples of $\pi/5$ . . . 28
13.1 The sentence . . . 29
14 The pentagon, in trigonometry . . . 30
15 The triangle and the pentagon: why the strong sector does not violate *CP* . . . 31
15.1 Strong *CP*, dissolved in passing . . . 31
16 *C*, *P*, *T* located — and why *CPT* is exact . . . 31
17 The fates, in one table . . . 33

Part IV. The Helix: A Picture Built, Tested, and Corrected 35
18 The hypothesis worth testing . . . 35
19 Building it quantitatively . . . 35
20 The pitch test — and it fails . . . 37

21 The retraction, formalized: three survivors . . . 37
22 What the failure teaches . . . 39

Part V. The Quark Sector: One Area for the Whole Graph 40
23 The Cabibbo angle, and why the golden ratio is in it . . . 40
24 One area, and what it means that there is only one . . . 41
25 The two areas, drawn to scale . . . 42
26 A remark on where this leaves the quark sector . . . 42

Part VI. Why the Quark Sector Cannot Pay for Us 44
27 Sakharov's conditions, located . . . 44
28 The Standard Model's two failures . . . 44
28.1 The crossover . . . 44
28.2 The measure, and its temperature . . . 45
29 Three doors, and which ones are welded . . . 46

Part VII. The Majorana Door: The Only Way In 48
30 Lepton number is winding number . . . 48
31 The geometric seesaw: a scale that is already in the ledger . . . 49
31.1 The light spectrum, and normal ordering as an output . . . 50
32 δ_{PMNS}: the leptonic phase and the integer inside it . . . 50
33 The Majorana pair: parity data, not dials . . . 51
33.1 The apparent paradox, and where its resolution actually lies . . . 52
34 The shadow the door casts: $m_{\beta\beta}$ . . . 53

Part VIII. Leptogenesis at Four TeV 55
35 The Davidson–Ibarra wall, and why everyone else is behind it . . . 55
36 Where the gate is . . . 56
37 The chain, step by step . . . 56
38 The resonance question, faced: what H-ϕ_{geo}-1 is, and why it is hard . . . 59
38.1 The question, in one sentence . . . 60
38.2 Why the obvious answer fails, and why the peak fails too . . . 60
38.3 The hypothesis, in its sharp form . . . 61
38.4 The candidate mechanism: the bath as a flavor-measuring apparatus . . . 61
38.5 The closed form, and why a ledger integer appears in it . . . 62
38.6 The regulator landscape: why this is falsifiable rather than decorative . . . 63
38.7 Why it is hard — five reasons, none of them evasions . . . 63
38.8 The six conditions . . . 64
38.9 Three routes by which the chain kills itself . . . 65
39 The washout, flavor by flavor: the electron door . . . 66
39.1 Why each flavor has its own washout . . . 66
39.2 The numbers, and the asymmetry between them . . . 67
40 The Boltzmann system, integrated . . . 67
40.1 The equations, term by term . . . 68
40.2 Validation before results . . . 68
40.3 What the integration shows . . . 69
41 Zero dilution . . . 71

Part IX. From Lepton Number to Baryon Number: The Predicted η_B 72
42 The door that converts . . . 72
43 The bookkeeping: what s is, and why every yield is measured against it . . . 72
43.1 The problem: everything is diluting at once . . . 73
43.2 The answer: entropy density . . . 73
43.3 Yields . . . 73
44 The six factors, one at a time . . . 74
44.1 Factor 1 — how many heavy neutrinos there were to decay: $Y_N^{\rm eq} \simeq 1/g_*$ . . . 74
44.2 Factor 2 — how biased each decay was: $\epsilon^{\rm res}$ . . . 74
44.3 Factor 3 — how much of it survived: κ . . . 75
44.4 Factor 4 — how much of the lepton excess became a baryon excess: $c_{\rm sph} = 28/79$ . . . 76
44.5 Factor 5 — how many species were sharing the bath: $g_* = 106.75$ . . . 76
44.6 Factor 6 — converting from entropy to photons: 7.04 . . . 77
45 Assembling the number . . . 77
45.1 Which factors could move, and by how much . . . 78
46 The error bar, derived . . . 79
46.1 The budget, before and after . . . 79
46.2 Three levels, three statements . . . 79
46.3 How to quote it . . . 80
46.4 What would tighten it further . . . 81
47 Why one part in a billion? . . . 82
48 The one seam we do not close . . . 82

Part X. Absolutely Stable Protons, Absolutely Stable Electrons 84
49 The apparent contradiction . . . 84
50 The theorem . . . 84
50.1 The subtle step, and why it is not automatic . . . 85
51 A prediction that points the wrong way on purpose . . . 86
52 The electron: stability as an absence of destinations . . . 86
52.1 The symmetry of the two arguments . . . 87

Part XI. What Is New Here, and What Is Not 89
53 Two claims that appear to be without precedent . . . 89
53.1 First: Majorana leptogenesis at a scale of a few TeV, with the scale not chosen . . . 89
53.2 Second: a quantitative baryon-to-photon ratio from an integer . . . 90
53.3 Third, and less often noticed: the compatibility . . . 90
54 The two joints where it can fail . . . 90
55 What is inherited, and from where — the cohort's own history . . . 91
55.1 GG-A9-SM: the flavor connection and the door . . . 92
55.2 GG-A10-Cosmos: the transport, the number, *and* the mechanism . . . 92
55.3 GG-P5-CPT: the hypothesis, the grading, the roadmap . . . 92
55.4 GG-W2-Holonomy: the anatomy . . . 93
56 What this study adds . . . 93

Part XII. The Falsification Program, 2026–2040 95
57 How this Part is organized . . . 95
58 $\delta_{\rm PMNS}$: DUNE and Hyper-Kamiokande . . . 95
59 The mass ordering . . . 96

60 Neutrinoless double beta decay: the Majorana door itself . . . 96
61 The neutrino mass sum: the tightest live tension . . . 97
62 The lock cluster at HL-LHC . . . 99
63 Proton decay and n–$\bar{n}$ oscillation . . . 100
64 Electric dipole moments, and a direct sphaleron handle . . . 100
65 The decision calendar . . . 100
66 A remark on the shape of this program . . . 100

Part XIII. Discussion 103
67 Sixty years of the question . . . 103
68 Where this chain differs . . . 103
69 The intuitive core, restated once . . . 104
70 Open items, stated as such . . . 105
71 What would make this work look different in ten years . . . 105

Part XIV. Conclusions 106
72 What has been established . . . 106
73 Closing . . . 108
Acknowledgments . . . 108

Part XV. Two Talks: Why Are We Here? 109
74 Division 1: A colloquium — the geometric origin of matter . . . 109
75 Division 2: A lecture for everyone — why is there anything? . . . 118

Part XVI. For the Curious Reader 125
76 The questions . . . 125

Part XVII. Technical Appendices 129
77 Appendix A. Dictionary of symbols and names . . . 129
78 Appendix B. Worked problems . . . 129
79 Appendix C. Numerical audit . . . 131
80 Appendix D. The three script files . . . 133

Part XVIII. References 134

Preface

Every physicist has one question underneath all the others. Mine has been the same since I first walked into the Kamiokande hall: where did the matter that makes us come from?

I have been fortunate to spend a career at the four experimental corners of exactly that question. At Kamiokande we watched tanks of water for the decay of the proton — the question of whether matter, once made, is permanent. At KTeV we measured direct *CP* violation in the kaon system — the first of the two laboratory windows on the asymmetry between matter and antimatter. At BaBar we watched the B mesons violate *CP* and, eventually, time reversal itself — the second window, and the confirmation that the Standard Model's single phase is real. And at the LHC we searched the TeV scale — the energy range where, if this study is right, the matter of the universe was actually decided, at 10^{-13} seconds, by heavy neutral particles a present-day machine can almost reach.

Four experiments, four pillars of this paper: proton stability (Part X), quark-sector *CP* violation and its beautiful insufficiency (Parts V–VI), the discrete symmetries and their geometric fates (Part III), and the TeV gate itself (Part VIII). I did not plan the career as a syllabus for this question, but it turned out to be one.

This study is the place where I try to put the whole answer together, under the discipline this program has always demanded: every number computed, every grade carried honestly, every way to be wrong written down with a date on it. If the answer survives the instruments now under construction, we will know where we came from. If it does not, we will know that too, and the question will still be the best one I know.

Notation

A handful of names appear in the Abstract and Executive Summary before the text can define them. They are collected here, one line each. The full dictionary is Appendix 77.

- **The bulk.** $\mathcal{M}_6 = \mathcal{M}_4 \times I_r \times S^1_\theta$. $I_r = [0, K]$ is the *scale interval*, the fifth dimension, with coordinate $r = \ln(E_{\rm GUT}/\mu)$. $S^1_\theta \cong S^1$ is the *phase circle*, the sixth dimension, with coordinate θ.
- **The ledger.** $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, with $J = 2N+1$ and $K = LJ = M^2+N^2+L^2$.[1] $N = 3$ counts families, $L = 5$ the block rank and the angle alphabet, $J = 7$ the chiral steps, $K = 35$ the depth of the descent.
- **Scales.** $M_{\rm Pl} = 1.22 \times 10^{19}$ GeV; $E_{\rm GUT} = M_{\rm Pl} e^{-2\pi} \simeq 2.28 \times 10^{16}$ GeV; $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-K} \simeq 3.85$ TeV, the *lock scale*; $v = 246$ GeV, with $\bar{v} = v/\sqrt{2} = 174$ GeV where the leptogenesis literature's convention requires it.

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

- **Phases.** $\delta_{\rm CKM}$ and $\delta_{\rm PMNS}$ are the Dirac *CP* phases of the quark and lepton mixing matrices; $(\alpha_{21}, \alpha_{31})$ are the two Majorana phases; $J_{\rm CKM}$, $J_{\rm PMNS}$ are the corresponding Jarlskog invariants.
- **The chain.** $\epsilon^{\rm res}$ is the resonant decay asymmetry per heavy-neutrino decay; κ the washout efficiency; $K_{\rm wash}$ the washout strength; $c_{\rm sph} = 28/79$ the sphaleron transfer coefficient; $g_* = 106.75$ the relativistic degrees of freedom at the TeV bath; $\eta_B = (n_B - n_{\bar{B}})/n_\gamma$ the baryon-to-photon ratio.
- **Grades** (the cohort's honesty register, used unchanged here). *Derived*: forced by a named theorem with every factor traceable. *Derived-conditional*: a genuine theorem supports the structure, but one named lemma or hypothesis is added and is named in place. *Named input*: a quantity the chain requires and does not derive. *Asserted / flagged*: a closed form whose role assignment is a working hypothesis. *Postulate*: stated as such. Nothing in this study upgrades any grade it inherits.
- **Typography.** Cohort papers are named in sans serif (GG-A9-SM, GG-A10-Cosmos, GG-P5-CPT, GG-W2-Holonomy). A boxed equation is a result carried forward; a shaded box is a caution or an open item.

Part I
The Question, and the Shape of the Answer

Why there is anything at all to be made of, and what a geometric theory is obliged to say about it

1 One survivor in a billion

Take the number seriously before taking anything else seriously. The cosmic microwave background fixes the ratio of baryons to photons at

$$\eta_B \equiv \frac{n_B - n_{\bar{B}}}{n_\gamma} = (6.12 \pm 0.04) \times 10^{-10} \tag{1}$$

(Planck Collaboration, 2018). Read that backwards. In the hot early universe, quarks and antiquarks were made and unmade in vast numbers at equilibrium; when the bath cooled below the point where pair creation could keep up, essentially all of them annihilated into photons. What survived the annihilation is what we are made of. For every 1.6 billion photons in the sky today there is one baryon that had no partner. Every carbon atom in the reader's hand is a leftover of that accounting error.

There are two ways to be unsatisfied with 6.1×10^{-10}. The first is that it is not zero: a universe with equal parts matter and antimatter would annihilate to a photon bath and no observer. The second is that it is not one: something suppressed the surplus by nine orders of magnitude, and a number nine orders down from unity is the signature of a mechanism, not of an initial condition. A theory that claims to fix the parameters of the world owes an account of both.

1.1 Sakharov's three conditions, and what the Standard Model does with them

Sakharov (1967) identified the three necessary ingredients: baryon number must be violated; C and CP must be violated; and the relevant processes must occur out of thermal equilibrium. The Standard Model possesses all three in principle. It possesses none of them in sufficient measure.

Baryon-number violation exists in the Standard Model as the electroweak anomaly, and Kuzmin et al. (1985) showed that at high temperature the associated sphaleron transitions are fast, not exponentially suppressed. That much works.

Departure from equilibrium does not. For a Higgs mass of 125 GeV the electroweak transition is a smooth crossover (Kajantie et al., 1996), not the first-order phase transition with expanding bubble walls that electroweak baryogenesis requires. There is no departure to speak of.

CP *violation* fails by the widest margin of the three. The Standard Model's single CP-violating phase enters observables only through the Jarlskog invariant, and the resulting dimensionless measure at the electroweak scale is of order 10^{-20} (Gavela et al., 1994, Huet and Sather, 1995) against the 10^{-10} required. Part VI does this arithmetic in full, including a feature of it that is usually glossed: the measure carries T^{-12}, so the shortfall is not a single number but a range, ten to twelve orders of magnitude depending on where in the crossover one evaluates it. Ten orders is already fatal.

So the Standard Model predicts, to the accuracy of its own machinery, that η_B is essentially zero, and the sky says otherwise. That is not a small discrepancy at the edge of a precision measurement. It is the largest quantitative failure in the Standard Model's ledger, and it is about us.

2 What a geometric theory is obliged to say

GG-Theory claims that the parameters of the Standard Model and of cosmology are not inputs but readings of a six-dimensional geometry governed by five small integers. A claim of that kind cannot decline the question of the matter excess. If CP violation is a geometric angle, then the angle is fixed and the excess it drives is a *prediction*, not a parameter. If the mechanism needs a heavy mass scale, that scale must already be in the ledger or the claim is void. And if the mechanism needs baryon-number violation, the theory must explain in the same breath why the proton does not decay — because the same exact symmetry is doing both jobs.

This study assembles that account. It derives nothing for the first time. Every equation, number and grade is carried from a named cohort paper at that paper's own grade, and where a paper flags a step, the flag is carried verbatim rather than absorbed. The results assembled here are owned by GG-A9-SM, GG-A10-Cosmos, GG-P3-Electron, GG-P5-CPT and GG-W2-Holonomy; the structure they stand on — the integer ledger, its anomaly closure and the descent machinery — is owned by GG-A2-GG6, GG-A3-BI6, GG-A6-GRIP and GG-P2-N3. All nine are listed in the references, each with its own page and its permanent identifier. What is new is the assembly — the argument end to end in one place, with the intuitive layer written out at the length the physics deserves — together with one construction original to this study, whose principal result is negative and is reported as such.

2.1 The chain in one paragraph

The sixth dimension is a compact circle and the fifth is an open interval; loops exist only on the circle, so every physical phase is a holonomy of the circle. The circle's block integer $L = 5$ admits only the pentagonal lattice $(\pi/5)\mathbb{Z}$, and the Higgs doublet's retention of $2/L$ forces the branch phase $3\pi/5$ — which is not a multiple of π, so CP violation is compulsory. The quark sector receives $\delta_{\rm CKM} = 2\pi/5$, completed to $13\pi/35$; the lepton sector receives $\delta_{\rm PMNS} = \pi + 2\arcsin(1/7)$ and, from the rank-two Majorana block, the parity pair $(\alpha_{21}, \alpha_{31}) = (0, \pi)$. The quark sector's Jarlskog area, throttled by mass splittings, falls ten orders short, so the excess must be written in the lepton sector. The compact fiber admits a Majorana pairing only at weight zero, which makes the neutral sector the unique door; the heavy Majorana states sit at the lock scale $M_{\rm lock} = \frac{1}{2}M_{\rm Pl}e^{-K} \simeq 3.85$ TeV that the ledger had already fixed; quasi-degeneracy of order $1/K$ puts the mechanism in the resonant regime; the throughput reconstructs $e^{-K/2}$; and sphaleron transfer with entropy dilution lands $\eta_B \approx 7 \times 10^{-10}$ (central 7.34). The parity $\alpha_{31} = \pi$ supplies the sign, through the flavor-difference mechanism rather than through any single amplitude. The proton, meanwhile, is absolutely stable, because the same exact baryon number permits change only in units of three.

2.2 What is inherited, and from where

Table 1 records the provenance of every load-bearing element, so that a reader can go to the owning paper for the derivation and to this one for the argument. The discipline is worth stating plainly, because it is easy to lose in a long chain: *no claim in this study is stronger than the weakest grade it inherits.*

2.3 What GG-P5-CPT and GG-W2-Holonomy added after GG-A9-SM and GG-A10-Cosmos

Two of the five source papers are recent, and the reader who knows the older two should be told what changed. GG-A9-SM established the flavor connection, the pentagonal lattice as a clause of T-SM-46.2, the Majorana door, the geometric seesaw, and the light spectrum — and then explicitly delegated the heavy sector: "the heavy-sector realization and its leptogenesis are carried, in completed form," by the cosmology companion. GG-A10-Cosmos built the transport: the Boltzmann system, the washout, the assembled η_B, and the honest register in which it sits.

GG-P5-CPT then did three things neither of those papers had done. First, it turned the pentagon from a clause into a Part: the branch phase is derived from the Higgs doublet's retention of $2/L$, the interior angle $3\pi/5$ and the central angle $2\pi/5$ are distinguished, and the statement *CP violation exists because* $3\pi/5$ *is not a multiple of* π is made in the form of a theorem with its rigidity hypothesis named. Second, it proved the door theorem T-P5-4 in the form "Majorana if and only if weight zero" and drew the consequence that the two Majorana phases are parity data rather than dials — and then located the sense in which $\alpha_{31} = \pi$ drives the excess: *CP*-conserving in the flavor sum, *CP*-violating in the flavor differences. Third, and most consequentially for the honesty of the whole chain, it *sharpened* the resonance condition from a wish into a named hypothesis with an acceptance band, bracketed from both sides, decomposed into six conditions C1–C6 with the crux identified, and equipped with three explicit routes by which it can be killed from within.

GG-W2-Holonomy supplied the anatomy. It is the paper that says *which fiber is doing which job*: that the interval carries no holonomy and can therefore carry no phase; that the circle carries all of it; that *CP* breaking is a property of the content pressed into the circle rather than of the connection, which is why its parameters neither run nor drift; and that the preferred sense of reading — without which a chiral alphabet would be meaningless — is manufactured jointly by chiral content and a one-way read head, with neither ingredient being itself asymmetric. It also located the leptogenesis argument in the anatomy: the neutrino sector is the pivot not by preference but by intersection, being the only sector that can write the surviving book, standing at the only gate the descent opens, holding the only letter that drives.

Those are the developments this study leans on hardest, and Part XI returns to them when assessing what is new.

Table 1: Provenance of the load-bearing elements. Every row is carried at its owning paper's grade; this study upgrades nothing. The one row original here is the helix construction of Part IV, whose principal result is negative. Read the Grade column first and the Element column second: the study's reach is the reach of the weakest grade in the chain it uses, and two rows carry a register no amount of downstream care can improve — the on-resonance condition, which is a named input, and the throughput exponent, which is a postulate.

Element	Owner	Grade as carried
The arena $\mathcal{M}_4 \times I_r \times S^1_\theta$	GG-W2-Holonomy	Derived (topological; premise (T))
Ledger $(1,3,5;7,35)$	GG-A2-GG6, GG-A3-BI6	Derived
Flavor connection; loop holonomies are the physical data	GG-A9-SM T-SM-46.2	Derived
Pentagonal lattice $(\pi/5)\mathbb{Z}$	GG-P5-CPT T-P5-3(iv)	Derived-conditional (rigidity of the holonomy moduli)
Branch phase $3\pi/5$; $\delta^{(1)}_{\mathrm{CKM}} = 2\pi/5$	GG-A9-SM, GG-P5-CPT	Derived-conditional
QD completion $\delta_{\mathrm{CKM}} = 13\pi/35$	GG-P5-CPT	Asserted, flagged post-hoc
$\delta_{\mathrm{PMNS}} = \pi + 2\arcsin(1/7)$	GG-A9-SM, GG-P5-CPT	Derived-conditional (leptonic assignment lemma)
Majorana door: pairing iff weight zero	GG-A9-SM T-SM-44.1; GG-P5-CPT T-P5-4	Derived
$(\alpha_{21}, \alpha_{31}) = (0, \pi)$	GG-A9-SM, GG-P5-CPT	Derived (structural parity datum)
Geometric seesaw $m_\star = \pi v^2/E_{\mathrm{GUT}}$	GG-A9-SM	Derived-conditional (Weinberg-coefficient lemma)
Light spectrum; normal ordering as output	GG-A9-SM	Derived-conditional
$M_{\mathrm{lock}} = \frac{1}{2} M_{\mathrm{Pl}} e^{-K}$	GG-A2-GG6	Derived
Heavy triplet at the lock; transport; Boltzmann system	GG-A10-Cosmos App. I	Derived-conditional
On-resonance condition	GG-P5-CPT C-P5-C(i)	**Named input**
$\epsilon^{\mathrm{res}} \simeq 2.5 \times 10^{-6}$	GG-P5-CPT	Derived-conditional on H-ϕ_{geo}-1
Throughput exponent $K/2$	GG-P5-CPT C-P5-C(ii)	**Postulated**
$\eta_B \approx 7 \times 10^{-10}$ (central 7.34)	GG-A10-Cosmos, GG-P5-CPT T-P5-5	D*, theory band $\times 3^{\pm 1}$
Proton absolute stability, $\Delta B \in 3\mathbb{Z}$	GG-A9-SM T-SM-60.1	Theorem, conditional on the (H2) baryon-descent step
Electron absolute stability	GG-P3-Electron, GG-P2-N3	Theorem within GDOS given GFF-ST7
Helix picture and pitch test	*this study*	Construction; result negative
Flavored Boltzmann integration; the derived error bar	*this study*	Candidate grade, conditional on C4

3 How to read this study

The paper runs on two levels throughout, and the reader is invited to take either or both. Each Part opens with the physical picture in plain language and closes with the arithmetic; the derivations are complete where they are this study's business to state, and are cited to the owning paper where they are not. Parts II and III are the geometry; Part IV is a construction and a negative result; Parts V and VI are the quark sector and its failure; Parts VII through IX are the lepton sector, the mechanism and the number; Part X is proton and electron stability; and Parts XI through XII are the assessment and the experimental program. Part XVI is a question-and-answer Part for the reader who wants the physics without the machinery, and Part XVII holds the worked problems, the dictionary, and a numerical audit in which every figure typeset in this paper is recomputed from its definition.

One editorial commitment governs the whole. Where the framework's own papers flag a step, this one flags it in the same words and in the same place, and never in a footnote when the parent put it in the text. A chain with one postulate and one named hypothesis is a good chain; a chain with one postulate and one named hypothesis *that are hard to find* is a bad paper.

Part II
The Arena: One Fiber that Remembers, One that Forgets

Why a phase can live in only one of the two extra dimensions, and what follows from that the moment it is admitted

---

4 The bulk, and the asymmetry inside it

Everything in this study follows from a single structural asymmetry between the two extra dimensions, so it is worth stating before anything else and then not letting go of it. The six-dimensional bulk of GG-Theory is

$$\mathcal{M}_6 \;=\; \mathcal{M}_4 \times I_r \times S^1_\theta, \qquad I_r = [0, K], \quad K = 35, \qquad S^1_\theta \cong S^1, \tag{2}$$

where $\mathcal{M}_4$ is the ordinary spacetime stage, I_r is the *scale interval* — the fifth dimension, the renormalization-group coordinate

$$r \;=\; \ln \frac{E_{\text{GUT}}}{\mu}, \qquad E_{\text{GUT}} \;=\; M_{\text{Pl}}\, e^{-2\pi} \;\simeq\; 2.28 \times 10^{16}\ \text{GeV}, \tag{3}$$

made geometry — and S^1_θ is the *phase circle*, the sixth dimension, the gauge phase made geometry. On this bulk acts one connection with gauge group

$$\text{Spin}(1,5) \;\times\; \text{SU}(5) \;\times\; \text{U}(1)_B \;\times\; \text{U}(1)_{Qi}. \tag{4}$$

Unification here is not a hope deferred to high energy. It is the statement that Eq. (4) acts on Eq. (2) as *one* object, and everything laboratories distinguish is that one object read on different fibers.

The asymmetry is this. S^1_θ is compact: every loop on it closes, and $\pi_1(S^1_\theta) = \mathbb{Z}$. I_r is an interval: it is contractible, $\pi_1(I_r) = 1$, and it contains no closed loop at all. Everything in Parts III through IX is a consequence.

4.1 Physical picture: a grooved record and a one-way stylus

The image that carries the most physics for the least machinery is a gramophone record. The circle is a *tape* on which the ledger's integers are pressed like a groove; the interval is the *read head*, drawn across the tape once and in one direction. Two features of that picture do real work later.

The first is that the groove is fixed. Nothing in the reading changes what was pressed. This is why the parameters of *CP* violation neither run nor drift: they are holonomy-class data, exact at

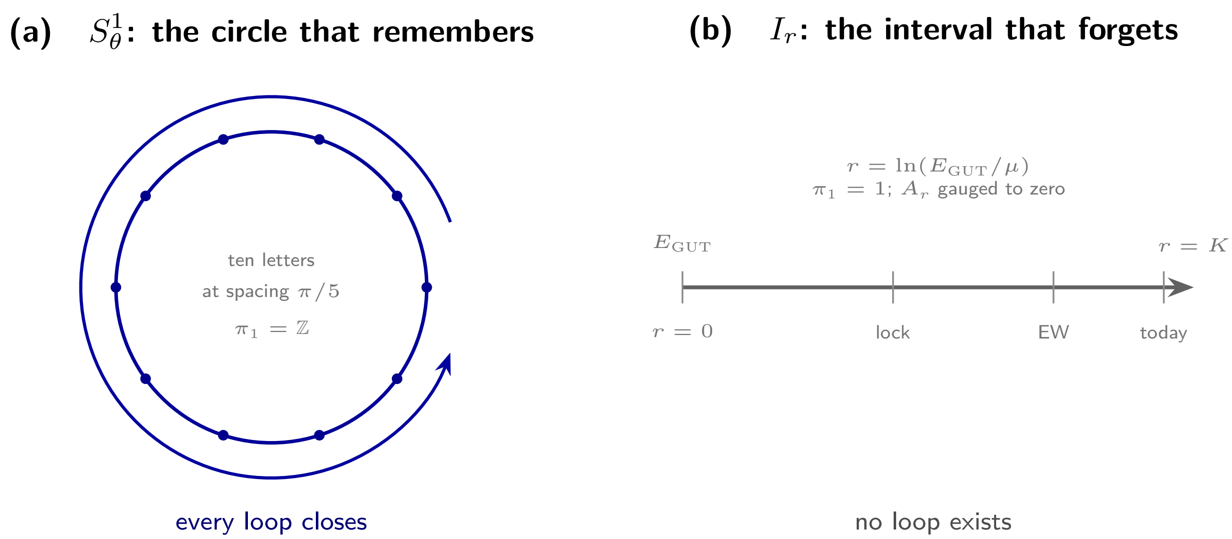

Figure 1: The arena. The sixth dimension (a) is a circle: loops close, the fundamental group is $\mathbb{Z}$, and a connection on it carries a gauge-invariant holonomy. The fifth dimension (b) is an interval: it is contractible, the dr component of the connection can always be gauged to zero, and there is no closed circuit on which to evaluate anything. Every phase in this study lives on (a). Every *amount* — every scale, every suppression, every rate — lives on (b). *The orientations come from the circle; the amounts come from the interval.* After GG-W2-Holonomy.

every scale, measured as angles. When a formula in this framework mixes an exact integer with a continuous scale, the integer is the circle speaking and the scale is the interval.

The second is subtler and is the answer to an objection a careful reader will raise in Part III. A chiral alphabet presupposes that "clockwise" and "counterclockwise" differ — but the circle itself is unoriented, and the laws are even-handed. Where does the preferred sense come from? It is manufactured, and by neither ingredient alone: the ledger's integers press a pattern into the circle whose two readings are not congruent, and the runtime reads the tape in the one direction the interval possesses. The composite — chiral content, one-way reading — produces an effective orientation without either the circle being oriented or the laws being asymmetric. *The laws do not prefer; the tape and the read head jointly do.*

5 A one-page primer: holonomy, from the bench to the bulk

The word "holonomy" does one job in this paper: it names the only kind of quantity that can be a physical phase. Three rungs suffice.

5.1 Rung one: the Aharonov–Bohm phase

Send a charged particle around a solenoid it never enters. The particle sees no field anywhere on its path, and yet the interference pattern shifts by

$$\Delta\varphi \;=\; \frac{q}{\hbar}\oint_C A\cdot dl \;=\; \frac{q\,\Phi}{\hbar}. \tag{5}$$

Three features of Eq. (5) are the whole subject. The charge q is protected — an integer times a unit, not a continuously adjustable coupling. The contour C must *close*: no measurement along an open segment reveals $\Delta\varphi$, because along an open segment the vector potential can be gauged to whatever one likes. And $\hbar$ is the exchange rate between a geometric quantity (flux) and a physical one (phase). What Eq. (5) describes is memory that survives deformation of the path and is readable only on closure.

5.2 Rung two: the Wilson loop

Promote the Abelian phase to a non-Abelian one and the content is the path-ordered exponential

$$W(C) \;=\; \mathcal{P}\exp\!\left(\frac{iq}{\hbar}\oint_C A\cdot dl\right). \tag{6}$$

Wu and Yang (1975) made the point this study needs: the physically meaningful content of a gauge field is neither the potential (too much, gauge-dependent) nor the field strength (too little, it misses the Aharonov–Bohm effect) but exactly the set of non-integrable phase factors — the holonomies. Wilson (1974) made the same identification the working definition of a gauge theory. *A connection is its holonomies, no more and no less.*

5.3 Rung three: two algebras, and only one of them a group

Transport around the circle composes as a group:

$$U(\theta_2)U(\theta_1) = U(\theta_1+\theta_2), \qquad U(\theta)^{-1} = U(-\theta) \text{ always exists.} \tag{7}$$

Transport along the interval composes as a semigroup only:

$$T(r_2)T(r_1) = T(r_1+r_2) \text{ for } r_1, r_2 \geq 0 \text{ only,} \qquad T(r)^{-1} \text{ does not exist.} \tag{8}$$

Coarse-graining loses information; there is no inverse to a renormalization step. The slogan is worth keeping: *what has inverses can promise; what lacks them can happen.* Eq. (7) gives the circle its exact, non-running labels; Eq. (8) gives the interval its arrow.

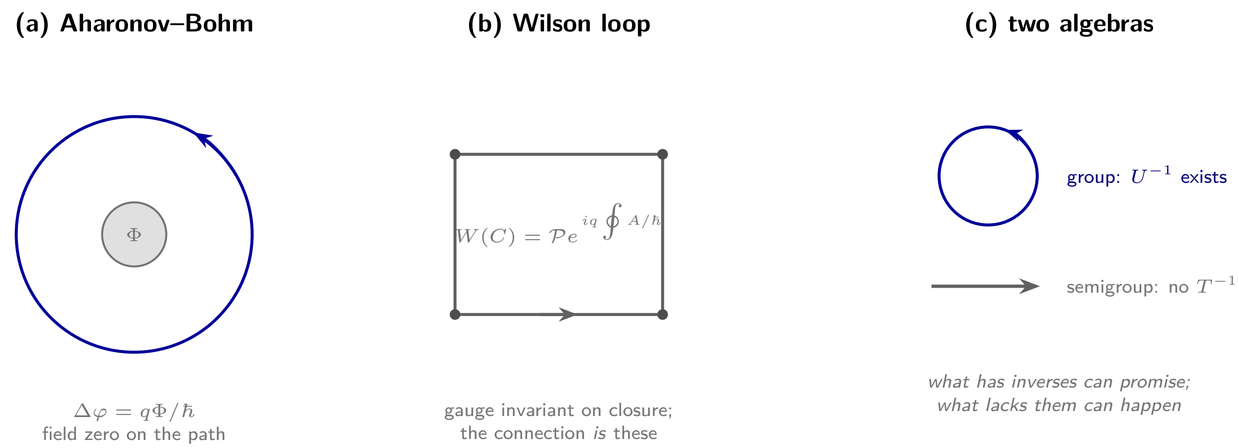

Figure 2: Holonomy, in three rungs. (a) The phase is loop data: it is invisible along any open segment and survives every deformation of the closed path. (b) The gauge-invariant content of a connection is precisely its holonomies. (c) The two extra dimensions carry different algebras — a group on the circle, a semigroup on the interval — and only the group side can hold an exact, non-running label. After GG-W2-Holonomy, §primer.

6 The decisive computation: the connection, fiber by fiber

Decompose the connection on Eq. (2) by fiber,

$$A \;=\; A_\mu(x,r,\theta)\,dx^\mu \;+\; A_r(x,r,\theta)\,dr \;+\; A_\theta(x,r,\theta)\,d\theta, \tag{9}$$

and ask what invariant each term can carry.

(i) The dr component carries nothing. Because I_r is an interval, one may always choose the axial gauge $A_r = 0$; no residual invariant survives, because there is no closed circuit in r on which to evaluate one. This is the technical form of "the interval forgets," and it is the single most consequential line in this Part.

(ii) The dθ component carries a modulus that cannot be gauged away. The Wilson line around the compact direction,

$$W_\theta \;=\; \mathcal{P}\exp\!\left(\frac{iq}{\hbar}\oint A_\theta\,d\theta\right) \;=\; e^{i\alpha}, \tag{10}$$

is defined modulo $\alpha \to \alpha + 2\pi n$, and the residual $\alpha \in [0, 2\pi)$ is a physical, gauge-invariant label. This is the object Hosotani (1983) made dynamical, and the reason a compact extra dimension can break a symmetry with no elementary scalar anywhere in sight.

(iii) The moduli spaces differ as topology says they must. Flat connections modulo gauge form $\mathrm{Hom}(\pi_1(M), G)/G$; for $M = S^1$ this is $G/\mathrm{Ad}(G)$, and for $M = I$ it is a single point. There is nothing on the interval to be a label of.

One boundary, marked here and honored throughout. The exactness claims of this framework trace to the *integrality* in (ii) and (iii) — the quantized shifts and the winding classes

— and not to the continuous residual α, which is a modulus and does move. Where this study says a phase is exact, it is a lattice point of the winding data that is meant, never a value of α read off a potential. GG-W2-Holonomy states the same boundary in the same place, and we do not step over it.

7 Holonomy on the bench: three worked examples

The word "holonomy" should not be taken on faith, and the fastest way to own it is to compute three of them. Each example below is complete; a reader with first-year graduate mechanics and electromagnetism can verify every line, and each teaches exactly one lesson the bulk construction reuses.

7.1 A vector walked around a cone of latitude

Take the unit sphere and walk a tangent vector around the circle of latitude θ (polar angle from the pole), keeping it parallel to itself as honestly as the curved surface allows — parallel transport. When the loop closes, the vector has rotated by the *deficit angle*

$$\Delta\varphi \;=\; 2\pi(1-\cos\theta) \;=\; \Omega_{\text{enclosed}}, \tag{11}$$

the solid angle enclosed by the loop. At $\theta = 60^\circ$: $\Delta\varphi = 2\pi(1-\frac{1}{2}) = \pi$. Walk a gyroscope around that circle and it comes home pointing backwards.

Three lessons. *The rotation is nobody's fault*: at every step the transport was locally parallel; the twist lives in the loop as a whole, not in any segment. *It measures enclosed curvature*: flatten the sphere and it vanishes. *It is quantized by nothing*: θ is continuous, so this holonomy takes any value. Remember that last point — the flavor alphabet will be different, and the difference is the whole game.

7.2 The Aharonov–Bohm phase, with laboratory numbers

Now the electromagnetic version, Eq. (5). Take a solenoid of radius 5 μm carrying an internal field $B = 1$ T, and fly an electron around it at any distance. The enclosed flux is $\Phi = B\pi r^2 = 7.9\times10^{-11}$ Wb; the flux quantum is $\Phi_0 = h/e = 4.14\times10^{-15}$ Wb; so the interference pattern shifts by

$$\frac{\Delta\varphi}{2\pi} \;=\; \frac{\Phi}{\Phi_0} \;\simeq\; 1.9\times10^{4}\ \text{turns}, \tag{12}$$

of which only the fractional part is observable. The electron never touches the field. The lesson beyond §5: the phase is *per unit charge* — double the charge, double the winding. Charge is the gear ratio between geometry and phase, which is why, in the bulk construction, every charge is a winding number and winding numbers cannot change continuously. Electric charge conservation, baryon-number conservation, and lepton-number conservation are all, in this framework, the single statement that an integer cannot drift.

7.3 The two transport algebras, with proofs the size of a postcard

Part II asserted that circle transport is a group and interval transport is a semigroup. Both halves deserve their one-line proofs.

Circle. Transport by angle θ is $U(\theta) = e^{iq\theta\hat{A}}$ for a fixed generator: composition adds exponents, $U(\theta_2)U(\theta_1) = U(\theta_1 + \theta_2)$, and $U(-\theta)$ is always available because $-\theta$ is as good a rotation as θ. Nothing is lost by rotating; rotations are reversible; labels carried by rotations are exact.

Interval. Transport along r is renormalization: coarse-graining modes between μ and $\mu - d\mu$ into effective couplings. Composing two coarse-grainings is a bigger coarse-graining — the semigroup law — but the inverse would require recovering the discarded microstates from the effective couplings, and that map does not exist: infinitely many microphysical theories flow to the same infrared. The non-invertibility is not a technicality; it *is* the second law, written in the scale direction. *What has inverses can promise; what lacks them can happen.*

The consequence used throughout this study: any quantity whose custodian is circle transport (a phase, a winding, a parity) is exact and scale-independent; any quantity whose custodian is interval transport (a coupling, a mass ratio, a rate) runs. When a formula mixes an exact integer with a running quantity — as $\Delta M/\Gamma^{\text{eff}} = 1/(35\, c_\tau y_\tau^2)$ does, one ledger integer over one running Yukawa — each factor is reporting from its own fiber.

8 The Kaluza–Klein tower, done honestly

One more piece of standard machinery, because Part VIII lives at its first floor. A field on $\mathcal{M}_4 \times S^1_\theta$ can be Fourier-expanded around the circle,

$$\phi(x,\theta) \;=\; \sum_{n\in\mathbb{Z}} \phi_n(x)\, e^{in\theta}, \tag{13}$$

and the θ-derivative in the six-dimensional kinetic term hands each mode a four-dimensional mass,

$$m_n^2 \;=\; m_0^2 + \frac{n^2}{R^2}, \tag{14}$$

with R the circle's radius. The picture to keep: the circle is a resonating cavity, the integer n counts nodes around it, and each harmonic weighs more than the last. Three consequences matter here.

The zero modes are the world. Everything in the Standard Model is an $n = 0$ mode — flat around the circle, massless as far as the geometry is concerned, with masses supplied later by the Higgs.

The first floor is the lock. The $n \neq 0$ modes sit at the compactification scale, and in this framework that scale is not free: it is $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-K} = 3.85$ TeV. The Kaluza–Klein excitations of the gluon and of the geometry itself are the lock cluster of §62 — the reason the falsifier is a *pattern* (several states, one address) rather than a single bump.

The heavy Majorana states live there because they must. The weight-zero sector that Theorem 30.1 singles out is the neutral direction of the fiber; its heavy excitations sit at the first floor of the tower with splittings set by the ledger's own spacing, $\Delta M = M_N/K$. The quasi-degeneracy that resonant leptogenesis needs — and that the literature has always had to engineer — is, here, the statement that consecutive rungs of one ladder are close together because $K = 35$ is a large integer. Nothing was arranged.

9 From the bench to the bulk: what is the same, what is new

Now the assembly, with the examples in hand.

The flavor connection of Part III is the sphere example with the graph in place of the sphere: rephasing freedom is the choice of frames along the path, loop angles are the deficit angles, and the CP criterion — every loop angle in $\{0, \pi\}$ — is the statement that the connection's holonomies are all their own mirror images. The Aharonov–Bohm example supplies the reading of charges as windings, hence the exactness of B and L on every charged species, hence Theorem 30.1's monopoly for the neutral one. The two-algebras proof locates every exact quantity on the circle and every running one on the interval. And the Kaluza–Klein tower places the heavy machinery of Part VIII at a computable address.

What is *new* relative to the bench — the one genuinely non-textbook input — is the rigidity of §13: on this bulk the flavor holonomies cannot take the continuous values the sphere's could, because the arena offers no continuous modulus for them to slide along; they are pinned to the pentagonal lattice. Every laboratory consequence in this study — the Cabibbo angle's golden ratio, the 66.9°, the 196.4°, the parity pair, and through Part VIII the amount of matter in the universe — descends from that single non-textbook fact, carried at its owner's grade.

10 The ledger, and why five integers

The integers that will fix every angle and every scale in this study are

$$(M, N, L;\, J, K) \;=\; (1,\, 3,\, 5;\, 7,\, 35), \qquad J = 2N + 1, \qquad K \;=\; LJ \;=\; M^2 + N^2 + L^2. \tag{15}$$

They are not five independent inputs. M counts the time direction. $L = 5$ is fixed by single-tensor Green–Schwarz consistency and is the rank structure of the unified block — and, for this study, the *angle alphabet*: with one circle and one five-block, and no second flavor torus on which an independent angle could float, every first-order physical phase is assembled from steps of $\pi/L = \pi/5$. $N = 3$ follows from the anomaly audit's dual-lock condition $L(2N + 1) = M^2 + N^2 + L^2$, whose roots at $L = 5$ are $\{3, 7\}$, with minimality selecting three; it counts the families, and it is why exactly one quark CP phase survives. $J = 2N + 1 = 7$ counts the chiral steps of the flavor alphabet, and it is the integer inside $\arcsin(1/J)$ that will set the leptonic phase. $K = LJ = 35$ is the depth of the descent, and it sets both the lock scale $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-K} \simeq 3.85$ TeV and the throughput exponent $e^{-K/2}$ of the matter chain.

Two consistency checks are worth doing out loud, because they cost one line each and they are the reason a reader should take the ledger seriously enough to follow it into a baryogenesis argument. $K = LJ = 5 \times 7 = 35$ and $M^2 + N^2 + L^2 = 1 + 9 + 25 = 35$: the same integer twice, along one audited chain. And $\alpha^{-1}(0) = 2^J + N^2 = 128 + 9 = 137$: the fine-structure constant's reciprocal at zero momentum, from two of the same five integers.

Table 2: The ledger, and where each integer acts in this study. Every quantity in the right-hand column is used below; none of them is adjusted. The point of the table is arithmetic rather than rhetorical: five integers, fixed upstream by GG-A2-GG6 and GG-A3-BI6, appear in the matter chain in eleven separate places, and there is no sixth integer and no continuous dial anywhere in the columns to the right.

Integer	Role	Where it appears in the matter chain
$M = 1$	the time direction	the single causal orientation; the quantum-defect numerator $M/(2^J L)$
$N = 3$	families	one surviving quark CP phase; $\Delta B \in N\mathbb{Z} = 3\mathbb{Z}$, hence proton stability; the color triangle's $2\pi/3$
$L = 5$	block rank; the **angle alphabet**	the lattice $(\pi/5)\mathbb{Z}$; the branch phase $\pi(1 - 2/L) = 3\pi/5$; $\delta^{(1)}_{\text{CKM}} = 2\pi/5$
$J = 7$	chiral steps	$\delta_{\text{PMNS}} = \pi + 2\arcsin(1/J)$; $\alpha^{-1}(0) = 2^J + N^2$; the corridor switch e^J
$K = 35$	depth of the descent	$M_{\text{lock}} = \frac{1}{2}M_{\text{Pl}}e^{-K}$; heavy splittings $\delta \sim 1/K$; the QD step π/K; the throughput $e^{-K/2}$; the vacuum's e^{-2K}

11 What this Part has bought

Three statements, and the rest of the study is their consequences.

1. **Every physical phase is a circle-side object.** There is no other place in the geometry for one to live. In particular, no phase is generated by motion along r — a claim Part IV tests quantitatively and confirms by finding the alternative excluded.
2. **Circle-side labels do not run.** They are holonomy classes, fixed by integrality, not by a coupling evaluated at a scale. This is why δ_{CKM} and δ_{PMNS} can be quoted as exact lattice points and compared with measurements made at wholly different energies.
3. **Every amount is an interval-side object.** Scales, suppressions, rates, the arrow of time, the departure from equilibrium. When the chain of Part VIII needs a number rather than an orientation, it will be reading the interval.

The next Part turns the first of these into the origin of CP violation.

Part III
The Holonomic Origin of *CP* Violation

Why a chiral alphabet on a circle makes *CP* violation compulsory in the weak sector and impossible in the strong one

12 Flavor as a connection on a graph

The move that makes everything else possible is to stop thinking of a Yukawa matrix as a table of numbers and start thinking of it as a connection.

Draw the flavor graph. The vertices are fields — three up-type families on one side, three down-type on the other — and each vertex carries a phase dial, the freedom to redefine that field by an arbitrary phase. The edges are couplings: the block y_{ij} connecting up-family i to down-family j, written as a magnitude times a twist, $|y_{ij}|\,R(\varphi_{ij})$. The twist on an edge is a fiber rotation. The graph is bipartite, and it has $N^2 = 9$ edges and $2N = 6$ vertices.

Theorem 12.1 (T-SM-46.2 / T-P5-3: flavor holonomy — rephasing invariants are loop angles; carried from GG-A9-SM, GG-P5-CPT)**.** *Regard a Yukawa texture as a* U(1)*-valued connection on the flavor graph. Then:* (i) *the rephasing group is the group of per-family dial re-zeroings, acting as* $\varphi_{ij} \mapsto \varphi_{ij} + \alpha_i - \beta_j$*;* (ii) *in a spanning-tree gauge every tree edge has angle zero and the residual physical data are exactly the loop holonomies, whose number for the* $N \times N$ *bipartite graph is* $N^2 - (2N-1) = (N-1)^2$*;* (iii) *the quartet invariants* $\mathrm{Im}\,(V_{ij}V_{kl}V_{il}^*V_{kj}^*)$ *are the plaquette sines, and for* $N = 3$ *unitarity ties all of them to a single Jarlskog invariant;* CP *violation holds if and only if some loop holonomy lies outside* $\{0, \pi\}$*;* (iv) *on the GG-6 bulk — one compact circle, no second flavor torus, block integer* $L = 5$ *— the first-order loop holonomies are valued in the pentagonal lattice* $(\pi/5)\mathbb{Z}$*.*

Clauses (i)–(iii) are ordinary flavor physics restated in the right language, and they are worth a moment because the restatement is what makes clause (iv) sayable at all.

On (i). Redefining each field by a phase shifts every edge angle by the difference of its endpoints' dials. That is precisely a gauge transformation of a U(1) connection on a graph. The physical content is what survives it.

On (ii). Choose a spanning tree of the graph — $2N - 1$ edges touching every vertex without closing a loop — and use the dials to set every tree angle to zero. That exhausts the freedom (the global common shift acts trivially, which is why the count is $2N - 1$ and not $2N$). What is left is $N^2 - (2N-1) = (N-1)^2$ angles, one per independent loop, and no further rephasing can touch them. For $N = 3$ that is four. *The physical content of a Yukawa texture is its loop angles.*

On (iii). The smallest closed loop in a bipartite graph has four edges: up-family i to down-family j to up-family k to down-family l and back — a *plaquette.* Its holonomy is the phase of the quartet

$V_{ij}V_{kl}V_{il}^*V_{kj}^*$, the smallest rephasing- invariant product with an imaginary part. For $N = 3$, unitarity ties the nine plaquettes together: all their sines are equal up to sign, and the common magnitude is the Jarlskog invariant (Jarlskog, 1985),

$$\left|\mathrm{Im}\,(V_{ij}V_{kl}V_{il}^*V_{kj}^*)\right| \;=\; J \quad \text{for every plaquette.} \tag{16}$$

One area for the whole graph — geometrically, twice the area of the unitarity triangle, whichever of the six triangles one draws. This is why the global fit's consistency is a one-number statement, why all of quark *CP* violation is priced by a single angle, and why the four loop angles collapse, for observables, to one invariant phase and three magnitudes.

And the *CP* criterion in this language is as clean as it ever gets. *CP* conjugates the texture by complex conjugation; the theory is invariant if and only if the negated texture is a per-family re-zeroing of the original; and that holds if and only if every closed-loop angle lies in $\{0, \pi\}$, the two angles that are their own negatives modulo a re-zeroing.

The criterion, stated once and used everywhere below. *CP* is a symmetry of a flavor texture if and only if every loop holonomy of that texture lies in $\{0, \pi\}$. *CP* violation is therefore not a property of the dynamics. It is a property of the *content* — the set of angles pressed into the circle — and the question of whether our world violates *CP* is the question of whether the available alphabet contains a letter outside $\{0, \pi\}$.

Nothing imaginary has entered at any step. The famous "complex phase" of the CKM matrix is a closed-loop angle on a real circle — the flavor-space analogue of an Aharonov–Bohm holonomy.

13 The alphabet: why the letters are multiples of $\pi/5$

Clause (iv) is where the geometry pays. The argument is an exhaustion, and the honest residue is named at the end of it.

A loop holonomy is a closed transport on the single circle S^1_θ. Ask what could set its value. *Could it be a free continuous parameter?* Only if the fiber carried a continuous modulus for the flavor connection to slide along; the arena is rigid and the holonomy moduli of the flavor connection are set by the block structure. *Could it live on a second flavor torus with its own alphabet?* No: the arena has exactly one compact direction — one circle, hence one alphabet of angles. *Then what angles does that alphabet contain?* The block geometry's own. The unified block carries the integer $L = 5$; the Higgs doublet retains the fraction $2/L$ of that block; the branch phase forced by that retention is

$$\varphi_{12} \;=\; \pi\left(1 - \frac{2}{L}\right) \;=\; \frac{3\pi}{5} \;=\; 108^\circ; \tag{17}$$

and per-family re-zeroings move angles by the step $\pi/L = \pi/5$. Every first-order holonomy is therefore a multiple of $\pi/5$: the pentagonal lattice. Absent a second torus there is no continuous direction in which a holonomy could float off it. Equivalently, the alphabet on the circle has

$$\Delta\theta \;=\; \frac{2\pi}{2L} \;=\; \frac{\pi}{5}, \qquad 2L = 10 \text{ letters}, \tag{18}$$

and the alphabet is *chiral*: read clockwise it does not match itself read counterclockwise.

The grade, carried unchanged. That *no* further continuous deformation of the holonomy moduli exists — "rigidity" — is graded **derived-conditional** by GG-P5-CPT, and it is the single hinge on which clause (iv) turns. This study inherits the clause with the hinge attached and does not lean on it harder than its owner does.

13.1 The sentence

Granted the rigidity, the conclusion is worth stating twice.

$$\boxed{\textbf{\textit{CP} violation exists because } 3\pi/5 \textbf{ is not a multiple of } \pi.} \tag{19}$$

The flavor pentagon is chiral, and no choice, anywhere, was available to make it otherwise. That is a different kind of statement from the one the Standard Model makes. In the Standard Model, δ_{CKM} is a number one measures and could as easily have been zero; the phrase "*CP* violation" names an experimental fact with no explanation attached. Here the fact is compulsory, and its *absence* would have been the thing requiring explanation.

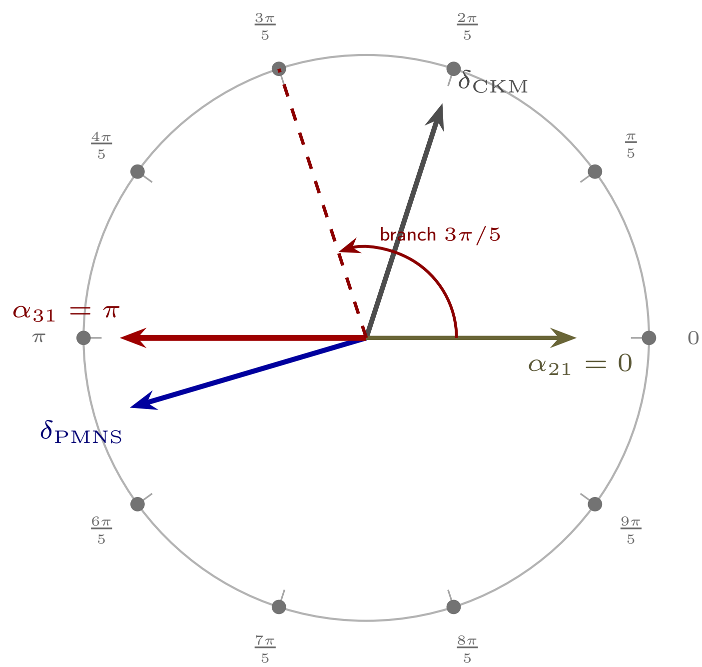

Figure 3: The chiral alphabet on the circle. Ten letters at spacing $\pi/5$, set by the block integer $L = 5$; the alphabet does not match itself read backwards, and that single property is the structural origin of *CP* violation. The forced branch phase $3\pi/5$ (dashed) is the pentagon's interior angle; one unphysical lattice step leaves the physical quark phase $2\pi/5$. Of the four physical letters, $\alpha_{21} = 0$ is inert, δ_{CKM} is a bystander in cosmology, δ_{PMNS} is a spectator at leading order, and the Majorana parity $\alpha_{31} = \pi$ — achiral in the flavor sum, chiral in the flavor differences — is the driver of the matter excess. After GG-W2-Holonomy Fig. "The chiral alphabet".

14 The pentagon, in trigonometry

The pentagon is not decoration. Its two characteristic angles do two different jobs, and conflating them is the commonest way to get this wrong — the present author has done it in conversation and corrects it here in print.

The regular pentagon has *interior* angle $108^\circ = 3\pi/5$ — the angle at each corner — and *central* angle $72^\circ = 2\pi/5$ — the angle each side subtends at the center. In this framework the interior angle is the forced branch phase, Eq. (17), and the central angle is the physical *CP* phase. The pentagon's trigonometry is golden,

$$\cos 36^\circ = \frac{\phi_g}{2} = \frac{1+\sqrt{5}}{4}, \qquad \cos 72^\circ = 2\cos^2 36^\circ - 1 = \frac{\sqrt{5}-1}{4}, \tag{20}$$

with ϕ_g the golden ratio, and Part V will find that golden coefficient sitting inside the measured Cabibbo angle — which is the sharpest available evidence that the pentagon is physically present and not a mnemonic.

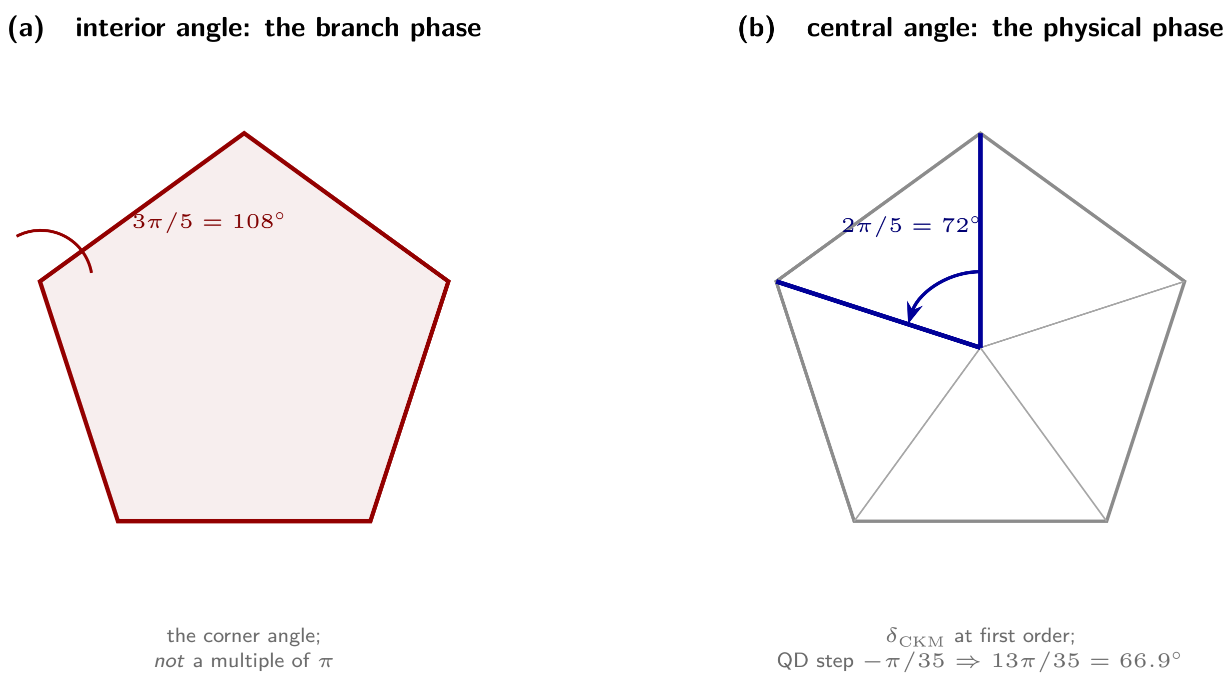

Figure 4: The pentagon's two angles, and why they must not be confused. (a) The interior angle $3\pi/5$ is the branch phase the Higgs doublet's retention of $2/L$ forces; that it is not a multiple of π is the whole of Eq. (19). (b) The central angle $2\pi/5$ is the physical Dirac phase after one unphysical lattice step is removed: raw mismatch three pentagonal steps, one basis choice, two physical. Interior angle in, central angle out.

15 The triangle and the pentagon: why the strong sector does not violate *CP*

Here is the payoff that makes the picture worth having, because it explains a *non*-effect as cleanly as the effect.

The flavor sector presses a pentagon onto the circle: five steps of $2\pi/5$, set by $L = 5$, whose circuit carries a chiral holonomy. The color sector presses the complementary figure — a *triangle*, three steps of $2\pi/3$, set by $N = 3$, the triality of $\mathrm{SU}(3)_C$ — and the triangle differs from the pentagon in the one respect that decides everything: *its circuit is trivial.* Three color steps compose to the identity. The loop closes with nothing left over. The holonomy is 1, which lies in $\{0, \pi\}$, and by the criterion of §12 *CP* is exact.

$$W(\triangle) \;=\; 1 \qquad \text{versus} \qquad W(P_5) \;=\; e^{i\delta}. \tag{21}$$

Closure without residue confines; closure with residue violates.

The same triviality that protects *CP* in the strong sector is what makes color confine: a state must close its color circuit, and the only way to close it with nothing left over is to be a singlet. Confinement and strong-*CP* conservation are two readings of one closure condition. That is an economy the Standard Model does not have, where they are unrelated facts.

Electromagnetism is the limiting case. It is what the fiber that remembers looks like when it simply *holds*: an Abelian phase, one charge, no flavor structure, and therefore no loop in a flavor graph to carry an angle at all — chiral-blind, because a memory with a single letter has no direction of reading.

15.1 Strong *CP*, dissolved in passing

The Standard Model's strong-*CP* problem is the question why the angle $\bar\theta$ multiplying $\mathrm{tr}\, G\widetilde{G}$ is smaller than 10^{-10} when nothing forbids it from being of order one. In this picture $\bar\theta$ is a circle-side angle whose value is not an input but an *occupation*: the compact-angle sector relaxes, by the GRIP cascade, to the stratum's ground attractor, and the attractor floor is the *CP*-conserving point — the same selection that stabilizes the Kaxion. The alignment is not a new mechanism bolted on for the purpose; it is the geometric realization of the Peccei–Quinn mechanism by a field the constitution already forced to exist.

The contrast with the weak sector is exact and is the cleanest statement of why the two fates differ. The $\mathrm{U}(1)_B$ Wilson line is *locked* at a lattice point, $\alpha_B = 2\pi/5$, and a locked angle cannot relax: *CP* violation must exist. The $\mathrm{U}(1)_{Qi}$ Wilson line is *free*, and it rolls to the Vafa–Witten minimum: $\bar\theta$ must vanish. Two Wilson lines, two fates, one mechanism.

16 *C*, *P*, *T* located — and why *CPT* is exact

Each discrete operation reverses an orientation, and in this geometry one can say *which* orientation.

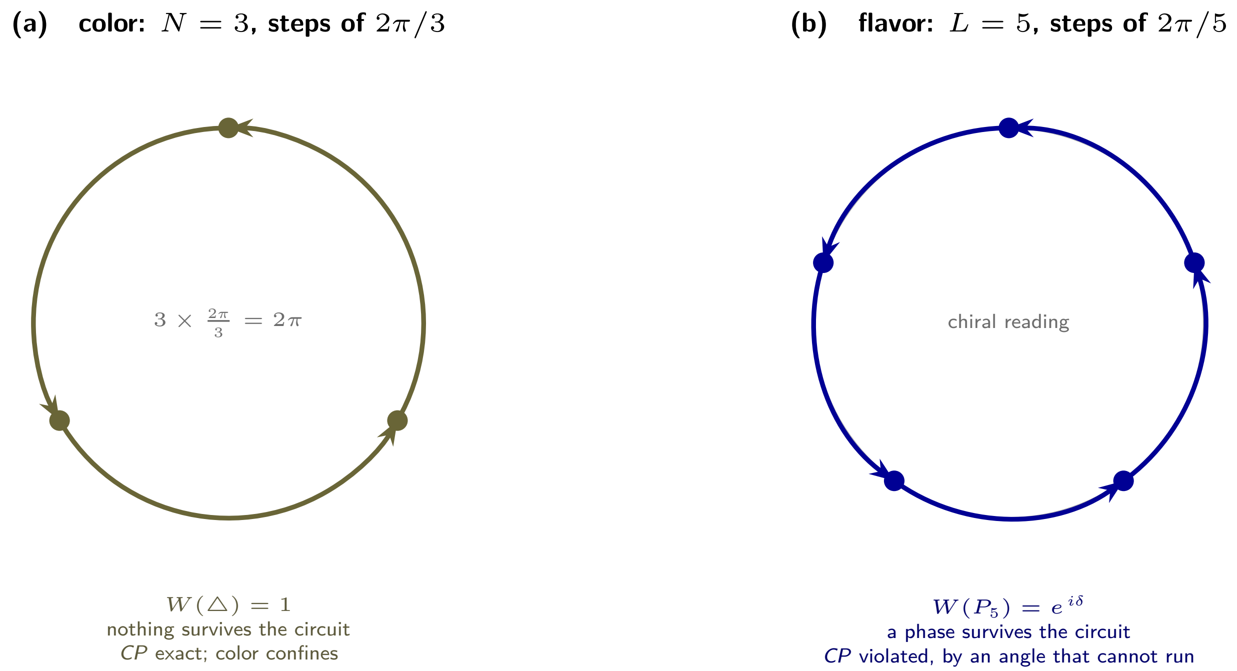

Figure 5: Two figures on one circle, two fates. The color triangle's three steps compose to the identity and its holonomy is trivial; the flavor pentagon's circuit leaves a residue. Confinement and strong-*CP* conservation are the same closure statement read twice; *CP* violation in the weak sector is the failure of that closure. After GG-W2-Holonomy.

$$C: \ \theta \mapsto -\theta, \qquad P: \ x \mapsto -x, \qquad T: \ t \mapsto -t. \tag{22}$$

Table 3: The five discrete operations and the structure each one reverses. The table is the reason *CPT* is exact and *T* is broken at runtime, and it is the dictionary Part VII uses when it needs to know what an antiparticle *is*. Each operation is matched to the piece of the geometry that owns the orientation it reverses, so the fates in the next table are read off this one rather than assumed: there are three orientations and no fourth.

Operation	What it reverses	Owning structure
C	the sense of the circle, $\theta \mapsto -\theta$	S^1_θ, the sixth dimension
P	the handedness of the stage	$\mathcal{M}_4$
T	the sense of the clock	the causal orientation the interval feeds
CP	circle and stage together	mixed
CPT	all three orientations at once	the total geometry

Why CPT *is exact.* Breaking *CPT* would require the bulk to possess a fourth orientation it does not have. The three hypotheses of the *CPT* theorem — locality, Lorentz invariance, positive energy — are, on this geometry, three ways of saying that the arena is the one Eq. (2) describes. There is nothing left to reverse. This is a structural prediction with no parameter attached: a particle–antiparticle mass splitting at *any* precision would end the arena, not merely embarrass it.

Why T *is broken at runtime.* A semigroup has no element that undoes another; a reversed clock would require exactly such an element; so T is broken wherever the runtime acts. Note that this is a *different* T violation from the weak sector's. The microscopic T violation seen in kaons and confirmed directly by BaBar Collaboration (2012) is the circle's contribution — an angle outside $\{0, \pi\}$. The macroscopic irreversibility of the second law is the interval's. The two are independent in origin, and Part VIII will need both at once: the circle supplies the CP asymmetry, the interval supplies the departure from equilibrium.

An antiparticle, in this language, is not a second object but a second *reading*: the same circle-side class traversed in the opposite sense of S^1_θ. Three consequences follow immediately and are all observed. Masses are exactly equal, because an address is orientation-blind. Charges are exactly negated, because every charge is a winding number. And no particle flips its electric charge without flipping its baryon and Qi charges in lockstep, because C is geometric rather than internal.

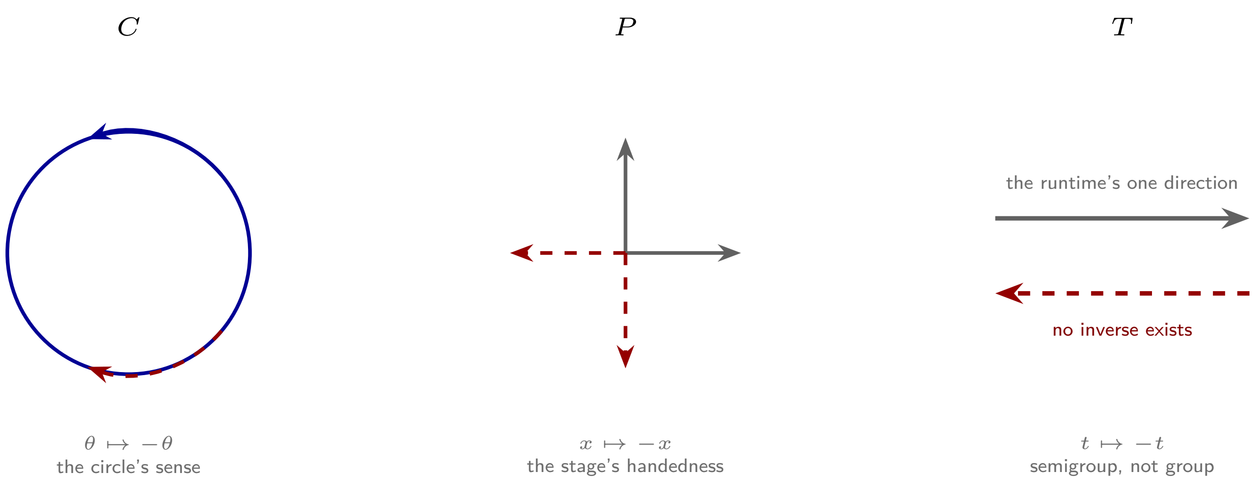

Figure 6: The three orientation reversals. Each discrete operation reverses the orientation of one structure. C acts on the circle, P on the stage, T on the clock the interval feeds. CPT reverses all three at once, which is why it is exact: there is no fourth orientation to violate it with. After GG-P5-CPT and GG-W2-Holonomy.

17 The fates, in one table

The Part closes with the sentence the anatomy gives for the fate of CP: *the connection is symmetric; the tape is not.* CP breaking is not a property of the dynamics but of the groove — the ledger's integers, pressed into the circle as a chiral lattice — and that is why its parameters neither run nor drift. They are holonomy-class data, exact at every scale, measured as angles.

Table 4: Why each symmetry has the fate it has. The pattern is not four independent facts but one structural reading: symmetries whose custodian is the circle's *connection* survive; the one whose custodian is the circle's *content* does not. The last two rows are the ones that cost something to state, because they say the framework predicts *no* strong-*CP* problem and *no* electric dipole moment from QED, and both are places where a measurement could contradict it.

Symmetry	Fate	Because
CPT	exact	no fourth orientation exists to break it
C, P	broken separately	the stage and the circle are independently reversible, and the content is chiral
CP	**broken**, by angles that do not run	the groove is chiral; its letters are holonomy classes, exact at every scale
T	broken at runtime	the interval's transport is a semigroup and has no inverse
CP in QCD	**exact**	the color triangle's circuit is trivial; $\bar{\theta}$ relaxes to the *CP*-conserving floor
CP in QED	exact	one charge, no flavor loop, nothing to carry an angle

Part IV
The Helix: A Picture Built, Tested, and Corrected

A construction original to this study, and the negative result that settles what the phase is

18 The hypothesis worth testing

There is a picture of the weak interaction that is almost irresistible once one has the arena of Part II in mind, and it deserves to be built properly and then tested rather than admired.

Consider the $\tau \to \mu \to e$ decay chain. It is a single descent: the weak interaction is the unique flavor-connecting channel of the boundary gauge group, and the cascade runs from a shallow basin to a deeper one to the ground state. Along that descent the energy scale μ falls, so the fifth coordinate $r = \ln(E_{\mathrm{GUT}}/\mu)$ *increases*. And the sixth dimension is a phase circle on which the flavor alphabet is written. Put the two together and the worldline of a decaying flavor is a *helix*: advancing along r, turning on θ.

The hypothesis that picture suggests is sharper than the picture itself:

$$\textbf{(H-helix)} \qquad \frac{d\theta}{dr} = \text{const}, \qquad \text{the same constant in every sector.} \tag{23}$$

If CP violation were generated by transport along the descent — if the phase were *swept out* rather than *read off* — then there would be a universal pitch, and the phase of a sector would be its pitch times its interval of r. That is a real, falsifiable statement, and the arithmetic to test it is entirely in hand.

We build the helix, test Eq. (23), and report the outcome. It fails, decisively, and the failure is worth more than the picture would have been.

> **Provenance of this Part.** The helix is a construction original to this study. GG-W2-Holonomy contains no helix — the word does not occur in it — and neither GG-A9-SM nor GG-P3-Electron assigns a value of r to an individual quark or lepton. What is inherited is the r axis of Eq. (3) and the lattice-valued phases of Part III. What is constructed here is the pairing of the two, and the r values below are computed in this study from each fermion's own mass, $r = \ln(E_{\mathrm{GUT}}/m)$. They are a construction, not a cohort result, and are labeled as such wherever they appear.

19 Building it quantitatively

Two cascades, four species in one and three in the other.

The quark cascade $b \to c \to s \to u$ runs from $r = 36.24$ to $r = 43.80$, an interval $\Delta r = 7.57$, and ends at the physical quark phase $\delta_{\mathrm{CKM}} = 13\pi/35 = 66.86°$.

The lepton cascade $\tau \to \mu \to e$ runs from $r = 37.09$ to $r = 45.24$, an interval $\Delta r = 8.15$, and ends at the physical lepton phase $\delta_{\mathrm{PMNS}} = \pi + 2\arcsin(1/7) = 196.43°$.

Figures 7 and 8 draw them to scale. Each species sits *on* a thin guide circle at its own r-slice, with a dotted radius showing its accumulated phase; the helix is the locus traced if Eq. (23) held within a sector.

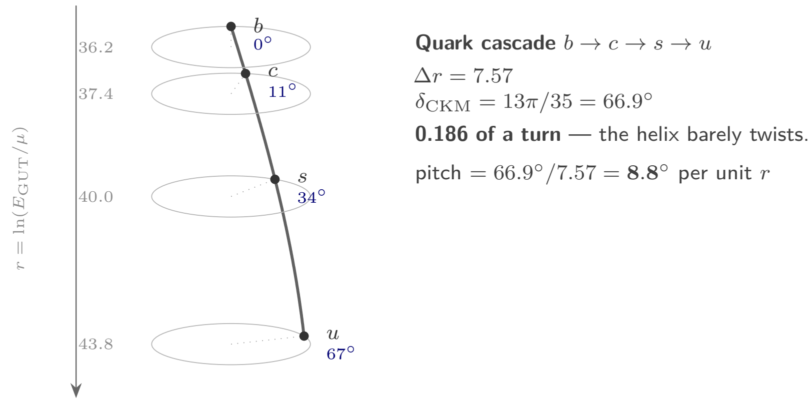

Figure 7: The quark cascade as a helix. Each quark sits on the rim of its own thin guide circle at its own r-slice; the dotted radius marks its accumulated phase, and the figure at the left margin is $r = \ln(E_{\mathrm{GUT}}/m)$. Over more than seven e-folds of scale the phase advances by less than a fifth of a turn. Construction of this study; the r values are computed here from the quark masses and are not a cohort result.

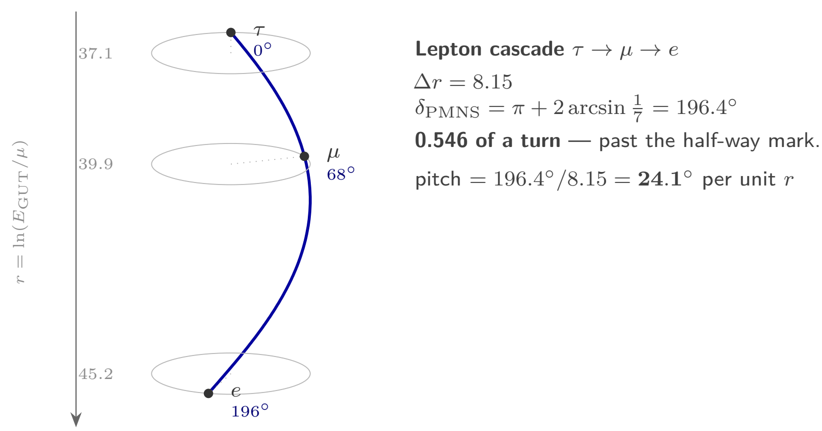

Figure 8: The lepton cascade as a helix. The same construction for the charged leptons. The interval of r is within eight per cent of the quark cascade's, but the phase advances almost three times as far. Compare Figure 7: the two pictures are drawn to the same scale, and the difference between them is the content of §20.

20 The pitch test — and it fails

Two sectors, nearly the same interval of r, wildly different phases.

Table 5: The pitch test. If *CP* phases were swept out along the scale direction at a fixed rate, the two sectors would share a pitch over comparable intervals of r. They differ by a factor of 2.73. The hypothesis **(H-helix)** of Eq. (23) is excluded.

Cascade	r range	Δr	phase	pitch (deg per unit r)
quark $b \to c \to s \to u$	$36.24 \to 43.80$	7.57	66.9°	8.83
lepton $\tau \to \mu \to e$	$37.09 \to 45.24$	8.15	196.4°	24.09
			ratio	**2.73**

The intervals agree to within eight per cent. The pitches differ by a factor of 2.73. There is no reading of Eq. (23) that survives this: a universal pitch is excluded by the framework's own numbers.

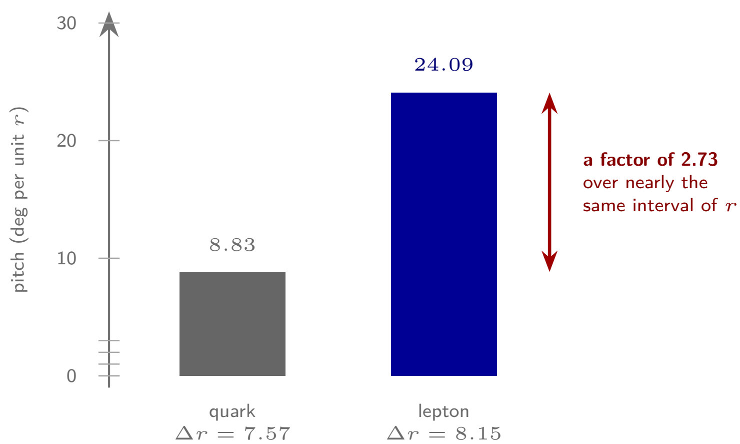

A universal pitch would have made these two bars equal. They are not. The phase is therefore *not* accumulated along the scale direction — it is a lattice angle, already fixed, read off at whatever depth the experiment happens to look. The helix is where the phase *lives*, not what *generates* it.

Figure 9: The pitch test, and what its failure buys. The construction of §19 makes a sharp prediction and the prediction is refuted. This is the outcome the framework requires: Part II proved that the dr component of the connection can always be gauged away, so no phase can be generated along r; the pitch test is the numerical confirmation of that theorem in the one place where a reader's intuition most wants to disagree with it.

21 The retraction, formalized: three survivors

The pitch test killed the mechanism reading of the helix numerically. It is worth also killing it *formally*, because the formal route says precisely which features of the picture are physical and which are decoration — and the inventory it produces is short.

The helix is a curve on the cylinder $I_r \times S^1_\theta$: a trajectory $\theta(r)$ assigning a phase to every depth. But the phase is a fiber coordinate of the $\mathrm{U}(1)_{Qi}$ bundle, and a gauge transformation is a smooth function $\alpha(r, \theta)$ that re-zeroes the fiber independently at every point of the base. Under it the trajectory re-draws itself wholesale:

$$\theta(r) \longmapsto \theta(r) + \alpha\big(r, \theta(r)\big). \tag{24}$$

Walk through what that destroys. The starting angle $\theta(r_0)$ of a cascade can be set to any value by choosing α at the top — so the initial phases marked on Figures 7 and 8 are conventions. Every intermediate angle can be rotated arbitrarily by choosing α on the interior — so the shape of the winding between species is a convention. Even the local slope $d\theta/dr$, the very quantity the pitch test interrogated, can be reshuffled into any smooth profile with the same endpoints — which is *why* the pitch test had to fail: a universal pitch would have been a gauge-dependent statement masquerading as a physical one, and the framework's own arena (§6) had already said the dr component of the connection carries nothing.

Exactly three features of the picture survive Eq. (24), and the reader should check each against the machinery of Part II:

1. **The holonomy.** The closed transport around S^1_θ — equivalently, the lattice letter the cascade carries — is invariant under every gauge transformation continuous around the circle. This is Eq. (10)'s content, and it is where δ_{CKM} and δ_{PMNS} live: not on the trajectory, but in the class the trajectory belongs to.

2. **The reading depth.** The scale at which an experiment reads the letter is a physical address on the interval — the lock crossing at $r = 29.41$ for the matter chain, the laboratory depths $r \simeq 36$–45 for the cascades — fixed by masses, not by gauge. Where a phase is *read* is physical; how the picture interpolates between readings is not.

3. **The winding number.** The integer part of the phase's advance around the circle is a homotopy invariant. For both cascades it is zero — the fractional parts, 0.186 of a turn for the quarks and 0.546 for the leptons, are the lattice letters again, read modulo a full turn. That the two sectors' fractions differ is a statement about two holonomy classes, not about two rates.

The two routes agree, and that is the point. The pitch test (§20) is the *numerical* form of this argument: it computed the one gauge-dependent quantity a mechanism reading needs — a universal $d\theta/dr$ — and found the two sectors disagree by a factor of 2.73. The retraction above is the *formal* form: the slope was never physical to begin with, and the only quantities that survive the gauge freedom are the holonomy, the reading depth, and the winding — exactly the three the rest of this study uses, and none of which is a rate. A picture stripped to its survivors is still worth drawing. It is just not a machine.

Provenance. The formal retraction was prompted by an external cross-check of this material, which reached the gauge-freedom argument independently; it is restated here in the cohort's own language and anchored to the cohort's own reading depths. The numerical pitch test remains this study's, and the two arguments are logically independent — either alone retires the mechanism reading.

22 What the failure teaches

A negative result is only worth reporting when it decides something. This one decides three things.

First, it confirms the arena. Part II showed on purely topological grounds that the interval carries no holonomy: the dr component of the connection can always be gauged to zero, and there is no closed circuit in r on which to evaluate a phase. The pitch test is that theorem's numerical shadow. Had the pitches agreed, the framework would have had a problem — a phase accumulating along a fiber that cannot carry one.

Second, it identifies what the phases are. They are lattice angles, fixed by $L = 5$ and $J = 7$ before any particle decays, read off at whatever depth an experiment happens to look. This is why a value measured in B-meson decays at a few GeV and a value measured in neutrino oscillations at a few hundred MeV can be compared with a single prediction: circle-side labels do not run. The difference between 8.8 and 24.1 degrees per unit r is not a physical rate at all. It is the ratio of two fixed angles divided by two irrelevant intervals.

Third, it fixes the status of the helix. The picture survives — as a picture. The worldline of a decaying flavor really does advance along r while carrying a phase on θ, and drawing it as a helix is an honest way to see where the two dimensions meet. What the picture may *not* be used for is mechanism. It is a map of where the phase lives, not a machine that makes it.

The correction, stated plainly. An earlier reading of this material, in conversation, held the lattice unit to be $\pi/5$ and not $2\pi/5$. Both are right, for different things, and conflating them is the error this Part exists partly to prevent: $\pi/5$ is the lattice *step*, the spacing of the ten letters in Figure 3, and $2\pi/5$ is δ_{CKM} *at first order*, the pentagon's central angle. The step is how far the alphabet moves; the central angle is where the quark sector lands.

Part V

The Quark Sector: One Area for the Whole Graph

From the pentagon's interior angle to a Cabibbo angle with a golden ratio in it, and to the single number that prices all of quark *CP* violation

---

23 The Cabibbo angle, and why the golden ratio is in it

The quark sector is where the pentagon can be caught in the act, because one of its consequences is a four-figure number that was measured long before anyone drew a pentagon.

The chain has six graded steps. We give them all, with the grade of each, because the flags matter more than the agreement.

Step 1 (inherited; parent grade ***Asserted****, chosen post-hoc, carried without upgrade).* The sector mixing angles are the square roots of Yukawa ratios, and on the first-order quark codebook these are pure powers of $e^{-\pi}$:

$$\theta_{12}^{(u)} = \sqrt{y_u/y_c} = e^{-\pi}, \qquad \theta_{12}^{(d)} = \sqrt{y_d/y_s} = e^{-\pi/2}. \tag{25}$$

The base $e^{-\pi}$ is the Bogomolnyi instanton action on the family projector moduli, not an adjustable small parameter; the integer depths are the freedom, and GG-A9-SM grades those eight depths *Asserted* and names the missing per-fermion computation as its dominant open work. We carry that grade.

Step 2 (derived-conditional): the branch phase. The two sector angles do not interfere in phase. The Higgs doublet retains two of the block's five directions, and the branch phase is π diminished by the retained fraction:

$$\varphi_{12} \;=\; \pi\left(1-\frac{2}{L}\right) \;=\; \frac{3\pi}{5} \;=\; 108^\circ, \tag{26}$$

the pentagon's interior angle.

Step 3 (derived, given Steps 1–2): the interference. Combine them:

$$\sin\theta_{12} \;=\; \left|\,e^{-\pi/2} - e^{\,i\,3\pi/5}\,e^{-\pi}\,\right| \;=\; 0.225019, \qquad \sin^2\theta_{12} \;=\; e^{-\pi} + e^{-2\pi} + \frac{\sqrt{5}-1}{2}\,e^{-3\pi/2}. \tag{27}$$

The measured value is 0.22501 ± 0.00068. The agreement is at the fourth figure, a pull of $+0.01$.

Look at what is sitting in the middle of Eq. (27). The cross-term coefficient is $2\cos(3\pi/5)$ carried through the modulus, which is $(\sqrt{5}-1)/2$ — the golden ratio's reciprocal. That is not decorative numerology. It is the trigonometric trace of the pentagonal holonomy implied by $L = 5$, appearing in a quantity measured to four figures in B-meson and kaon decays. If the pentagon were a mnemonic, there would be no reason for $\sqrt{5}$ to be in the Cabibbo angle.

Step 4 (derived-conditional): from branch phase to physical phase. The raw mismatch is three pentagonal steps; one of them is basis choice and two are physical:

$$\delta^{(1)}_{\rm CKM} \;=\; \frac{3\pi}{5} - \frac{\pi}{5} \;=\; \frac{2\pi}{5} \;=\; 72^\circ. \tag{28}$$

Interior angle in, central angle out.

Step 5 (asserted; flagged): the quantum-defect completion. The completion applies the shortest allowed throat slip, $-\pi/K$ with $K = 35$. Since $2\pi/5 = 14\pi/35$, one step down is

$$\boxed{\delta^{\rm QD}_{\rm CKM} \;=\; \frac{2\pi}{5} - \frac{\pi}{35} \;=\; \frac{13\pi}{35} \;=\; 1.16688\ \text{rad} \;=\; 66.86^\circ.} \tag{29}$$

The measured value is 1.147 ± 0.026 rad, a pull of $+0.70$.

The flag, carried verbatim. The first-order value Eq. (28) is the forced result. Against the experimental error alone it sits at $+4.2\sigma$; with a one-lattice-step theory width $\pi/35 = 0.0898$ added in quadrature the pull is $+1.17$. The QD step of Eq. (29) is a bounded lattice correction whose direction and size are ledger-determined but whose *necessity is asserted*: the parent codebook flags it "post-hoc," and GG-P5-CPT carries the flag verbatim rather than absorbing it. So does this study. A reader entitled to be skeptical should be skeptical here and not elsewhere: the $+1.17\sigma$ presumes the theory width, and without it the raw lattice point is excluded at four-sigma class. That is precisely why the QD completion carries a flag and not a footnote.

Step 6 (derived, given the above): the area. With PDG-class moduli $s_{12} = 0.2250$, $s_{23} = 0.0418$, $s_{13} = 0.0037$ and the prefactor $c_{12}c_{23}c_{13}^2 = 0.9735$,

$$J_{\rm CKM} \;=\; c_{12}c_{23}c_{13}^2\, s_{12}s_{23}s_{13}\, \sin\delta^{\rm QD}_{\rm CKM} \;=\; 3.1 \times 10^{-5}, \tag{30}$$

against the measured $(3.12^{+0.13}_{-0.12}) \times 10^{-5}$. Across the framework's evaluation layers — first-order versus QD phase, all-first- order moduli versus measured moduli — the same area evaluates between 2.7 and 3.2×10^{-5}. That layer spread, not an error bar, is the band we quote.

24 One area, and what it means that there is only one

Equation (16) said that for $N = 3$ every plaquette of the flavor graph carries the same invariant magnitude. It is worth dwelling on how strong a statement that is.

A general 3×3 unitary matrix has four physical parameters — three angles and one phase — and one might expect CP-odd observables to probe them in several independent combinations. They do not. Unitarity collapses every CP-odd quantity in the quark sector to a single number, geometrically twice the area of the unitarity triangle, and it does not matter which of the six triangles one draws: they all have the same area. This is why the global fit's consistency is a one-number statement. It is also why the question "is there enough CP violation to make the universe?" has a definite answer rather than a range of model-dependent ones: there is exactly one dial, and it has been measured.

Table 6: The quark sector, predicted against measured. The pull column uses the theory width where the framework declares one, and says so. The two δ_{CKM} rows are the same prediction at two orders and are meant to be read together: the first-order $2\pi/5$ survives only with the one-lattice-step width $\pi/35$ carried in quadrature, and against experimental error alone it is a +4.2 pull, which the footnote states rather than hides.

Quantity	Predicted	Measured	Pull
$\sin\theta_{12}$ (Cabibbo)	0.225019, Eq. (27)	0.22501 ± 0.00068	+0.01
δ_{CKM}, first order	$2\pi/5 = 1.2566$ rad	1.147 ± 0.026 rad	$+1.17^{\dagger}$
δ_{CKM}, QD completion	$13\pi/35 = 1.1669$ rad	1.147 ± 0.026 rad	+0.70
J_{CKM}	3.1×10^{-5} (band 2.7–3.2)	$(3.12^{+0.13}_{-0.12}) \times 10^{-5}$	consistent
$\sin 2\beta$	0.708 (layer span 0.707–0.727)	0.709 ± 0.011	−0.1

† with the one-lattice-step theory width $\pi/35$ in quadrature; against experimental error alone, +4.2.

25 The two areas, drawn to scale

The lepton sector's area is derived in Part VII; we draw the comparison here because it is the visual form of the argument that follows in Part VI, and because a reader who sees it once does not need to be persuaded of it again.

With the current global-fit lepton mixing angles and the lattice phase,

$$J_{\text{PMNS}} = c_{13}^2 s_{13} c_{12} s_{12} c_{23} s_{23} \sin\delta_{\text{PMNS}} \simeq -9.4 \times 10^{-3}, \qquad \sin\delta_{\text{PMNS}} = -\frac{8\sqrt{3}}{49} = -0.283, \tag{31}$$

where the closed form for $\sin\delta_{\text{PMNS}}$ at first order follows from $\delta_{\text{PMNS}} = \pi + 2\arcsin(1/7)$ by the double-angle identity — a small check anyone can do in a line, and a reassuring one, because it means the lepton area inherits the ledger integer $J = 7$ exactly.

The ratio is

$$\frac{|J_{\text{PMNS}}|}{J_{\text{CKM}}} \simeq \frac{9.4 \times 10^{-3}}{3.1 \times 10^{-5}} \simeq 3.0 \times 10^2. \tag{32}$$

Three hundred. And even that understates the disparity, for a reason Part VI makes quantitative: what enters baryogenesis is not the area alone but the area *times* a mass-suppression factor that is catastrophic in the quark sector and absent in the lepton one.

26 A remark on where this leaves the quark sector

It is worth being explicit about what the quark sector *is* in this account, because the answer is unusual and is easy to mistake for a dismissal.

The quark sector is where the theory is most sharply exposed to laboratory data. $\sin\theta_{12}$ at the fourth figure, J_{CKM} within its layer band, $\sin 2\beta$ at a pull of −0.1, and δ_{CKM} within a lattice step of the world fit — these are the tests that are being run continuously, right now, by LHCb and

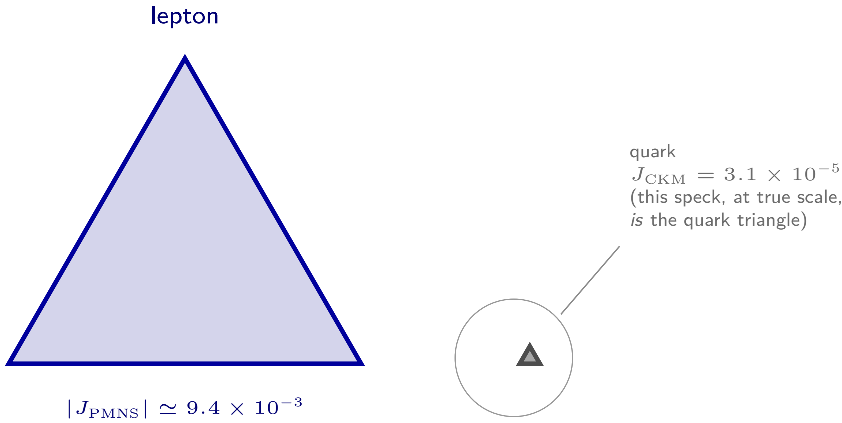

Both areas are real and both are measured. The unitarity triangles are drawn here to *true relative area*: the lepton area is **three hundred times** the quark area. Part VI shows that the disparity relevant to baryogenesis is far larger still, because the quark contribution is throttled a second time by the product of quark mass-squared splittings.

Figure 10: The two unitarity triangles, drawn to true relative area. A single number prices all of quark *CP* violation and a single number prices all of lepton *CP* violation, and they differ by a factor of three hundred. The lepton phase is large because δ_{PMNS} sits near 196° with $|\sin\delta_{\mathrm{PMNS}}| = 0.28$ and, decisively, because the lepton mixing angles are large where the quark mixing angles are small.

Belle II, and a converged world fit excluding $13\pi/35$ at five sigma would kill the completion outright. The quark sector is the framework's laboratory.

What the quark sector is *not* is the origin of us. Its area is real and it is a hundred times smaller than the lepton sector's, and by the time the thermal history is through with it, it is a bystander at the level of ten orders of magnitude. The next Part does that arithmetic carefully, including a temperature sensitivity that is usually skipped and that makes the conclusion stronger, not weaker.

Part VI
Why the Quark Sector Cannot Pay for Us

Sakharov's three conditions located in the geometry, the Standard Model's two failures, and the ten orders of magnitude that force the lepton route

27 Sakharov's conditions, located

Sakharov (1967) gave three conditions. In this framework each is carried by a different structure, and none of them is imported.

Baryon-number violation is carried by the circle's topology meeting the weak sector. The bulk baryonic ledger is exactly conserved — a gauged $\mathrm{U}(1)_B$ — and what changes is the projected subsystem's budget. The one exact non-conservation is the $\mathrm{SU}(2)_L$ anomaly, which changes baryon number in units of $N_{\text{gen}} = 3$ and conserves $B - L$ identically. That single sentence will do two jobs: it opens the door in Part IX, and it welds it shut against the proton in Part X.

C *and* CP *violation are carried by the chiral alphabet* of Part III. The *CP*-odd input is a real, lattice-locked angle of the flavor geometry — a loop holonomy of the coupling texture, the flavor analogue of an Aharonov–Bohm angle. Not a dial.

Departure from equilibrium is constitutional. The boundary law is a forward semigroup; openness is the default, not an achievement. As Eq. (8) put it, what lacks inverses *can happen.* The universe is out of equilibrium by topology, not by tuning.

That is a satisfying allocation. It is also not, by itself, an answer: having all three conditions in principle is exactly the Standard Model's position, and the Standard Model fails anyway. The question is always *how much.*

28 The Standard Model's two failures

28.1 The crossover

Electroweak baryogenesis needs the electroweak transition to be first-order: bubbles of the broken phase nucleating in a symmetric background, with walls across which *CP*-violating reflection can bias the sphaleron traffic, and a strong enough transition that the sphalerons switch off inside the bubble before they wash the asymmetry out. For a Higgs mass of 125 GeV the Standard Model's transition is a smooth crossover (Kajantie et al., 1996). There are no bubbles, no walls, and nothing for a phase to reflect off. This failure is structural: no amount of *CP* violation would rescue it, and it is the reason electroweak baryogenesis in every surviving form requires new physics in the Higgs sector.

28.2 The measure, and its temperature

The second failure is the one this study needs quantitatively, because its size is what forces the lepton route.

Because unitarity collapses all quark CP violation to a single invariant, and because that invariant is multiplied by every mixing angle it needs, the CP-odd effect in the quark sector at temperature T is governed by the dimensionless combination

$$\delta_{CP} \;=\; \frac{J_{\rm CKM}}{T^{12}}\,(m_t^2-m_c^2)(m_t^2-m_u^2)(m_c^2-m_u^2)\,(m_b^2-m_s^2)(m_b^2-m_d^2)(m_s^2-m_d^2), \tag{33}$$

the standard construction of Gavela et al. (1994) and Huet and Sather (1995). The reason for the extra factors is worth stating in words, because it is the physical heart of the failure. A CP-odd observable must be antisymmetric under exchanging any two quarks of the same charge; if two up-type quarks were degenerate, one could rotate between them and remove the phase entirely. So every CP-odd effect must vanish when *any* same-charge pair degenerates, which forces the product of all six mass-squared differences. The quark masses span five orders of magnitude, and their splittings are compared to a temperature of order 100 GeV, twelve times over.

Evaluating Eq. (33) with $J_{\rm CKM} = 3.12 \times 10^{-5}$ and $\overline{\rm MS}$ masses:

Table 7: The Standard Model's CP measure at the electroweak epoch. The T^{-12} means this is a range, not a number, and the range is part of the honest statement. Every entry is recomputed in Appendix 79. The shortfall column is the reason this study exists: at every temperature in the electroweak window the Standard Model's own CP measure falls short of the observed baryon asymmetry by ten orders of magnitude or more, and no choice within the window rescues it.

Evaluation temperature	δ_{CP}	shortfall against $\eta_B = 6.12 \times 10^{-10}$
$T = 100$ GeV	1.2×10^{-19}	9.7 orders
$T = 132$ GeV (crossover)	4.4×10^{-21}	11.1 orders
$T = 160$ GeV	4.2×10^{-22}	12.2 orders

Ten to twelve orders of magnitude. The Standard Model's quark sector delivers a CP-odd measure between 10^{-19} and 10^{-22}, against the 6×10^{-10} the sky requires. It does not matter which end of that range one takes, and it does not matter that the mapping from δ_{CP} to η_B carries $\mathcal{O}(1)$ factors we have not tracked. *The Standard Model predicts, essentially,* $\eta_B = 0$. That is the largest quantitative failure in its ledger.

The temperature sensitivity is worth flagging as a point of method rather than of physics. A single number quoted for "the Standard Model's δ_{CP}" is a number evaluated at a temperature the quoter did not state, and the twelfth power means the choice matters by two and a half orders. We state the range. It changes nothing about the conclusion and it prevents a false precision that would be embarrassing to defend.

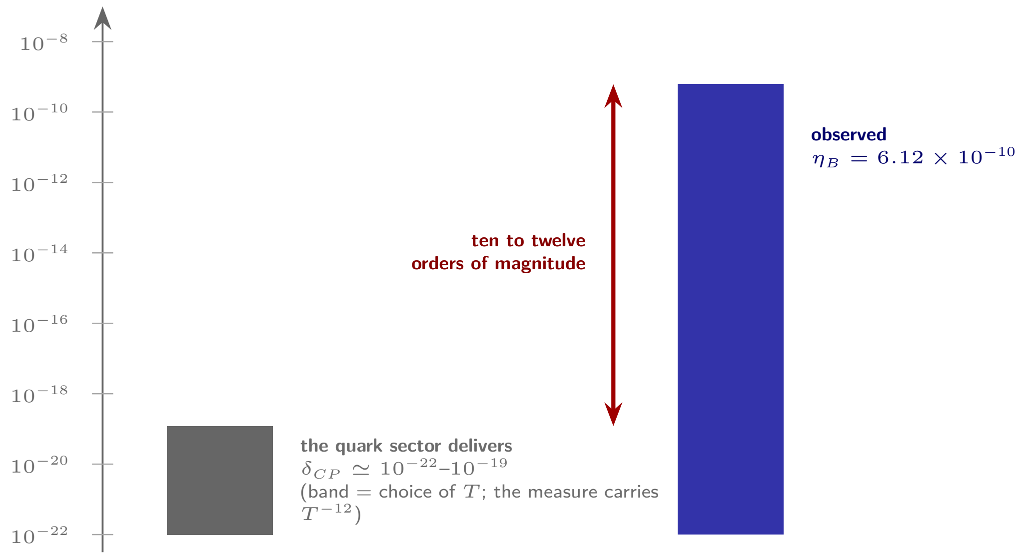

The Jarlskog area is not the whole story. The quark contribution is throttled a second time by the product of quark mass-squared splittings over T^{12} — a factor that exists because a CP-odd effect must vanish whenever any two same-charge quarks degenerate. That second throttle, not the area, is what makes the shortfall fatal.

Figure 11: Why the quark sector cannot pay. The gap is not marginal and cannot be closed by $\mathcal{O}(1)$ improvements in the transport calculation. Together with the crossover of §28, it removes electroweak baryogenesis from consideration inside the Standard Model, and it is the sentence that forces the lepton route.

29 Three doors, and which ones are welded

There are, in the standard taxonomy, three ways to write a baryon asymmetry. In this framework two of them are shut, and the closures are structural rather than circumstantial.

The grand-unified door is welded by exact baryon number. In ordinary grand unification the heavy X and Y bosons mediate baryon-number-violating decays, and one can arrange an asymmetry at the GUT scale. Here that door does not exist: the X, Y are modes of the six-dimensional bulk and are absent from the four-dimensional boundary spectrum, and the bulk $\mathrm{U}(1)_B$ descends to act on the split fields as genuine, component-resolved baryon number. Every operator with $\Delta B \neq 0$ has an identically vanishing coefficient. GUT baryogenesis is not disfavored here; it is forbidden. And the same sentence forbids proton decay, which is the subject of Part X.

The electroweak door is closed by the crossover and by the shortfall, as §28 showed. Both failures, not one.

What remains is the lepton route (Fukugita and Yanagida, 1986). And the framework does not merely find it as the last option standing; it finds it as an intersection of three independent facts, each of which would be enough on its own to point at the neutral sector.

1. **The door.** Sphalerons process $B+L$ while conserving $B-L$ exactly. An asymmetry written in L therefore survives into B; an asymmetry written in $B+L$ does not survive at all. Whatever

writes the surviving book must write it in $B - L$, and the only sector that can is the one with a $\Delta L = 2$ structure.

2. **The address.** The chain needs a once-only, out-of-equilibrium release at a definite depth. The descent supplies exactly one gate at the right depth — the lock crossing at $M_{\text{lock}} = 3.85$ TeV — and it supplies it for reasons that have nothing to do with baryogenesis.
3. **The letter.** The driver of the excess turns out to live among the *Majorana* phases (Part VII), and Majorana structure is native to the neutral sector alone. The charged fermions cannot carry it, because they carry non-zero fiber weight.

The pivot is not a model preference but an intersection: *the only sector that can write the surviving book, standing at the only gate the descent opens, holding the only letter that drives.*

The next Part opens that door and shows that it is the *only* door the geometry admits.

Part VII
The Majorana Door: The Only Way In

Why a Majorana mass exists if and only if the fiber weight vanishes, why the two Majorana phases are parity data rather than dials, and why one of them is the driver of the matter excess

30 Lepton number is winding number

Start from a question that sounds like bookkeeping and turns out to be geometry: *why* is lepton number conserved, and where could a violation of it possibly come from?

In this framework every species carries a weight n under the compact phase direction S^1_θ — the integer counting how its field winds as θ goes once round. Lepton number *is* that winding number. It is exact on every dial-carrying species for the same reason electric charge is exact: a winding number cannot change continuously. So the charged-lepton ledger conserves L identically, and there is no perturbative correction, no loop, no operator of any dimension that can change it there.

The exception, if there is one, must therefore live where the winding number is zero.

Theorem 30.1 (T-SM-44.1 / T-P5-4: the dial-free door — Majorana if and only if weight zero; carried from GG-A9-SM, GG-P5-CPT)**.** *Let a fermion multiplet carry weight n under the compact phase circle S^1_θ (spinor and family indices commuting with the fiber action). A Lorentz-invariant, gauge-invariant Majorana-type self-pairing — equivalently a $\Delta L = 2$ bilinear — exists if and only if $n = 0$. On every $n \neq 0$ sector the unique invariant symmetric self-pairing is the charge-conserving Dirac pairing. And the trilinear LHN is the unique fiber-neutral portal touching the weight-zero sector linearly.*

Proof, in four lines. Write the weight-n pair as a real doublet (a, b) rotating with weight n under θ. Its symmetric square decomposes into a rotation-invariant piece $a^2 + b^2$ — which is exactly the charge-conserving Dirac pairing — and a doublet piece rotating with weight $2n$. The doublet piece is the candidate Majorana bilinear, and it is invariant if and only if $2n\alpha \equiv 0$ for every fiber rotation α, i.e. if and only if $n = 0$. There is no third possibility, because the symmetric square has no third piece. ■

Read the theorem the right way round and it is startling. It does not say that Majorana masses are *allowed* in the neutral sector and that one may choose to write them. It says that the neutral sector is the *only* place the geometry admits a $\Delta L = 2$ structure at all, and that the structure there is unique. As GG-A9-SM puts the consequence: the Majorana channel is the minimal and the unique constitutional option, not the minimal of several.

The bet, and its size. A Dirac-only neutrino sector would require leaving a constitutionally permitted, uniquely minimal term unused — an unforced suppression the framework's economy clause forbids. So the framework does not merely prefer Majorana neutrinos: it stakes itself on them. If the neutrinos are Dirac particles, the driver of Part VIII is deleted and the matter chain retires in its present form. The bet is forced, sized, and on the table.

31 The geometric seesaw: a scale that is already in the ledger

Having a unique door is not yet having a number. The number comes from what the door costs.

The standard seesaw writes the light neutrino mass as $m_\nu = v^2 Y_\nu^2 / M_N$ with M_N a new heavy scale, and its notorious weakness is that M_N is free: one may put it anywhere from a TeV to the GUT scale and adjust Y_ν to compensate. In this framework there is no such freedom, because BI-6 forbids any new intermediate mass scale and any second independent compact loop. The Majorana channel must reuse the already-derived ultraviolet scale.

The unique dimension-five operator is therefore

$$\mathcal{L}_\nu^{(5)} \;=\; \frac{2\pi}{E_{\rm GUT}}\,(LH)\,\Xi_\nu\,(LH) \;+\; \text{h.c.}, \tag{34}$$

in which the suppression scale is the seam $E_{\rm GUT}$ because no new arbitrary scale is admitted, and the prefactor 2π is one full compact traversal of S^1_θ — the natural closed Majorana loop. The neutrino mass scale that follows is

$$\boxed{m_\star \;=\; \pi\,\frac{v^2}{E_{\rm GUT}} \;=\; \pi\,\frac{(246\ \text{GeV})^2}{2.28\times 10^{16}\ \text{GeV}} \;=\; 8.35\ \text{meV}.} \tag{35}$$

That is one of the cleanest formulas in the whole cohort, and it is worth saying why. The neutrino mass scale is not imported from observation and it is not fitted. It is the electroweak scale squared divided by the grand-unification seam, times the unique compact angle constant. Two scales that were already in the ledger for other reasons, and π. And it lands within a factor of two of the measured atmospheric splitting's square root.

A distinction that must not be blurred. GG-A9-SM states, three times and emphatically, that the *light* masses require "no independent see-saw or right-handed-neutrino sector." They come from Eq. (34) directly. The heavy Majorana triplet that Part VIII needs is the *ultraviolet completion* of that same operator, supplied by GG-A10-Cosmos, and it does not introduce a new scale either — it sits at the lock, which the ledger had already fixed. A reader who takes the heavy triplet to be an added ingredient has misread the architecture: the operator is primary and the completion is its realization.

31.1 The light spectrum, and normal ordering as an output

The family projector $\Xi_\nu^{(1)} = P_2 + 2NP_3$, with $2N = 6$ the smallest ledger-coded enhancement that separates the third mass from the second without introducing a new integer, gives at first order

$$(m_1, m_2, m_3)^{(1)} \;=\; (0,\; 1,\; 2N)\, m_\star \;=\; (0,\; 8.36,\; 50.17)\ \text{meV}, \tag{36}$$

and with the quantum-defect completion $m_1^{\text{QD}} = m_\star/2^J$, $m_2^{\text{QD}} = (1 + 1/K)\, m_\star$, $m_3^{\text{QD}} = 2N(1 + 1/2^J)\, m_\star$,

$$\boxed{(m_1, m_2, m_3)^{\text{QD}} \;=\; (0.0653,\; 8.592,\; 50.513)\ \text{meV}, \qquad \Sigma m_\nu = 59.17\ \text{meV}.} \tag{37}$$

The implied splittings are $\Delta m^2_{21} = 7.382 \times 10^{-5}\ \text{eV}^2$ and $\Delta m^2_{31} = 2.552 \times 10^{-3}\ \text{eV}^2$, at pulls of -0.13 and $+1.66$ against the global fit.

Two consequences are worth naming, because both are near-term kill criteria and Part XII returns to them.

Normal ordering is an output, not an assumption. The three light states are assigned to the three Morse levels of the descent, with ν_1 the ground state of the Majorana stratum; the depth assignment *requires* $m_1 < m_2 < m_3$. An experimentally established inverted ordering falsifies the assignment outright, and the framework has no fallback spectrum.

The mass sum is a prediction with no slack in it. $\Sigma m_\nu = 59$ meV sits essentially at the normal-ordering floor. It cannot be lowered by any adjustment, because there is nothing to adjust. This is currently the tightest live tension in the entire chain, and we give it a section of its own in Part XII.

32 δ_{PMNS}: the leptonic phase and the integer inside it

The lepton chain runs on a different ledger integer from the quark chain, and the difference is itself structural. Where the quark alphabet is the block integer $L = 5$, the lepton alphabet is the odd chirality counter $J = 2N + 1 = 7$ — the same integer that counts the generations' doubled hands plus the axis, and the same 7 that appears in $\alpha^{-1}(0) = 2^J + N^2$.

Step 1 (inherited). The minimal neutrino mixing operator carries the reactor-block angle

$$\phi_\nu \;=\; \arcsin\frac{1}{2N+1} \;=\; \arcsin\frac{1}{7}, \tag{38}$$

the smallest angle the odd-integer alphabet writes, appearing where the smallest lepton mixing demands it.

Step 2 (derived-conditional, on the leptonic assignment lemma). The physical Dirac phase is read against the parameterization's real reference channel, and that channel enters the structural form with a negative sign: the measured phase is the block phase shifted by π. Hence

$$\boxed{\delta_{\text{PMNS}} \;=\; \pi + 2\arcsin\frac{1}{7} \;=\; 3.42829\ \text{rad} \;=\; 196.43^\circ, \qquad \delta_{\text{PMNS}}^{\text{QD}} = 196.06^\circ.} \tag{39}$$

A dead branch, pre-registered. A subdominant assignment degeneracy $\{\pi, 3\pi/2\}$ exists at the lemma level and is resolved in favor of the structural form. GG-P5-CPT pre-registers the $3\pi/2$ branch as **dead**: a measurement near 270–286° may *not* be reinterpreted under the alternative branch. The assignment lemma is frozen in the form of Eq. (39), and its failure is the chain's failure. We restate that here because a pre-registration is worth nothing if the companion paper quietly loosens it.

Step 3 (the confrontation). The lattice point sits within 1.0σ of the current normal-ordering global fit and is already disfavored at about 3σ under inverted ordering. The joint T2K–NOvA analysis places the normal-ordering best fit at 203° with a 68% interval of $[166°, 241°]$ (T2K and NOvA Collaborations, 2025), and the NuFIT 6.0 global fit gives $212°\,^{+26}_{-41}$ with the atmospheric sample included and $177° \pm 19°$ without it (Esteban et al., 2024). The lattice point lies inside the one-sigma region of all three.

What that agreement does and does not mean. It means the prediction is not excluded and sits near the center of the current allowed region. It does *not* mean *CP* violation has been established in the lepton sector: $\delta = 180°$, the *CP*-conserving value, is also inside the current intervals in normal ordering. Nobody has measured leptonic *CP* violation, and this study does not pretend otherwise. What the framework offers is a sharp number for a quantity that is about to be measured properly, at $\pm 15°$ and then ± 7–$10°$.

33 The Majorana pair: parity data, not dials

Now the two phases that do not appear in oscillations at all, and that carry the whole of the matter excess.

A Majorana mass matrix for three light states admits, beyond the Dirac phase, two further physical phases $(\alpha_{21}, \alpha_{31})$. In the standard treatment these are two free continuous parameters, unmeasured and unmeasurable except through $0\nu\beta\beta$, and their freedom is precisely why leptogenesis in the literature can accommodate any η_B one likes. Here they are not free.

The constitutionally forced *rank-two* Majorana block fixes one state of each intrinsic parity — the two surviving heavy states are a *parity pair*, not a phase continuum — so

$$\boxed{(\alpha_{21},\ \alpha_{31}) \ = \ (0,\ \pi).} \tag{40}$$

Parity data, not dials. 0 means same parity; π means opposite. There is no continuum to choose from, and no third value.

Where this stands. GG-P5-CPT points forward to T-P5-4 for the rank-two structure, and T-P5-4 as stated is a weight-zero theorem rather than a rank statement; the rank-two structure *per se* is asserted by reference rather than separately derived there. GG-A9-SM grades $\alpha_{31} = \pi$ "Derived (structural)" — a discrete opposite-parity datum of the rank-two block — and $\alpha_{21} = 0$

"Derived (convention)": the simplest admissible value on an as-yet-unmeasured sector. GG-A9-SM's own objection section calls both Asserted. We record the disagreement rather than picking the flattering side, and note that $\alpha_{31} = \pi$ is load-bearing for Part VIII while $\alpha_{21} = 0$ is inert.

33.1 The apparent paradox, and where its resolution actually lies

There is an apparent paradox here that a careful reader will already have noticed. It has a resolution, and it is worth being exact about which step does the work, because the plausible answer is the wrong one.

The criterion of Part III said that CP is conserved if every loop angle lies in $\{0, \pi\}$. But $(\alpha_{21}, \alpha_{31}) = (0, \pi)$ consists of exactly those two *achiral* values. How can the letter that carries the matter excess be a CP-even letter?

First, a kinematic fact — which is true, and is not the answer. A Majorana letter is read at *half angle*. The door's weight-zero pairing splits every angle in two: the $\Delta L = 2$ bilinear is quadratic in the field, so a mass-term phase enters amplitudes through $e^{i\alpha/2}$ and not through $e^{i\alpha}$. Therefore

$$\boxed{\alpha_{31} = \pi \qquad \Longrightarrow \qquad e^{i\alpha_{31}/2} \;=\; e^{i\pi/2} \;=\; i.} \tag{41}$$

This is standard and it is used throughout Part VIII. It is *not*, however, a demonstration that CP is broken. A factor of i in an amplitude is not CP violation; in the standard classification a Majorana phase of exactly π — which is exactly this factor of i — is the signature of *opposite intrinsic* CP *parity* between the mass eigenstates, and opposite intrinsic parity is a CP-conserving configuration. The half angle tells us what the letter does to an amplitude. It does not tell us whether CP is broken, and a reader who knows the classification will say so immediately.

Second, the resolution — which is a statement about *which quantity* vanishes. The textbook statement "Majorana phases of 0 and π conserve CP" is a statement about the *flavor-summed* asymmetry. That sum does vanish here, and it vanishes exactly. The individual flavored asymmetries do not.

Computed at this study's own inputs, in a standard type-I seesaw with a real orthogonal Casas–Ibarra matrix and both the vertex and self-energy contributions at one loop:

$$\epsilon_e : \epsilon_\mu : \epsilon_\tau \;=\; 8.52 \times 10^{-14} \;:\; -4.75 \times 10^{-13} \;:\; 3.90 \times 10^{-13}, \qquad \sum_\alpha \epsilon_\alpha \;=\; 5 \times 10^{-29}. \tag{42}$$

The three are of order 10^{-13}; their sum is sixteen orders of magnitude smaller, which is zero to the precision of the arithmetic. The structural reason is one line: with opposite CP parity between ν_3 and $\nu_{1,2}$ and a real orthogonal R,

$$\mathrm{Im}[Y^*_{\alpha i} Y_{\alpha j}] \;\propto\; \sqrt{m_1 m_3}\, V_{\alpha 1} V_{\alpha 3}\, \left(R_{i3} R_{j1} - R_{i1} R_{j3}\right), \tag{43}$$

which is manifestly nonzero flavor by flavor and sums to zero over α for exactly one reason: unitarity gives $\sum_\alpha V_{\alpha 1} V_{\alpha 3} = 0$. The cancellation is column orthogonality of the mixing matrix. It is not dynamics, which is why it is exact.

$$\boxed{\alpha_{31} = \pi \text{ is } CP\text{-conserving in the flavor sum, and } CP\text{-violating in the flavor differences.}} \tag{44}$$

A zero-sum set of flavored asymmetries is not a zero asymmetry, because the three flavors are not erased at the same rate. That is the whole resolution, and Part VIII carries it out: the electron door is the leaky one — because $\sin^2\theta_{13} = 3/140$, the smallest mixing element there is and a derived one, decouples the electron flavor from the heaviest state — and §39 finds ninety-six per cent of the survivor in it. This study's own Boltzmann integration in §40 is a numerical demonstration of exactly this: three components summing near zero at $z \simeq 1.4$, two of them ground down, one left standing.

Why this is the strongest paragraph in the study rather than the weakest. It is checkable: the flavor sum is a number one can compute and watch vanish, and it does. It explains *why* a flavored treatment is mandatory here rather than merely more careful — an unflavored calculation of this chain returns exactly zero and would have had to be reported as a failure. And it promotes the flavor-sum cancellation, which would otherwise read as a technical nuisance, to the load-bearing step of the argument.

And the consequence for $\delta_{\rm PMNS}$. None of this used the Dirac phase. Moving $\delta_{\rm PMNS}$ across its whole range moves the asymmetry by a few parts in a thousand; switching it off entirely changes the answer by +0.2 per cent, in the wrong direction for it to be the source. Switching the Majorana parity off instead costs a factor of twenty, and switching both off gives 1.3×10^{-29} — the control that shows the computation knows what CP conservation is. The Dirac phase is a spectator in this chain, and the Majorana parity is invisible to every oscillation experiment ever proposed, because Majorana phases cancel out of oscillation probabilities identically.

Answering the question directly. $\delta_{\rm PMNS}$ is unrelated to leptogenesis not because it is small, and not because it lives at the wrong scale, but because it is the *wrong letter*. Oscillations measure the Dirac phase; the matter excess is written by the Majorana parity. The two are different physical quantities, and the only laboratory that can see the second is neutrinoless double beta decay.

34 The shadow the door casts: $m_{\beta\beta}$

The one laboratory observable that touches the Majorana phases directly is the effective mass of neutrinoless double beta decay,

$$m_{\beta\beta} \;=\; \Big| m_1 c_{12}^2 c_{13}^2 + m_2 s_{12}^2 c_{13}^2 e^{\,i\alpha_{21}} + m_3 s_{13}^2 e^{\,i(\alpha_{31} - 2\delta_{\rm PMNS})} \Big|. \tag{45}$$

With the derived spectrum Eq. (37), the parity pair Eq. (40), the lattice Dirac phase Eq. (39), and current global-fit mixing angles, this evaluates to

$$m_{\beta\beta} \;=\; 1.72 \text{ meV (first order)}, \qquad 1.79 \text{ meV (QD layer)}. \tag{46}$$

Note what this is: a *number*, not a range. Every ingredient of Eq. (45) is fixed — three masses, three angles, three phases, none of them free. In the standard treatment $m_{\beta\beta}$ for normal ordering is a band spanning from a few meV down through a cancellation funnel to zero, because the Majorana phases can be chosen to cancel. Here the phases cannot be chosen.

Part XII confronts Eq. (46) with the tonne-scale program and reaches an uncomfortable conclusion honestly: 1.8 meV is *below* the reach of every approved experiment, so this prediction is one-sided. A discovery anywhere in the tonne-scale window of 9–20 meV kills it. An absence proves nothing.

Table 8: The lepton sector: what is fixed, and by which ledger integer. The right-hand column is the point of the Part — there is no continuous parameter anywhere in it. Every entry in the Value column is an integer expression in (N, L, J, K) or a consequence of one, and the Grade column marks all but two of them Derived-conditional — conditional, and therefore falsifiable, on the lemmas the owning papers name.

Quantity	Value	Fixed by	Grade
$\sin^2\theta_{12}$	$L^2/(3N^3) = 25/81 = 0.3086$	L, N	Derived-cond.
$\sin^2\theta_{13}$	$N/(4K) = 3/140 = 0.02143$	N, K	Derived-cond.
$\sin^2\theta_{23}$	$\frac{1}{2}(1+1/J) = 4/7 = 0.5714$	J	Derived-cond.
δ_{PMNS}	$\pi + 2\arcsin(1/J) = 196.43°$	J	Derived-cond.
α_{21}	0	parity of the rank-two block	Derived (convention)
α_{31}	π (**the driver**, read as i)	parity of the rank-two block	Derived (structural)
$m_\star$	$\pi v^2/E_{\text{GUT}} = 8.35$ meV	v, E_{GUT}	Derived-cond.
(m_1, m_2, m_3)	(0.065, 8.59, 50.5) meV	Ξ_ν, QD layer	Derived-cond.
Σm_ν	59.17 meV	the above	Derived-cond.
$m_{\beta\beta}$	1.72–1.79 meV	the above	Derived-cond.
J_{PMNS}	-9.4×10^{-3}	angles $\times \sin\delta_{\text{PMNS}}$	Derived-cond.

Part VIII
Leptogenesis at Four TeV

Why the mechanism sits where the ledger already put a scale, how it evades the bound that puts every other version of it out of reach, and exactly what is being assumed

35 The Davidson–Ibarra wall, and why everyone else is behind it

Leptogenesis is old and it works (Fukugita and Yanagida, 1986): heavy Majorana neutrinos decay out of equilibrium, their decays are CP-asymmetric, the resulting lepton asymmetry is reprocessed by sphalerons into a baryon asymmetry, and $B - L$ carries it safely through. The mechanism is not in doubt. What has always been in doubt is whether it can be *tested*.

The obstacle has a name. For hierarchical heavy neutrinos the CP asymmetry per decay is bounded above by the light and heavy masses (Davidson and Ibarra, 2002),

$$|\epsilon_1| \lesssim \frac{3}{8\pi} \frac{M_1 \, m_3}{\bar{v}^2}, \qquad \bar{v} = 174 \text{ GeV}, \tag{47}$$

so demanding enough asymmetry to make the observed universe forces

$$M_N \gtrsim 10^9 \text{ GeV}. \tag{48}$$

That is a beautiful result and a devastating one. It places the mechanism nine orders of magnitude above anything a collider will ever reach, and it decouples the heavy phases from the measured PMNS phases entirely. Standard thermal leptogenesis is, by its own arithmetic, permanently unfalsifiable in the laboratory.

The known escape is *resonant* leptogenesis (Pilaftsis, 1997, Pilaftsis and Underwood, 2004). If two heavy states are quasi-degenerate, the self-energy contribution to the decay asymmetry is resonantly enhanced and the Yukawa suppression cancels:

$$\epsilon \propto \frac{\Delta M \, \Gamma}{\Delta M^2 + \Gamma^2/4} \sin(2\phi) \qquad \xrightarrow{\Delta M = \Gamma/2} \qquad \tfrac{1}{2} \sin(2\phi). \tag{49}$$

On resonance the asymmetry is order unity regardless of how small the Yukawas are, the bound Eq. (48) evaporates, and TeV-scale heavy neutrinos become admissible. The price, in the standard literature, is that the degeneracy must be *engineered*: one puts the two masses within a decay width of each other by hand, because nothing else does it.

The claim of this Part, stated before the arithmetic. The GG-6 chain sits in exactly the resonant regime, but with the inputs inverted. The TeV scale is not chosen for accessibility — it is the lock scale $\frac{1}{2}M_{\rm Pl}e^{-35}$ that the ledger fixed for its own, entirely unrelated reasons. And the degeneracy question is not engineered away but faced in public, as a named hypothesis with an acceptance band and three routes to its own refutation. *The bound is evaded where the literature evades it; what is new is who pays, and how visibly.*

36 Where the gate is

The thermal history has one feature the mechanism needs and exactly one place to find it.

Reheating after the inflationary phase deposits energy at $T_{\rm rh} \simeq 3\times10^{10}$ GeV — comfortably above the standard thermal-leptogenesis window, comfortably below the GUT edge so that monopoles stay diluted. The bath then cools, and at $t \sim 10^{-13}$ s it crosses the lock at

$$M_{\rm lock} \;=\; \tfrac{1}{2}M_{\rm Pl}\,e^{-K} \;=\; \tfrac{1}{2}M_{\rm Pl}\,e^{-35} \;=\; 3.85\ \text{TeV}, \qquad r = \ln\frac{E_{\rm GUT}}{M_{\rm lock}} = 29.41, \tag{50}$$

producing the Kaxion and the Kaluza–Klein sector as ordinary cooling thresholds. That crossing is the one gate the descent opens at the right depth. Sphalerons then continue to operate down to the electroweak seam at $\mu \sim 100$ GeV, $r \simeq 33.1$, where the crossover shuts the channel off and the door is, as GG-W2-Holonomy puts it, welded cold.

37 The chain, step by step

Proposition 37.1 (T-P5-5: the assembled η_B, at D* grade; carried from GG-P5-CPT)**.** *On the inherited ledger — heavy masses at the lock, the dial-free door, the seesaw-matched Yukawa, the holonomy* CP *phase — and conditional on the on-resonance condition (Conjecture C-P5-C, whose sharp form is the partial-resonance hypothesis H-$\phi_{\rm geo}$-1), the boundary transport yields* $\eta_B \simeq 5.9\times10^{-10}$*, within the factor-$\sim 3$ theory band of the observed* $(6.12\pm0.04)\times10^{-10}$.

Nine steps. We give each with its grade, and we do the arithmetic for each, because the whole value of this chain lies in there being nothing hidden in it.

1. Heavy scale (derived). $M_{N_i} = M_{\rm lock}(1+\delta_i)$ with $M_{\rm lock} = \frac{1}{2}M_{\rm Pl}e^{-K} = 3.85$ TeV and splittings $\delta_i = \mathcal{O}(1/K)$. Numerically the splitting is $\Delta M \simeq M_{\rm lock}/K = 110$ GeV. Reheating at $T_{\rm rh} \simeq 3\times10^{10}$ GeV populates the heavy sector thermally, so the mechanism is *selected by the launch*, not chosen.

2. Portal (derived). Theorem 30.1: the trilinear LHN is the unique fiber-neutral portal. There is no competing operator to consider and no branching fraction to assume.

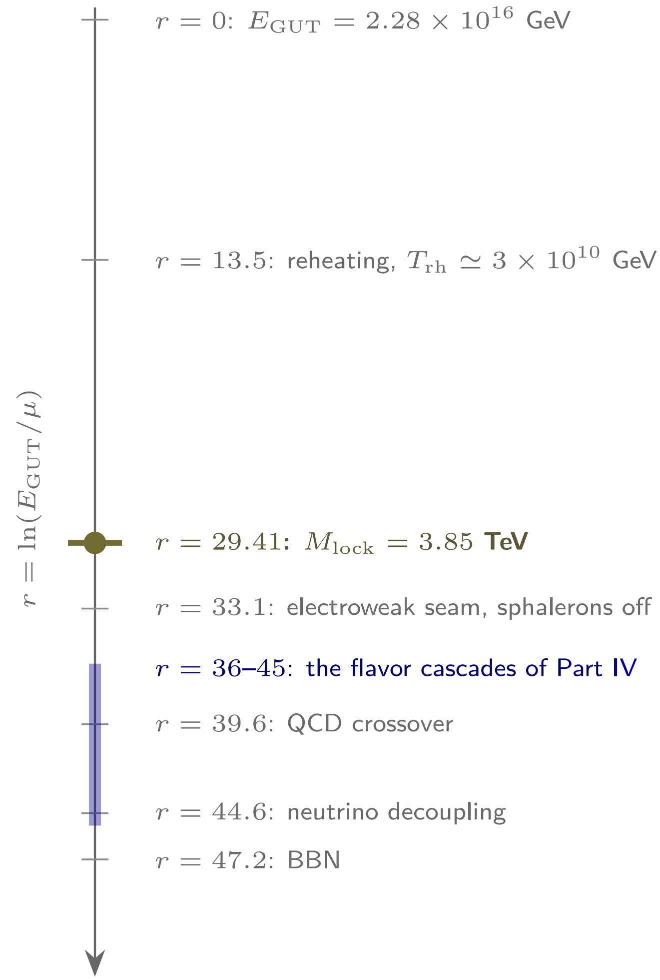

Figure 12: Where on the ladder the excess is paid. One axis, one descent, and one gate at the right depth. The lock crossing at $r = 29.41$ was fixed by $K = 35$ long before any question of baryogenesis arose; that the mechanism needs a gate exactly there is the coincidence the framework claims is not one. **Leptogenesis happens at the lock scale.** Heavy Majorana states, quasi-degenerate at order $1/K$, decay out of equilibrium as the bath crosses. This is *resonant* leptogenesis, evading the Davidson–Ibarra bound $M_N \gtrsim 10^9$ GeV that governs the hierarchical case. Throughput $\epsilon^{\mathrm{res}}\,\kappa = e^{-K/2} = e^{-17.5} = 2.5 \times 10^{-8}$ — the same $K = 35$ that sets the lock depth and the Kaxion's frequency window. **The asymmetry is made before the cascades, not by them.** Leptogenesis is finished by $r \simeq 29$; the helices of Part IV run from $r \simeq 36$ down. What the cascades carry is the *same* pentagonal alphabet, read later and lower — which is why the laboratory phase and the cosmic excess are related but not the same quantity.

3. Yukawa (derived). The seesaw at the lock, matched to the seam coefficient of Eq. (34), fixes the neutrino Yukawa with no freedom:

$$Y_\nu^2 \;\simeq\; \pi\,\frac{M_{\mathrm{lock}}}{E_{\mathrm{GUT}}} \;=\; 5.30 \times 10^{-13}, \qquad Y_\nu \;\simeq\; 7.28 \times 10^{-7}, \tag{51}$$

whence the vacuum decay width

$$\Gamma_N = \frac{Y_\nu^2 M_N}{8\pi} = 8.1 \times 10^{-11}\ \text{GeV}, \qquad \frac{\Gamma_N}{M_N} = 2.1 \times 10^{-14}. \tag{52}$$

It is worth pausing on Eq. (51). A neutrino Yukawa of 7×10^{-7} is not a tuned small number: it is within a factor of four of the electron's Yukawa, 2.9×10^{-6}. The seesaw here is not buying a hierarchy by fiat; it is reading one off two scales that were already in the ledger.

Two conventions, declared. The chain uses $\widetilde{m}$ in two distinct roles and they differ. The heavy-sector resonance width is evaluated at the *seam-matched* one-generation Yukawa, whose effective mass is the seam scale itself, $\widetilde{m} \simeq 8.4$ meV; the *washout* strength is evaluated at the light-neutrino mass the thermal bath actually sees, $\widetilde{m} \sim m_3 \simeq 50.5$ meV. The factor of $\simeq 13$ this induces in Y_ν^2 is declared, accounted for, and priced within the order-of-magnitude register by GG-A10-Cosmos's leptogenesis appendix. We flag it here rather than letting a reader discover it by arithmetic.

4. Resonance (the named input). At $|M_i - M_j| \simeq \Gamma_j^{\text{eff}}/2$ the Breit–Wigner factor of Eq. (49) saturates and $\epsilon^{\text{res}} \simeq \frac{1}{2}\sin(2\phi)$. This is the one quantity the chain requires and does not derive. §38 is devoted to it.

5. *CP* phase (derived-conditional on H-ϕ_{geo}-1). The unflavored *CP* invariant of the resonant formula is U_{PMNS}-independent — the Casas–Ibarra cancellation (Casas and Ibarra, 2001) — and vanishes exactly for the dial-free real-R door. At the TeV lock scale the chain is fully *flavored*, and the inherited phases reach η_B only through the flavored asymmetries $\epsilon_{i\alpha}$. The driver is the Majorana parity $\alpha_{31} = \pi$, which sets the sign of the matter excess — the numerics are J_{PMNS}-blind at the few-percent level — and the result is nonzero at $m_1 = 0$, so the asymmetry does not ride on the quantum-defect layer. The target value is

$$\epsilon^{\text{res}} \simeq 2.5 \times 10^{-6}. \tag{53}$$

6. Washout (derived-conditional). The equilibrium neutrino mass is

$$m_\star^{\text{wash}} = \frac{16\pi^{5/2}\sqrt{g_*}}{3\sqrt{5}} \frac{\bar{v}^2}{M_{\text{Pl}}} \simeq 1.07 \times 10^{-3}\ \text{eV}, \tag{54}$$

so with $\widetilde{m} \sim m_3 = 50.5$ meV the washout strength is $K_{\text{wash}} = \widetilde{m}/m_\star^{\text{wash}} \simeq 47$ — strong washout — and the efficiency is $\kappa \simeq 1/K_{\text{wash}} = 2.1 \times 10^{-2}$. The chain adopts $\kappa = 1.0 \times 10^{-2}$ within the uncomputed $\mathcal{O}(1)$. At the seam convention the numbers would instead read $K_{\text{wash}} \simeq 7.9$ and $\kappa \simeq 0.13$.

7. The throughput identity (the central claim). Multiply steps 5 and 6:

$$\boxed{\epsilon^{\rm res}\,\kappa \;\simeq\; (2.5\times10^{-6})\,(1.0\times10^{-2}) \;\simeq\; 2.5\times10^{-8} \;=\; e^{-K/2} \;=\; e^{-17.5}.} \tag{55}$$

An explicit seesaw-plus-washout chain, built from a Yukawa fixed by the seam and a washout fixed by the light spectrum, *lands on* the ledger's throat-depth exponential rather than asserting it.

Three things about Eq. (55) that must be said in the same breath as the equation. First, the displayed agreement uses the $\mathcal{O}(1)$ freedom in κ: the identity holds at the adopted $\kappa = 1.0\times10^{-2}$, a factor of two inside $1/K_{\rm wash} = 2.1\times10^{-2}$; at the unadjusted value it would read 5.0×10^{-8}. Second, the exponent $K/2$ is a **postulate** of the compression — geometrically suggestive as $\sqrt{e^{-K}}$, not derived — and fixing it from first principles is exactly what Conjecture C-P5-C(ii) must do. Third, $\epsilon^{\rm res}$ itself is carried at derived-conditional grade on H-$\phi_{\rm geo}$-1. What is new here is the reconstruction, not a proof.

8. Sphalerons, entropy, and the equilibrium yield (standard). Part IX.

9. The verdict. $\eta_B \approx 7\times10^{-10}$ (central 7.34), a +20% central deviation inside a factor-~ 3 theory band.

Table 9: The nine-step audit. One named input, one flavored projection carried derived-conditional on a named hypothesis, one postulated exponent, and zero fitted parameters. The audit is reproduced from GG-P5-CPT with the numbers recomputed in Appendix 79.

#	Quantity	Value	Route	Grade
1	M_N	3.85 TeV, $\Delta M \simeq 110$ GeV	$\frac{1}{2}M_{\rm Pl}e^{-K}$	Derived
2	portal	LHN, unique	T-P5-4 / T-SM-44.1	Derived
3	Y_ν	7.3×10^{-7}	$Y_\nu^2 = \pi M_{\rm lock}/E_{\rm GUT}$	Derived
4	resonance	$\lvert\Delta M\rvert \simeq \Gamma^{\rm eff}/2$	—	**Named input** (C-P5-C(i))
5	$\epsilon^{\rm res}$	2.5×10^{-6}	flavored $\epsilon_{i\alpha}$; driver $\alpha_{31}=\pi$	Derived-cond. (H-$\phi_{\rm geo}$-1)
6	κ	1.0×10^{-2}	$K_{\rm wash}\simeq 47$, strong washout	Derived-cond.
7	$\epsilon^{\rm res}\kappa$	$2.5\times10^{-8} = e^{-K/2}$	reconstruction	**Exponent postulated** (C-P5-C(ii))
8	transfer	$7.04\times(28/79)\times1.25/g_*$	Harvey–Turner; entropy; yield	Standard
9	η_B	$\approx 7\times10^{-10}$ (central 7.34)	assembled	D*, band $\times3^{\pm1}$

38 The resonance question, faced: what H-$\phi_{\rm geo}$-1 is, and why it is hard

This is the joint on which the chain turns, and the section exists so that no reader has to go looking for it. It is also the section a skeptical physicist should read first.

38.1 The question, in one sentence

Every resonant-leptogenesis calculation contains a Breit–Wigner factor of the form

$$F_{\rm BW} \;=\; \frac{\Delta M\,\Gamma^{\rm eff}}{\Delta M^2 + (\Gamma^{\rm eff})^2}, \qquad F_{\rm BW} \le \tfrac{1}{2}, \tag{56}$$

and the whole magnitude of the asymmetry rides on it. The mass splitting ΔM is known here — it is the ledger's own level spacing, $\Delta M = M_N/K = 110$ GeV. So everything reduces to a single question:

> ***What is*** $\Gamma^{\rm eff}$ ***?*** Which physical rate regulates the resonance — that is, how quickly is the quantum interference that builds the CP asymmetry destroyed?

That question has an obvious-looking answer which is almost certainly wrong, and a non-obvious answer which is the content of H-$\phi_{\rm geo}$-1.

38.2 Why the obvious answer fails, and why the peak fails too

The obvious answer is that $\Gamma^{\rm eff}$ is the heavy neutrino's own decay width, $\Gamma_N = Y_\nu^2 M_N/8\pi = 8\times10^{-11}$ GeV. That is the vacuum, S-matrix prescription, and it is what a naive reading of Eq. (49) suggests. It gives

$$\left.\frac{\Delta M}{\Gamma^{\rm eff}}\right|_{\rm vacuum} \;=\; \frac{110}{8\times10^{-11}} \;\simeq\; 1.4\times10^{12}, \qquad F_{\rm BW} \sim 10^{-12}, \tag{57}$$

twelve orders off resonance — and the resulting η_B is short of the observed value by seven to eight orders of magnitude. *This branch is fatal, and the framework says so.*

The opposite extreme — sitting exactly on the Breit–Wigner peak, $\Delta M \simeq \Gamma/2$, which is where the standard resonant literature places the mechanism by construction — gives $F_{\rm BW} \simeq \frac{1}{2}$ and $\epsilon \sim \mathcal{O}(1)$. That overshoots the required throughput by about 3×10^4, and would give a universe some ten thousand times more matter than it has. *This branch is fatal too.*

> **GG-A10-Cosmos's own words, quoted rather than softened.** The $\mathcal{O}(1/K)$ splittings motivate *proximity* to the resonant regime, but at vacuum numbers the system sits far off the Breit–Wigner peak — the splitting $d \sim 1/K \simeq 3\times10^{-2}$ against the vacuum width $\Gamma_j/M_j \sim 2\times10^{-14}$ — *so the on-resonance condition is the chain's named conjecture, not its natural consequence, and "natural rather than tuned" would overstate what vacuum numbers support.*

38.3 The hypothesis, in its sharp form

The named hypothesis H-ϕ_{geo}-1 (partial resonance). *There exists a finite-temperature transport computation of the lock-crossing decays in which the effective splitting-to-width ratio at the crossing satisfies $\Delta M/\Gamma^{\text{eff}} \simeq 3 \times 10^4$ — equivalently, a Breit–Wigner factor $\simeq 1.5 \times 10^{-5}$, with $\Gamma^{\text{eff}} \simeq 3$ MeV at $\Delta M \simeq 110$ GeV — so that the flavored, real-R throughput on the inherited phases lands within $\mathcal{O}(1)$ of $e^{-K/2}$.*

The name decodes as follows: **H** for hypothesis, $\boldsymbol{\phi_{\text{geo}}}$ for the geometric phase whose projection into the flavored asymmetries is what the hypothesis is about, and **1** because it is the first such named item in the chain's promotion roadmap. It is not a free parameter and not a fudge factor: it is a *prediction about the outcome of a calculation nobody has yet done*, stated with a number and an acceptance band of $\pm\mathcal{O}(1)$ in ϵ_{eff}, and it can be wrong.

The two bracket numbers quoted throughout this study now decode too, and they are not arbitrary:

$$\underbrace{\frac{3 \times 10^4}{1}}_{\text{peak overshoot } \sim 3\times10^4} \qquad \text{and} \qquad \underbrace{\frac{1.0 \times 10^{11}}{3 \times 10^4}}_{\text{vacuum undershoot } \sim 4\times10^6}, \tag{58}$$

where the second numerator is $\Delta M/\Gamma_N$ evaluated in the washout convention (the one in which Y_ν^2 carries the factor $\simeq 13$ of §37). *The bracket is a ratio of ratios, computed, not asserted.*

38.4 The candidate mechanism: the bath as a flavor-measuring apparatus

The identification that turns H-ϕ_{geo}-1 from a wish into a computation is due to GG-A10-Cosmos Appendix I.3, and it is worth setting out properly because it is the most physically interesting step in the whole chain.

The picture. A thermal plasma is a measuring apparatus. Any interaction that can distinguish one lepton flavor from another acts, on the flavor coherences that the N decays write into the plasma, exactly as a measurement acts on a quantum superposition: it destroys the coherence. So the question "how long does the CP-building interference live?" becomes "how fast does the bath *learn which flavor* was produced?"

The crucial asymmetry. The gauge interactions are enormously faster than anything else in the plasma — $g^2T \simeq 1.5 \times 10^3$ GeV at the lock — but they are *flavor-blind.* They cannot measure what they cannot distinguish. The fastest *honest* measurer of lepton flavor at $T \simeq 4$ TeV is the tau Yukawa interaction, and the muon and electron rates are down from it by $(y_\mu/y_\tau)^2 \simeq 3 \times 10^{-3}$ and below. So

$$\Gamma^{\text{eff}} \;=\; \Gamma_\tau \;=\; c_\tau\, y_\tau^2\, T, \qquad c_\tau = 5 \times 10^{-3}\text{–}10^{-2}, \tag{59}$$

with c_τ the kinetic coefficient of the flavor-equilibration literature.

Why this is not merely a substitution. It rests on a three-step argument, and each step is a numbered condition below.

Step 1 — with R real the source is purely flavor-odd. The flavor trace of the CP source is the unflavored invariant, and for the dial-free door it vanishes identically:

$$\mathrm{Tr}\, S \;=\; \sum_\alpha S_{\alpha\alpha} \;\propto\; \mathrm{Im}\big[(Y_\nu Y_\nu^\dagger)_{ij}^2\big] \;=\; 0 \qquad (R \text{ real}). \tag{60}$$

The entire CP content therefore lives in flavor *differences*: $\sum_\alpha \epsilon_\alpha = 0$ with each $\epsilon_\alpha \neq 0$.

Step 2 — flavor-blind damping cancels from a flavor-odd source. The decohering action of a thermal rate on a flavor matrix is its double commutator with the projector that rate measures. For a flavor-universal rate the projector is the identity, and

$$[\mathbf{1},[\mathbf{1},\delta n_\ell]] = 0 \;\text{ exactly, for every } \delta n_\ell; \qquad [P_\tau,[P_\tau,\delta n_\ell]] = \delta n_\ell \;\text{ on pure tau coherences.} \tag{61}$$

Gauge damping cannot regulate a source it does not distinguish. It drops out of the flavor-odd channel *however large it is* — which is why the biggest rate in the plasma is not the answer.

Step 3 — the two-level reduction. Reducing the heavy pair to a two-level system, the CP-odd weight accumulated by an interference oscillating at ΔM while its flavor coherence is measured at Γ_{fl} is

$$I(\Gamma_{\text{fl}}) = \Gamma_{\text{fl}} \int_0^\infty \mathrm{d}s\, e^{-\Gamma_{\text{fl}} s} \sin(\Delta M\, s) = \frac{\Gamma_{\text{fl}}\, \Delta M}{\Gamma_{\text{fl}}^2 + \Delta M^2} \;\leq\; \tfrac{1}{2}. \tag{62}$$

This is worth reading as physics, not algebra: the interference is a sine wave in proper time, the measurement is an exponential cutoff, and the accumulated asymmetry is the Laplace transform of one against the other. It is a Lorentzian partial-resonance envelope — the Breit–Wigner factor of Eq. (56) with Γ^{eff} in place of the vacuum width, up to a priced convention factor of two off-peak.

38.5 The closed form, and why a ledger integer appears in it

Now the step that makes this a GG-6 result rather than a generic one. At the generation temperature $T \simeq M_N$, the heavy mass *cancels* between the ledger splitting $\Delta M = M_N/K$ and the thermal rate $\Gamma_\tau = c_\tau y_\tau^2 T$:

$$\boxed{\frac{\Delta M}{\Gamma^{\text{eff}}} \;=\; \frac{M_N/K}{c_\tau y_\tau^2 T}\bigg|_{T=M_N} \;=\; \frac{1}{K\, c_\tau\, y_\tau^2} \;=\; \frac{1}{35\, c_\tau\, y_\tau^2} \;=\; (2.7\text{–}5.3)\times 10^4.} \tag{63}$$

Zero dials. The only inputs are the ledger integer $K = 35$ and the running Standard-Model tau Yukawa. No GG-6-sector parameter, no neutrino-sector input, no freedom beyond the order-unity kinetic coefficient c_τ survives into Eq. (63). With one-loop running, $y_\tau(3.85\ \text{TeV}) = 0.0104$, so

$$\Gamma^{\text{eff}} = c_\tau y_\tau^2 M_N = 2.1\text{–}4.1\ \text{MeV}, \qquad F_{\text{BW}} = (0.9\text{–}1.9)\times 10^{-5}, \tag{64}$$

against the hypothesized 3 MeV, 3×10^4, and 1.5×10^{-5}: the target sits inside the band, reproduced exactly at $c_\tau = 8.9\times 10^{-3}$.

Two consistency checks are worth recording because they could have failed. $\Gamma_\tau/H(T = M_N) = (1\text{–}2)\times 10^8$: the fully flavored regime holds with eight orders of margin, so the flavored treatment is not marginal. And $\Gamma_\tau/\Delta M \simeq 3\times 10^{-5} \ll 1$: the mechanism is genuinely *partial*, not an accidental return to the excluded peak.

Why the ledger integer is there at all. The splitting $\Delta M = M_N/35$ is the ledger's own level spacing. Once the bath fixes the measurement rate, the resonance quality $\Delta M/\Gamma^{\text{eff}}$ is *geometry over chemistry* — the same integer that sets the hierarchy e^{-K}, the lock mass $\frac{1}{2}M_{\text{Pl}}e^{-K}$, and the throughput target $e^{-K/2}$ also sets how far off peak the resonance sits, with the tau Yukawa — an object entirely of the Standard Model — supplying the only other factor. And the sign story is untouched: Γ^{eff} is CP-even, so the Majorana parity $\alpha_{31} = \pi$ still sets the sign of the matter excess.

38.6 The regulator landscape: why this is falsifiable rather than decorative

The identification would be worthless if every candidate prescription gave roughly the same answer. They do not. They span sixteen orders of magnitude, and exactly one lands in the decade H-ϕ_{geo}-1 requires.

Table 10: The regulator landscape at the lock crossing ($\Delta M = 110$ GeV, $T = M_N$): four candidate prescriptions for Γ^{eff}, spanning sixteen orders of magnitude in $\Delta M/\Gamma^{\text{eff}}$. **Exactly one lands in the required decade.** This table, more than any other in the study, is why H-ϕ_{geo}-1 is a scientific hypothesis: the alternatives are not nearby, they are catastrophically far away, and each is independently identifiable. Throughput column at the single-T evaluation. From GG-A10-Cosmos Appendix I.3.

Candidate	Γ^{eff} [GeV]	$\Delta M/\Gamma^{\text{eff}}$	F_{BW}	$\eta_B/\eta_B^{\text{obs}}$
(a) N widths (S-matrix prescription)	2×10^{-10}	6×10^{11}	8×10^{-13}	2×10^{-8}
(b) **flavor measurement** $c_\tau y_\tau^2 T$	$(2–4) \times 10^{-3}$	$(2.7–5.3) \times 10^{4}$	$(0.9–1.9) \times 10^{-5}$	**0.28–0.55**
(c) oscillation rate ΔM itself	1.1×10^{2}	1	0.40	1.4×10^{4}
(c′) gauge damping g^2T	1.5×10^{3}	0.07	0.14	1.7×10^{4}
(d) Hubble rate $H(T = M_N)$	2.1×10^{-11}	5×10^{12}	1×10^{-13}	3×10^{-9}

The landscape discriminates cleanly rather than scattering. The N-width class undershoots by $10^{7–8}$; the oscillation-rate and gauge-damping prescriptions re-land on the peak overshoot the flavored recomputation already excludes; the Hubble rate is worse than all of them. *One candidate, and only one, is in the required decade.*

38.7 Why it is hard — five reasons, none of them evasions

It is fair to ask why, if the answer is a two-line formula, nobody has simply computed it. Five reasons, and they compound.

One: the regulator of a resonance is not unique even in the literature. Different resummation schemes for resonant leptogenesis — the effective-Hamiltonian approach, closed-time-path/Kadanoff–Baym, and density-matrix methods — return *different* regulators: Γ_i, $\Gamma_i + \Gamma_j$, or a flavor-projected combination. When $\Delta M \sim \Gamma$ these differ by order-unity factors and nobody minds. When $\Delta M/\Gamma \sim 10^{12}$, as here, they differ by many orders of magnitude and the choice *is* the answer. This is a known open question in the field, not a difficulty peculiar to this framework.

Two: the objects are quasiparticles, not particles. At $T \simeq M_N$ the heavy neutrinos are not asymptotic states with a well-defined vacuum width. They are excitations of a hot plasma with

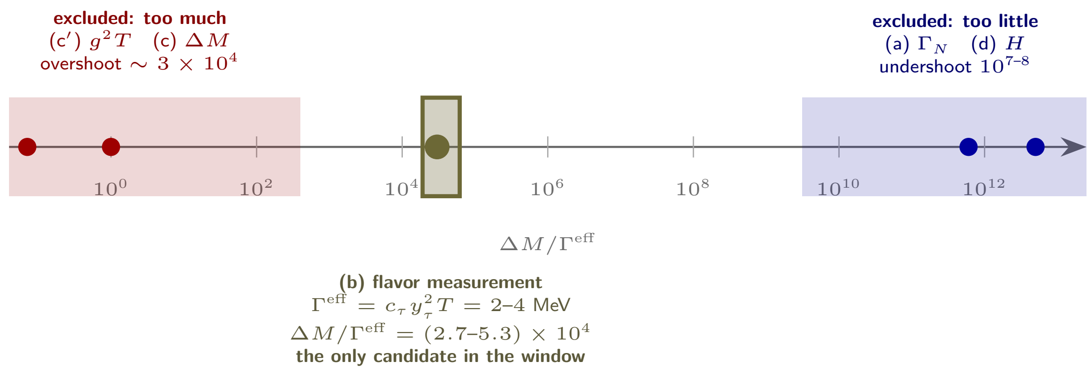

Sixteen orders of magnitude, four candidate prescriptions, and one window. The chain does not sit at a convenient point on a one-sided curve; it sits at a specific point on a scale that is excluded above and below, and the mechanism that puts it there is a closed form in the ledger integer $K = 35$ and a Standard-Model Yukawa, with no dial between them. If the finite-temperature calculation returns any candidate other than (b), the chain dies — and it dies in a direction one can state in advance.

Figure 13: The regulator landscape, drawn. The single most important figure in this study for anyone deciding how much to believe. Compare with the usual situation in resonant leptogenesis, where the degeneracy is imposed and no analogous discrimination exists. What the figure cannot show is a mechanism. It shows that one prescription lands where the observed asymmetry needs it and the others do not, which is a reason to take H-$\phi_{\rm geo}$-1 seriously and is not a derivation of it; the study names it a hypothesis for that reason, and the kill table treats it as one.

thermal masses and thermal widths, and the relevant analytic structure is a spectral function, not a pole.

Three: the coherence that carries the CP *is not in the N lines.* This is the physical heart of C4. What is oscillating is a coherence in the *lepton flavor* sector, sourced by the heavy pair; its lifetime is a decoherence time, not a decay time. Decay times and decoherence times are different objects and in this plasma they differ by seven orders of magnitude.

Four: the calculation is genuinely nonequilibrium. The lock crossing is a time-dependent event, the source and the washout are both active during it, and the required object is a two-time correlation function integrated across a moving threshold. A full Kadanoff–Baym treatment with flavor structure, thermal widths, and the crossing's time dependence is a substantial piece of work in its own right.

Five: the sensitivity is linear, so an order-of-magnitude error in $\Gamma^{\rm eff}$ is an order-of-magnitude error in η_B. There is no suppression of the mistake. Far off resonance $F_{\rm BW} \simeq \Gamma^{\rm eff}/\Delta M$, so the answer scales directly with the regulator, and there is nowhere for an error to hide.

38.8 The six conditions

The derivation above is honest about what it assumes, and each assumption is a numbered condition with a status. Their home is GG-A10-Cosmos Appendix I.3; GG-P5-CPT carries them as row 9 of its promotion roadmap.

Table 11: The six conditions C1–C6. Three are established, one is plausible, one is computed with its closure open, and **one — C4 — is the crux**. The outstanding content of C4 is a full Kadanoff–Baym resummation, and it is the single most valuable calculation anyone could do to this chain, in either direction.

	Condition	Status
C1	flavor-odd source	$\mathrm{Tr}\, S = 0$ exactly for real R, Eq. (60). *Established* (algebraic; machine-checked)
C2	flavor-blind cancellation	gauge damping enters through the identity projector and drops from the flavor-odd channel, Eq. (61). *Established* within the density-matrix framework; its persistence under full resummation folds into C4
C3	two-level reduction	the Lorentzian envelope and the $\lvert\mathcal{R}\rvert \leq \frac{1}{2}$ bound survive for every regulator value, Eq. (62). *Established* at reduction level, given C4
C4	**regulator identification (the crux)**	in the regime $\Gamma_\tau/\Gamma_N \sim 10^7$ with the flavor-summed source zero, the interference time of the flavor-odd source is $1/\Gamma_\tau$, not $1/\Gamma_N$: the vacuum N-width regulator is superseded in-medium. **Open at derivation level** — structurally supported, but the full Kadanoff–Baym/CTP resummation has not been carried out, and the standard fully-flavored Boltzmann limit would return the fatal N-width class instead
C5	the coefficient	$c_\tau \in [5 \times 10^{-3}, 10^{-2}]$ at $T \simeq M_N$; the kinetic-estimate, regime-boundary and CTP equilibration classes of the flavor-equilibration literature agree within a factor two. *Plausible*; the precise value open
C6	window throughput	the T-integrated window factor 0.16–0.27 is part of the honest chain, and the priced $\mathcal{O}(1)$ prefactors must supply a remaining one-sided factor $\simeq$ 2–20 *from below*. *Computed* at Boltzmann level; closure open

The status, exactly, in the parent papers' own words: *a candidate mechanism with a conditional derivation — nothing here upgrades the grade of the chain.* Established: the closed form Eq. (63) as the consequence of the identification, the landscape discrimination of Table 10, and the flavor-odd source structure (C1–C3). Derived-conditional: the effective regulator itself, on C4–C5; the throughput landing, additionally on C6. Open: the full Kadanoff–Baym resummation, the precise c_τ, and the closure of the window factor. *H-ϕ_{geo}-1 remains a named hypothesis — not retired, not closed — now sharpened from a numerical target into the conditions C1–C6.*

38.9 Three routes by which the chain kills itself

The reason H-ϕ_{geo}-1 is a scientific hypothesis and not a placeholder is that its failure modes are written down in advance, and each maps onto a row of Table 10.

1. **If the finite-temperature derivation returns the N-width class** — row (a) — the chain undershoots by $\sim 10^{7\text{–}8}$ and is falsified from within. Note that the *standard* fully flavored Boltzmann limit does exactly this, which is why C4 is the crux and not a formality.

2. **If it returns the flavor-blind** g^2T — row (c′) — the excluded peak overshoot stands, and the verdict is the same. This branch has in fact been examined and excluded: it lands $\Delta M/\Gamma^{\rm eff} \simeq 0.07$, reproducing the overshoot, and fails structurally besides.

3. **If C4 and C5 close but the window deficit of C6 survives** a full transport computation at its adverse end with the priced prefactors exhausted, H-$\phi_{\rm geo}$-1 fails by *throughput* rather than by regulator identity — a third, softer kill beside the two clean ones. This is the failure mode whose size §46 now measures: $E = 7.1$ [4.4, 11.4].

Success is a finite-temperature transport computation — thermal widths in the propagators, the lock crossing's time dependence in the rates — landing at the partial-resonance point, discharging both clauses of C-P5-C at once and promoting Proposition 37.1 from D* toward derived. Failure is the same computation showing the enhancement dead at the crossing for every candidate mechanism — no thermal rescue, no radiative degeneracy, no substructure — in which case the chain is falsified from within, the $e^{-K/2}$ agreement stands exposed as coincidence, and the account reverts to Sakharov custody plus an unexplained number. *Both outcomes are scientific; only the unexamined middle would not be.*

39 The washout, flavor by flavor: the electron door

Everything so far has treated the washout as a single number. It is not, and the difference is where the surviving universe actually comes from. This section computes the washout *per lepton flavor* from quantities GG-Theory has already fixed — the light masses and the PMNS matrix — and finds that the three flavors are washed at wildly different rates. That spread is not a nuisance. It is the mechanism's final, essential ingredient: because the flavor-summed *CP* asymmetry vanishes exactly (Lemma L-Cosmos-U, §38), *a net asymmetry survives only because the flavors are erased unequally.* No spread, no matter.

39.1 Why each flavor has its own washout

The washout processes — inverse decays $\ell_\alpha H \to N$ and their relatives — must pick out a definite lepton flavor α, because at $T \simeq 4$ TeV the tau and muon Yukawa interactions are fast enough to measure flavor (§38.4): the plasma knows an electron from a tau. The strength with which flavor α is washed is set by how strongly the heavy states couple to that flavor, summarized by the *flavor-projected effective mass*

$$\widetilde{m}_\alpha \;=\; \sum_i m_i \left|U_{\alpha i}\right|^2, \qquad K_\alpha \;=\; \frac{\widetilde{m}_\alpha}{m_\star}, \qquad m_\star = 1.07\times 10^{-3}\ \text{eV}. \tag{65}$$

Read Eq. (65) physically: $|U_{\alpha i}|^2$ is the probability that mass state i "is" flavor α, so $\widetilde{m}_\alpha$ is the share of the total seesaw coupling that flavor α owns, and K_α compares that flavor's erasure rate with the expansion rate. Everything on the right-hand side is a derived GG quantity: the masses from Eq. (37), the mixing matrix from Table 8, and $m_\star$ from Eq. (54).

39.2 The numbers, and the asymmetry between them

Putting in the derived spectrum (0.065, 8.59, 50.5) meV and the derived matrix elements:

Table 12: The flavored washouts, computed in this study from the derived masses and mixing matrix (script `s2_boltzmann.py`), against the values GG-A10-Cosmos Appendix I states on the same ledger. The agreement is a genuine cross-check: the two computations were done independently. The last two columns are the efficiency each flavor's asymmetry enjoys — the electron flavor survives an order of magnitude better than the other two.

Flavor	$\widetilde{m}_\alpha$ [meV]	K_α (this study)	K_α (GG-A10-Cosmos)	regime	$\kappa(K_\alpha)$
e	3.7	3.5	3.4	mild washout	0.137
μ	31.3	29.3	28.0	strong washout	0.0084
τ	24.1	22.5	22.4	strong washout	0.0116
sum = $\Sigma m_\nu = 59.2$ meV		55.3			

Why is the electron flavor so lightly coupled? Follow Eq. (65) through: the heaviest light state m_3 dominates everyone's $\widetilde{m}_\alpha$, and the electron row of the mixing matrix barely touches it — $|U_{e3}|^2 = \sin^2\theta_{13} = 3/140 = 0.021$, the smallest matrix element there is, and a *derived* one: $N/(4K)$, the ledger integers again. So $\widetilde{m}_e$ is carried almost entirely by m_2, an order of magnitude lighter. The reactor angle being small is, in this framework, why there is something rather than nothing:

The electron door. The flavor-summed asymmetry is exactly zero. What survives is the difference between how fast the flavors are erased, and that difference exists because $\sin^2\theta_{13} = 3/140$ decouples the electron flavor from the heavy third mass state. The matter in the universe leaks out through the electron flavor — the numerical integration below finds *96 per cent* of the final asymmetry in it — at a rate set by the same ledger integers that fix the reactor angle.

For the reader who wants the efficiency in closed form before the numerics: in strong washout the standard analytic fit is

$$\kappa(K) \;\simeq\; \frac{2}{K\,z_B(K)}\Big(1 - e^{-K z_B(K)/2}\Big), \qquad z_B(K) \;\simeq\; 2 + 4K^{0.13}e^{-2.5/K}, \tag{66}$$

whose physical content is one sentence: the asymmetry that survives is the asymmetry made in the *last* moment before the inverse decays shut off, at $z = z_B$, and by then only a fraction $\sim 1/(Kz_B)$ of the heavy states remain to make it. Everything earlier is erased; κ is the fraction made late.

40 The Boltzmann system, integrated

GG-A10-Cosmos's washout section closes with a sentence this study now answers: *"the $\mathcal{O}(1)$ prefactor . . . awaits a full Boltzmann integration."* This section performs that integration, at Boltzmann level, with every input either derived by GG-Theory or standard. It is the second of this study's two original contributions, and it is carried at the same grade as everything it rests on: **conditional on C4**, upgrading nothing.

40.1 The equations, term by term

Two coupled equations, written in the variable $z = M_N/T$ (time runs with z: early universe at small z, late at large) and in photon-normalized abundances $N_X = n_X/n_\gamma$:

$$\frac{dN_N}{dz} = -\,D(z)\,(N_N - N_N^{\rm eq}), \tag{67}$$

$$\frac{dN_{\Delta\alpha}}{dz} = \underbrace{\epsilon_\alpha(z)\,D(z)\,(N_N - N_N^{\rm eq})}_{\text{source}} \;-\; \underbrace{\tfrac{1}{4}\,K_\alpha\,z^3 K_1(z)\,N_{\Delta\alpha}}_{\text{washout}}. \tag{68}$$

Every symbol earns its place:

$N_N^{\rm eq}(z) = \frac{3}{4}\cdot\frac{1}{2}z^2 K_2(z)$ *per species* is the equilibrium abundance: 3/4 of the photon abundance while relativistic ($z \ll 1$), then Boltzmann-suppressed as e^{-z} once T falls below the mass. The Bessel function K_2 is just the correct relativistic bookkeeping between those limits.

$D(z) = K_{\rm tot}\,z\,K_1(z)/K_2(z)$ is the decay term. The ratio $K_1/K_2 \to z/2$ at small z is *time dilation*: a hot, fast-moving heavy neutrino lives longer in the plasma frame, so decays switch on only as the bath cools toward $z \sim 1$. At large z the ratio $\to 1$ and $D \simeq K_{\rm tot}\,z$: decays are fast and the population tracks equilibrium tightly.

The source says: every net decay (the excess of decays over inverse decays, which is exactly the rate at which N_N follows its falling equilibrium) deposits a *CP* asymmetry $\epsilon_\alpha(z)$ into flavor α. The asymmetry per decay is z-dependent because the regulator is thermal:

$$\epsilon_\alpha(z) \;=\; \bar\epsilon_\alpha\cdot 2R(z), \qquad R(z) \;=\; \frac{2\,\Delta M\,\Gamma^{\rm eff}(z)}{4\Delta M^2 + \Gamma^{\rm eff}(z)^2}, \qquad \Gamma^{\rm eff}(z) \;=\; c_\tau y_\tau^2\,\frac{M_N}{z}, \tag{69}$$

with $\bar\epsilon_\alpha = (+3.4,\,+3.1,\,-6.6)\times 10^{-3}$ the flavored pattern GG-A10-Cosmos computes on the ledger inputs (peak-normalized; note $\sum_\alpha \bar\epsilon_\alpha = 0$, the exact flavor-sum zero of L-Cosmos-U). Since $\Gamma^{\rm eff} \propto T$, the resonance factor $R(z)$ *falls* as the universe cools — the interference is measured more slowly, but the splitting stays put — and that decline is the physical origin of the window suppression that §40.3 measures.

The washout term erases each flavor at its own rate: $\frac{1}{4}K_\alpha z^3 K_1(z)$ is the standard inverse-decay kernel, peaking near $z \sim 2$–4 and dying exponentially once the bath can no longer make heavy neutrinos. K_α from Table 12: this is where the electron door opens.

Scope, stated plainly. The kernels are decays and inverse decays; $\Delta L = 1$ scatterings, spectator redistribution, thermal-mass corrections and the relativistic edge of the window are deliberately omitted and priced in §46 — each is a standard, named effect, and each of the omitted source-side effects points *up*. The equations are integrated from $z_0 = 1$ (scanned 0.5–2) to $z = 50$ with an exponentially-exact stiff stepper.

40.2 Validation before results

A numerical claim in this cohort must arrive with its checks. Three, none of them adjustable after the fact:

Table 13: Validation of the integration. The efficiency check runs the code with a constant unit asymmetry in a single flavor and compares with the standard analytic fit, Eq. (66) — the two were derived independently and agree to four per cent across the full washout range of this problem.

Check	This study	Reference	Agreement
K_α from derived U, masses	(3.5, 29.3, 22.5)	GG-A10-Cosmos: (3.4, 28.0, 22.4)	2–5%
$\kappa(K{=}3.5)$	0.1371	analytic fit: 0.1327	3%
$\kappa(K{=}22.5)$	0.0116	analytic fit: 0.0121	4%
$\kappa(K{=}29.3)$	0.0084	analytic fit: 0.0089	5%
single-T $\epsilon\kappa$ at $c_\tau = 8.9 \times 10^{-3}$	2.7×10^{-8}	ledger $e^{-K/2} = 2.5 \times 10^{-8}$	bookkeeping $\mathcal{O}(1)$

40.3 What the integration shows

Figure 14 is the history of the asymmetry, and it tells the whole story of this Part in one picture. Three results come out of the integration, in ascending order of importance.

First, the sign. Nothing in the integration was told which sign to prefer: the input pattern sums to zero. The net asymmetry emerges positive — matter — because the negative (τ-dominated) components sit in strong washout and are erased, while the positive electron component survives. The sign of every baryon traces through this figure to the Majorana parity, through the flavor filter. That the sign is an *output* of differential erasure, rather than an input, is worth a moment of appreciation.

Second, the window, measured. The single-temperature snapshot of the chain reproduces the ledger throughput; the honest question was always what the full time integration costs, and the compressed chain could only bracket it (0.16–0.27, crude; GG-A10-Cosmos now carries this integration in its Appendix I). The integration answers: at central parameters,

$$\eta_B^{\text{Boltzmann}} \;=\; 1.0 \times 10^{-10}, \qquad \frac{\eta_B^{\text{Boltzmann}}}{\eta_B^{\text{obs}}} \;=\; 0.17. \tag{70}$$

The deficit against observation — the quantity condition C6 named but could not compute — is now a *measured* number with an error bar, not a range spanning a decade: scanning the declared freedoms (§46),

$$\boxed{E \;\equiv\; \frac{\eta_B^{\text{obs}}}{\eta_B^{\text{Boltzmann}}} \;=\; 7.1\,, \qquad 68\% \text{ interval } [4.4,\ 11.4]} \tag{71}$$

against the prior bracket of 2–20, since retired in the parent as well. The target for the finite-temperature transport program is no longer a decade wide.

Third, where the deficit lives. Because the integration resolves time, it localizes what the crude window estimate could not: the asymmetry is generated at $z \simeq 1.5$–4, where the thermal regulator $\Gamma^{\text{eff}} \propto 1/z$ has already fallen from its $z = 1$ value — the resonance is sharpest exactly when the washout is also sharpest, and the survivors are made late, off the best part of the resonance. Every one of the omitted standard effects — the $\Delta L = 1$ scattering source, the relativistic production edge — moves net production *earlier*, where $R(z)$ is larger. This is why the omissions are one-sided, and it is now visible rather than asserted.

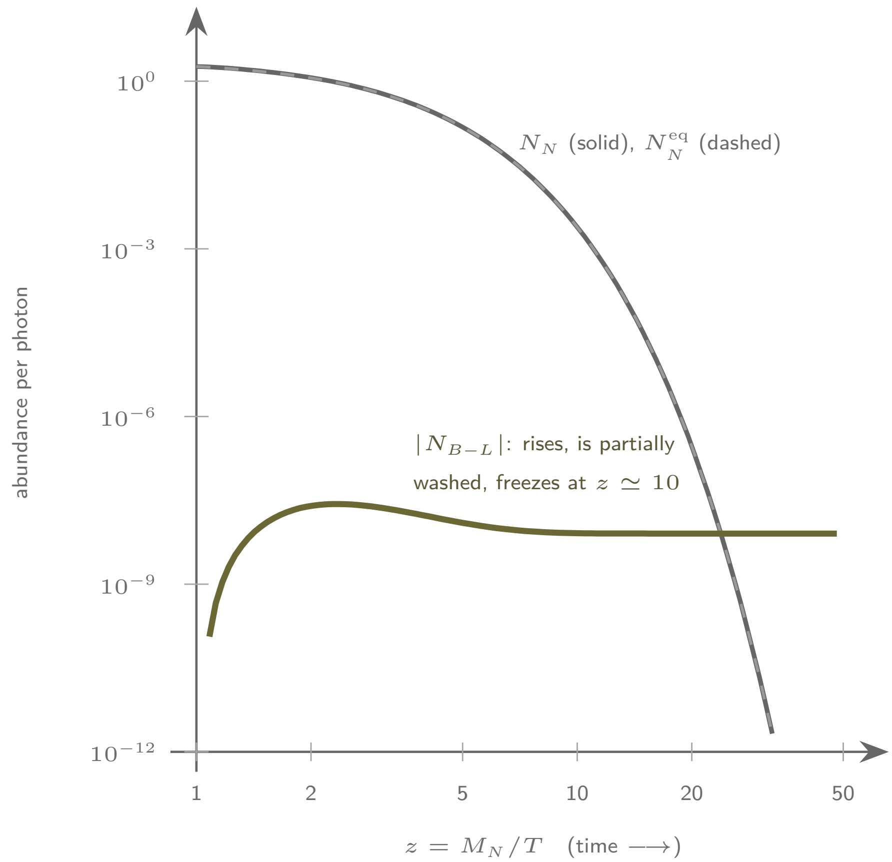

Figure 14: The birth of the asymmetry, integrated. Output of the flavored Boltzmann system Eqs. (67)–(68) at central parameters. The heavy neutrinos track equilibrium so closely that the two upper curves are indistinguishable — strong washout at work. The asymmetry (lower curve) is built between $z \simeq 1.5$ and 4, partially erased as the washout does its work, and frozen once the inverse decays die. At $z = 1.4$ the three flavor components are $(+2.0, +0.9, -2.2) \times 10^{-8}$ — summing nearly to zero, exactly as the flavor-sum theorem requires. By the end, the μ and τ components have been ground down by their strong washouts and **96 per cent of the surviving asymmetry is in the electron flavor**. The sign of the survivor is set by the μ–τ side dying faster: *matter*, exactly as the parity $\alpha_{31} = \pi$ demands on GG-A10-Cosmos's sign audit. This is the figure the six-box diagram of Part IX compresses: gate, letter, toll and door, now as functions of time.

What this integration is, and is not. It is the Boltzmann-level computation GG-A10-Cosmos declared as awaited, performed with pinned GG inputs and validated kernels, and it converts the window freedom and the convention freedom into computed numbers. It is *not* the Kadanoff–Baym closure of C4: the regulator identification $\Gamma^{\rm eff} = \Gamma_\tau$ is assumed throughout, at candidate grade, exactly as in §38. Nothing in this section upgrades any grade. If C4 closes on the N-width class, every number here dies with it, at the first kill route.

41 Zero dilution

One loose end that would otherwise ruin the arithmetic. If the lock crossing were a strongly first-order transition with substantial supercooling, the entropy released on completion would dilute whatever asymmetry had been made, by a factor that would have to be computed and would be a free parameter in all but name.

It is not. The crossing is prompt, on three independent grounds: the nucleation action sits some fifty-six orders of magnitude above the supercooling threshold; the saddle is barrier-free; and there is no large-N enhancement to change either statement. The dilution factor is therefore unity, and the same promptness is what makes the number of e-folds after the crossing $N_2 = 0$ — which is, incidentally, why the inflationary observables $n_s = 1 - 2/N^* = 0.9620$ at leading order (0.9634 exact) and $r = 12/N^{*2} = 0.0043$ come out where they do. One structural fact, two unrelated predictions, no adjustment in either.

Part IX
From Lepton Number to Baryon Number: The Predicted η_B

The sphaleron door, the bookkeeping, the six factors one at a time, and why the excess is one part in a billion

42 The door that converts

An asymmetry in lepton number is not yet an asymmetry in baryon number, and the step between them is where the geometry does its second piece of work.

The $\mathrm{SU}(2)_L$ anomaly gives

$$\partial_\mu J^\mu_B \;=\; N_{\rm gen}\,\frac{g^2}{32\pi^2}\,\mathrm{tr}\,W\widetilde{W}, \qquad \Longrightarrow \qquad \Delta(B+L)\neq 0, \qquad \Delta(B-L)=0. \tag{72}$$

At high temperature the associated sphaleron transitions are unsuppressed (Kuzmin et al., 1985): the door is open. Two features of Eq. (72) do the work.

$B-L$ is exactly conserved. An asymmetry written in $B-L$ cannot be erased by any sphaleron process at any temperature. This is why leptogenesis works at all: the lepton asymmetry made at the lock is a $B-L$ asymmetry, and it is safe.

Baryon number changes only in units of three. Each family contributes once to the anomaly, so $\Delta B = N_{\rm gen}\,n = 3n$. This is a selection rule, not a rate statement, and Part X turns it into absolute proton stability. Here it means only that the conversion happens in units of three baryons at a time — which is irrelevant to a bulk thermal reprocessing and decisive for a single nucleon.

When the bath reaches the electroweak seam near $\mu \sim 100$ GeV, the crossover shuts the channel off, and whatever the sphalerons had converted at that moment is frozen in. *The interval opened the door, the circle counted what passed, and the interval sealed it.*

43 The bookkeeping: what s is, and why every yield is measured against it

Before any factor can be assembled, one has to decide what to count *against.* This section fixes that, because the whole of §44 is bookkeeping in the units set here, and a reader who is unsure what s means will be unsure of everything downstream.

43.1 The problem: everything is diluting at once

The universe is expanding, so every number density falls. Write the comoving scale factor as $a(t)$; a species that is neither created nor destroyed has $n \propto a^{-3}$. That makes a raw number density a bad bookkeeping variable: it changes with time even when nothing physical is happening to the species. Worse, the photon number density does *not* simply scale as a^{-3} either, because whenever a particle species becomes non-relativistic and annihilates, its entropy is dumped into the remaining relativistic bath and *heats* the photons relative to the naive scaling. Between 4 TeV and today, that happens several times — the top quark, the Higgs, the W and Z, the bottom and charm, the QCD crossover, and finally e^+e^- annihilation.

So we need a reference density that (i) also falls as a^{-3} so that ratios are constant, and (ii) already knows about every reheating event along the way. There is exactly one such quantity.

43.2 The answer: entropy density

The entropy density of a relativistic thermal bath is

$$\boxed{s \;=\; \frac{2\pi^2}{45}\, g_{*s}(T)\, T^3} \tag{73}$$

where $g_{*s}(T)$ counts the relativistic degrees of freedom weighted for entropy (bosons with weight 1, fermions with weight 7/8). Entropy is s; it is not a rate, not a probability, not an efficiency. It is *the number of units of disorder per unit volume*, and to within factors of order one it is simply *the number density of relativistic particles* in the bath, counted with their statistical weights.

Its virtue is a conservation law. For adiabatic expansion the total comoving entropy is constant:

$$S \;=\; s\,a^3 \;=\; \text{const} \qquad \Longleftrightarrow \qquad g_{*s}(T)\,T^3 a^3 \;=\; \text{const.} \tag{74}$$

That single equation absorbs every reheating event automatically: when a species annihilates, g_{*s} falls and T rises to compensate, and the product stays fixed.

43.3 Yields

Therefore define, for any conserved-in-comoving-volume quantity X,

$$Y_X \;\equiv\; \frac{n_X}{s}, \tag{75}$$

the *yield*. Because both numerator and denominator fall as a^{-3}, Y_X is a pure number that does not change with expansion. It changes only when physics changes it — when a decay creates asymmetry, when a washout process destroys it, when sphalerons convert one kind into another. *A yield is the thing a Boltzmann equation actually computes.*

This is why the chain is written in yields from beginning to end. We compute Y_{B-L} at the lock at 4 TeV; it is untouched by fourteen orders of magnitude of subsequent expansion; and only at the very last step do we convert it into the quantity astronomers measure, which is a ratio to *photons* rather than to entropy.

In one sentence. s is the entropy density; $Y = n/s$ is a comoving-conserved abundance; the chain computes a Y at 4 TeV, carries it unchanged to recombination, and converts it to η_B at the end.

44 The six factors, one at a time

The predicted asymmetry, Eq. (85) below, is a product of six numbers. This section takes them one at a time: what each one is physically, where its value comes from, and what would change it. None of the six is fitted.

44.1 Factor 1 — how many heavy neutrinos there were to decay: $Y_N^{\rm eq} \simeq 1/g_*$

What it is. The number of heavy Majorana neutrinos per unit entropy, before they decay.

Where the value comes from. At $T \gg M_N$ the heavy states are relativistic and in thermal equilibrium with everything else, so they are simply one more species in the bath. Their equilibrium number density is $n_N^{\rm eq} = (3\zeta(3)/4\pi^2)\, g_N T^3$ with $g_N = 2$ for a Majorana fermion, and dividing by Eq. (73),

$$Y_N^{\rm eq} \;=\; \frac{3\zeta(3)}{4\pi^2}\,\frac{45}{2\pi^2}\,\frac{g_N}{g_{*s}} \;=\; \frac{0.42}{g_*} \qquad \text{per species.} \tag{76}$$

The intuition, which is the whole point. A species in equilibrium carries *its share* of the bath, and its share is its degrees of freedom divided by the total. That is why the answer is $\sim 1/g_*$ and not something requiring a calculation: with about a hundred species in the plasma, any one of them is about one per cent of it. The heavy neutrinos were never rare. What is rare is the *asymmetry* they leave behind.

The honest $\mathcal{O}(1)$. There are three quasi-degenerate heavy states, so the total is $3 \times 0.42/g_* = 1.25/g_*$. Writing $Y_L \simeq \epsilon^{\rm res}\kappa/g_*$, as the chain does, therefore absorbs a factor of 1.25 into the order-unity freedom. We state that rather than let a reader discover it: it is a 25% effect inside a factor-three band, and it points *up*.

44.2 Factor 2 — how biased each decay was: $\epsilon^{\rm res}$

What it is. $\epsilon^{\rm res}$ — the *resonant* CP *asymmetry per decay* — is the fractional excess of leptons over antileptons produced when one heavy neutrino decays:

$$\epsilon_i \;\equiv\; \frac{\Gamma(N_i \to \ell H) - \Gamma(N_i \to \bar{\ell} H^*)}{\Gamma(N_i \to \ell H) + \Gamma(N_i \to \bar{\ell} H^*)}. \tag{77}$$

It is a pure number between -1 and $+1$. If $\epsilon = 0$ the decays are perfectly even-handed and no asymmetry is made at all; if $\epsilon = 10^{-6}$, then for every million heavy neutrinos that decay, there is one extra lepton left over.

Why it is not zero. A decay rate and its *CP*-conjugate can differ only if two conditions hold at once: there must be a *CP*-odd phase in the couplings, and there must be two interfering amplitudes with a relative *absorptive* (on-shell, rescattering) phase. The first is supplied by the geometry — the Majorana parity $\alpha_{31} = \pi$, whose flavored asymmetries are nonzero even though their sum is not (§33.1). The second is supplied by the fact that the heavy states can mix through a loop with an on-shell intermediate state. Neither alone is enough. This is why CP violation in the couplings does not by itself guarantee an asymmetry, and why the value of ϵ is a dynamical question and not merely an algebraic one.

Why it is resonant. The interference between the two nearly-degenerate heavy states is governed by a Breit–Wigner envelope in the mass splitting, Eq. (49): the asymmetry grows as the two states approach each other and saturates at $\epsilon \leq \frac{1}{2}$ when the splitting reaches the width. This is what evades the Davidson–Ibarra bound. It is also, as §38 sets out at length, the one place in the chain where a number is assumed rather than derived.

The value carried. $\epsilon^{\mathrm{res}} \simeq 2.5 \times 10^{-6}$ — roughly one extra lepton per four hundred thousand decays — at derived-conditional grade on H-ϕ_{geo}-1.

44.3 Factor 3 — how much of it survived: κ

What it is. κ is the *washout efficiency*: the fraction of the asymmetry produced by decays that is still there when the dust settles. It runs from $\kappa = 1$ (nothing erased) down to $\kappa \ll 1$ (almost everything erased).

Why there is any erasure at all. Every process that makes an asymmetry can run backwards. If the heavy neutrinos are still being efficiently *re*-created by inverse decays $\ell H \to N$, and if $\Delta L = 2$ scatterings $\ell H \leftrightarrow \bar{\ell} H^*$ are fast, then the plasma is being continually pushed back toward equilibrium — and equilibrium has zero asymmetry. Only asymmetry generated *after* those processes have become slower than the expansion survives.

The parameter that decides it. Compare the heavy neutrino's decay rate with the expansion rate at the moment the decays matter:

$$K_{\mathrm{wash}} \equiv \left.\frac{\Gamma_N}{H}\right|_{T=M_N} = \frac{\widetilde{m}}{m_\star^{\mathrm{wash}}}, \qquad m_\star^{\mathrm{wash}} = \frac{16\pi^{5/2}\sqrt{g_*}}{3\sqrt{5}}\,\frac{\bar{v}^2}{M_{\mathrm{Pl}}} = 1.07 \times 10^{-3}\ \mathrm{eV}. \tag{78}$$

The second equality is worth reading twice, because it is what makes $m_\star^{\mathrm{wash}}$ meaningful rather than mysterious: $m_\star^{\mathrm{wash}}$ *is the value of the effective neutrino mass at which the decay rate exactly equals the Hubble rate.* Above it, decays beat expansion and the system stays in equilibrium; below it, expansion wins and the asymmetry freezes out. It is a competition between two clocks, and $m_\star^{\mathrm{wash}} \simeq 1$ meV is the tie.

Where our value lands. The effective mass the bath sees is the light neutrino scale, $\widetilde{m} \sim m_3 =$ 50.5 meV, so

$$K_{\mathrm{wash}} = \frac{50.5}{1.07} \simeq 47, \qquad \kappa \simeq \frac{1}{K_{\mathrm{wash}}} = 2.1 \times 10^{-2}, \tag{79}$$

and the chain adopts $\kappa = 1.0 \times 10^{-2}$ within the uncomputed order-unity factor. $K_{\mathrm{wash}} \simeq 47 \gg 1$ is the *strong washout* regime: the decays were fast, the plasma spent a long time near equilibrium, and about ninety-eight per cent of the asymmetry that was ever made was subsequently erased. That is not a defect of the mechanism — it is a feature of a universe whose neutrino masses are what they are, and it is why the surviving number is so small.

The two conventions, flagged again because they matter here. The resonance width in §44.2 is evaluated at the seam-matched Yukawa, whose effective mass is $\tilde{m} \simeq 8.4$ meV; the washout in this subsection is evaluated at the light-neutrino mass the bath actually sees, $\tilde{m} \sim m_3 \simeq 50.5$ meV. The two conventions differ by $\simeq 13$ in Y_ν^2, and GG-A10-Cosmos declares, accounts for, and prices the difference. At the seam convention one would instead get $K_{\text{wash}} \simeq 7.9$ and $\kappa \simeq 0.13$.

44.4 Factor 4 — how much of the lepton excess became a baryon excess: $c_{\text{sph}} = 28/79$

What it is. c_{sph} is the *sphaleron transfer coefficient*: the fraction of a conserved $B-L$ asymmetry that sits in the baryons rather than the leptons once the fast processes have reached chemical equilibrium.

Why a fraction, and not all of it. Sphalerons violate $B+L$ and conserve $B-L$. So they cannot change $B-L$, but they are free to *redistribute* it between the baryon and lepton sectors, and they will redistribute it until the chemical potentials balance. Solving the equilibrium conditions — sphaleron equilibrium, Yukawa equilibrium, hypercharge neutrality — for the Standard Model with three generations and one Higgs doublet gives (Harvey and Turner, 1990)

$$B \;=\; \frac{28}{79}\,(B-L), \qquad c_{\text{sph}} = \frac{28}{79} = 0.35443. \tag{80}$$

The intuition. Roughly a third of the conserved quantity ends up as baryons and roughly two-thirds as leptons, because there are more lepton-side degrees of freedom to share it among. The exact rational number is the solution of a small linear system, not a fit; the 79 in the denominator is the trace of that system's structure. Nothing in this framework modifies it: the doublet count is one and the spectrum is the Standard Model's. *We flag it as textbook rather than derived here*, because a chain claiming zero fitted parameters must be explicit about which of its factors are simply inherited from the literature.

44.5 Factor 5 — how many species were sharing the bath: $g_* = 106.75$

What it is. g_* is the effective number of relativistic degrees of freedom in the plasma — the count that appears in Eq. (73). Bosonic states count 1 each, fermionic states 7/8 each.

Where 106.75 *comes from.* Above the electroweak scale the whole Standard Model is relativistic. Counting: 8 gluons $\times 2$ helicities $= 16$; the photon and the three weak bosons before symmetry breaking, $4 \times 2 = 8$; the complex Higgs doublet, 4. Bosons total 28. Then 6 quark flavors $\times 3$ colors $\times 2$ spins $\times 2$ (particle/antiparticle) $= 72$; 3 charged leptons $\times 2 \times 2 = 12$; 3 neutrinos $\times 1 \times 2 = 6$. Fermions total 90, and

$$g_* \;=\; 28 + \tfrac{7}{8} \times 90 \;=\; 28 + 78.75 \;=\; 106.75. \tag{81}$$

Why it divides. It enters twice, and both times for the same reason. It sets how much entropy there is to divide into (§44.1), so the asymmetry per unit entropy is diluted by roughly $1/g_*$; and it sets the Hubble rate through $H \propto \sqrt{g_*}\,T^2/M_{\rm Pl}$, which is why it also appears inside $m_\star^{\rm wash}$. *The physical statement is that the excess was made by one species and then shared out among a hundred.*

44.6 Factor 6 — converting from entropy to photons: 7.04

What it is. Astronomers do not measure baryons per unit entropy; they measure baryons per photon, because photons are what the cosmic microwave background counts. The conversion is the ratio of the two reference densities *today*:

$$\frac{s}{n_\gamma}\bigg|_{T_0} = \frac{(2\pi^2/45)\,g_{*s}(T_0)}{(2\zeta(3)/\pi^2)} = \frac{0.43865\times 3.909}{0.24359} = 7.04, \qquad \eta_B = 7.04\,Y_B. \tag{82}$$

Here $g_{*s}(T_0) = 2 + \frac{7}{8}\times 6\times\frac{4}{11} = 3.909$: the two photon polarizations, plus six neutrino states which are colder than the photons by the factor $(4/11)^{1/3}$ because e^+e^- annihilation heated the photons after the neutrinos had already decoupled.

The intuition. There is more entropy around today than there are photons, by a factor of seven, because the relic neutrinos are still carrying their share of it. Nothing subtle is happening; it is a unit conversion, and it is the last step of the chain.

A number worth keeping. Inverting Eq. (82) gives the entropy per baryon, $s/n_B = 7.04/\eta_B \approx 1.2\times 10^{10}$. Ten billion units of entropy per particle of matter. By particle count the universe is overwhelmingly a photon gas with a trace contamination of baryons — and that enormous number is precisely why nucleosynthesis and recombination happen so "late," far below the relevant binding energies: with a billion photons per baryon there are still enough photons in the high-energy tail to break up deuterium, and later hydrogen, long after the mean photon energy has fallen below threshold.

45 Assembling the number

Now put the six together. The logic runs in four lines, each of which is one physical statement.

Line 1. How many decayed. The heavy neutrinos were in equilibrium, so their abundance per unit entropy was their share of the bath — Eq. (76), $Y_N^{\rm eq} \simeq 0.42/g_*$ per species, hence $\sum_{i=1}^{3} Y_{N_i}^{\rm eq} \simeq 1.25/g_*$ for the quasi-degenerate triplet. The sum deserves display rather than absorption: writing $1/g_*$ for it, as compressed treatments long did, understates the yield by a quarter of the final answer. (The Boltzmann integration of §40 always counted all three species.)

Line 2. How much asymmetry each decay left, and how much survived. Each decay left a fractional excess $\epsilon^{\rm res}$; a fraction κ of the total survived washout. So

$$Y_{B-L} \simeq \epsilon^{\rm res}\,\kappa\sum_i Y_{N_i}^{\rm eq} \simeq \frac{1.25\,\epsilon^{\rm res}\,\kappa}{g_*}. \tag{83}$$

Line 3. The throughput identity. The framework's claim is that the product of the two factors it does not simply inherit is the ledger's own exponential (Part VIII, Eq. (55)):

$$\epsilon^{\rm res}\,\kappa \;\simeq\; (2.5\times10^{-6})\,(1.0\times10^{-2}) \;\simeq\; 2.51\times10^{-8} \;=\; e^{-K/2} \;=\; e^{-17.5}. \tag{84}$$

Line 4. Convert. Sphalerons move 28/79 of it into baryons, and the entropy-to-photon factor turns the yield into the observed ratio:

$$\boxed{\eta_B \;\simeq\; 7.04\,c_{\rm sph}\,\frac{1.25\,\epsilon^{\rm res}\,\kappa}{g_*} \;=\; 7.04\left(\frac{28}{79}\right)\frac{1.25\,e^{-K/2}}{g_*} \;=\; 7.34\times10^{-10}.} \tag{85}$$

Table 14: The chain arithmetic, factor by factor, with the running product. Read down the last column to watch 10^{-8} become 7×10^{-10}. Every entry is recomputed in Appendix 79. Only the first row is the framework's own: steps 2 through 5 are standard thermal bookkeeping that any leptogenesis calculation performs the same way, so the agreement at the foot of the column is a test of the throughput factor and of nothing else.

Step	Factor	Value	Running product
1	throughput $\epsilon^{\rm res}\,\kappa = e^{-K/2}$	2.511×10^{-8}	2.511×10^{-8}
2	$\times$ triplet yield, $\sum_i Y^{\rm eq}_{N_i}\cdot g_*$	1.249	3.137×10^{-8}
3	$\div$ degrees of freedom, g_*	106.75	2.939×10^{-10}
4	$\times$ sphaleron transfer, 28/79	0.35443	1.042×10^{-10}
5	$\times$ entropy-to-photon, s/n_γ	7.04	$\mathbf{7.334\times10^{-10}}$
observed (Planck Collaboration, 2018)			$(6.12\pm0.04)\times10^{-10}$
central deviation			+20%
declared theory band			$\times3^{\pm1}$

What the +20% is, and is not. It is not a precision statement in either direction, and the display convention says so: at headline precision the prediction is $\eta_B\approx7\times10^{-10}$ with the factor-~3 band $(2.4\text{–}22)\times10^{-10}$, and the observation sits comfortably inside. Displaying the equilibrium yield of Eq. (76) matters for honesty in both directions: absorbing it silently would place the central within a few per cent of the observation — a closeness that would be an artifact of the undisplayed factor, which is exactly what an earlier compressed form of this chain did. The genuine claim is narrower and, we think, more interesting: *the independently derived washout efficiency and the candidate regulator's closed form bracket the geometric $e^{-K/2}$ ledger value*, with zero fitted parameters, one named input, and one postulated exponent.

45.1 Which factors could move, and by how much

The value of a prediction is decided by what could have been adjusted to produce it. Here is the complete inventory.

Read that table as the answer to the only question that matters about Eq. (85): *could it have come out otherwise?* Three of the seven rows have no freedom at all. Two have bounded order-unity freedom. And two — the exponent and the resonance — are the joints this paper flags everywhere

Table 15: What could move, and what could not. The only genuine freedom in Eq. (85) is the order-unity factor inside κ — and even that is bounded by $1/K_{\text{wash}}$, which the light neutrino masses fix. Four of the seven rows have no freedom at all, being standard thermal constants; the row that carries the weight is the first, where the exponent $K/2$ is postulated rather than derived, and the table exists to make that visible in one place.

Factor	Fixed by	Freedom
$e^{-K/2}$	$K = LJ = M^2 + N^2 + L^2 = 35$; exponent $K/2$ postulated	**the postulate**
ϵ^{res}	flavored asymmetries on the inherited phases; driver $\alpha_{31} = \pi$	on H-ϕ_{geo}-1
κ	$K_{\text{wash}} = m_3/m_\star^{\text{wash}} \simeq 47$	factor ~ 2
g_*	the Standard Model spectrum above 100 GeV	none
28/79	Harvey–Turner chemical equilibrium, one Higgs doublet	none
7.04	$g_{*s}(T_0) = 3.909$ and $2\zeta(3)/\pi^2$	none
Y_N^{eq}	equilibrium abundance, $0.42/g_*$ per species	the 1.25 absorbed into $\mathcal{O}(1)$

it can. If either is wrong, the answer is wrong by orders of magnitude, not by percent. That is the correct epistemic shape for a prediction: it could have failed loudly, and it did not.

46 The error bar, derived

A prediction without an uncertainty is not a prediction, and in v1 of this study the uncertainty was a declared bracket, $\times 3^{\pm 1}$ — a span of nine, inherited from the parent papers' order-of-magnitude register. This section replaces the bracket with a derivation. The method is the only honest one available: *pin everything that can be computed, scan only what genuinely remains, and let the distribution be the error bar.* The integration of §40 is what makes it possible.

46.1 The budget, before and after

46.2 Three levels, three statements

The result is best stated as three numbers of decreasing assumption and increasing exposure, each with its own error bar. All three are conditional on C4 — the regulator identification — which no amount of Boltzmann-level work can discharge.

Level A: the ledger identity. Unchanged from the parent papers: $\epsilon^{\text{res}}\kappa = e^{-K/2}$ gives $\eta_B = 7.34 \times 10^{-10} \times 3^{\pm 1}$, the exponent postulated. The integration adds one thing to Level A: the single-temperature snapshot of the chain reproduces it (Table 13, last row), so the identity is confirmed as a statement about the snapshot. What the snapshot omits is time.

Level B: the Boltzmann computation. The integration of §40, scanned over the four remaining freedoms:

$$\boxed{\eta_B^{\text{B}} \;=\; 0.86^{+0.52}_{-0.32} \times 10^{-10} \quad (68\%), \qquad \text{band width } \times 2.6} \tag{86}$$

Table 16: The uncertainty budget, before and after. Four of the seven v1 entries are now computed — pinned by the integration or resolved by keeping two physical roles distinct — and no longer contribute at all. What remains is scanned, and every scanned range is a declared freedom of the cohort, not a choice of this study.

Source	v1 status	v2 status	How it was closed
washout convention (seam vs m_3)	$\times 6.0$ priced	**resolved**	the two conventions were two *roles*: the seam Yukawa sets the resonance width, the light masses set the washout; the flavored K_α of Table 12 are unambiguous
κ, order-unity in strong washout	$\times 2.1$ adopted	**computed**	full integration, validated against the analytic fit to 4%
T-window factor	$\div(3.7\text{–}6.2)$ estimated	**computed**	the integration resolves time; the window is Eq. (71)'s business now
three-species abundance	$\times 1.25$ noted	**included**	in the kernels
Breit–Wigner convention	$\times 2$ priced	**pinned**	the integration uses GG-A10-Cosmos's exact R_{res} form; the Lorentzian alternative moves to the omissions inventory
c_τ (condition C5)	$\times 2.1$	scanned, $[5, 10] \times 10^{-3}$	remains a genuine freedom
pattern acceptance band	—	scanned, $\times[2/3,\ 3/2]$	the $\pm\mathcal{O}(1)$ band on ϵ_{eff}
washout-kernel $\mathcal{O}(1)$	—	scanned, $\times[0.7,\ 1.3]$	thermal corrections to the kernel
window edge z_0	—	scanned, $[0.5,\ 2.0]$	the relativistic boundary of validity

against the v1 bracket's width of $\times 9$: *the error bar has contracted by a factor of three and a half*, because most of what used to be bracketed is now computed. The price of the honesty is visible: Level B sits a factor 7.1 [4.4, 11.4] below observation. That deficit is not a new problem — it is C6's known one-sided gap, and the gain is that it is now measured instead of bracketed at 2–20.

Level C: the completed-inventory estimate. Level B deliberately omits four standard transport effects, each named by GG-A10-Cosmos in its enhancement inventory, each pointing up, none specific to GG. Folding them in at their literature ranges — $\Delta L = 1$ scattering source $\times[1.3,\ 2.5]$, Lorentzian-vs-Breit–Wigner convention $\times[1,\ 2]$, relativistic-window residue $\times[1,\ 1.5]$, spectators and the A-matrix $\times[0.85,\ 1.2]$ — gives the best-estimate prediction of this study:

$$\boxed{\eta_B \;=\; 2.8^{+2.0}_{-1.2} \times 10^{-10} \quad (68\%), \qquad 2.8^{+5.0}_{-1.9} \times 10^{-10} \quad (95\%)} \tag{87}$$

against $\eta_B^{\text{obs}} = (6.12 \pm 0.04) \times 10^{-10}$. On the logarithmic scale the observation sits 1.3σ above the Level-C median: *consistent, with a central deficit of a factor* 2.2 that the upper halves of the inventory ranges — or the finite-temperature completion of C4 itself — must supply.

46.3 How to quote it

Three formulations, in descending order of what they assume.

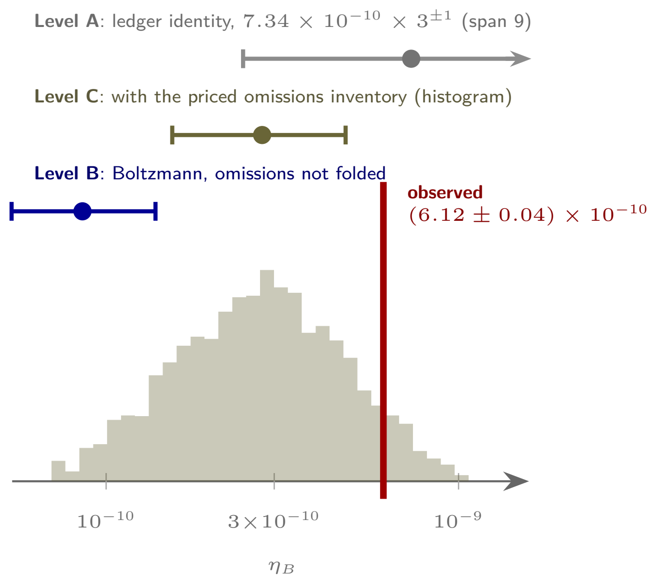

Logarithmic axis. The v1 statement was the wide bracket at the top, its right end running past the frame: a span of nine with the measurement inside it. The v2 statement is the histogram: a computed distribution one third as wide, whose median sits a factor 2.2 below the measurement, with the shortfall assigned to named, standard, one-sided physics. A narrower error bar that exposes a deficit is worth more than a wide one that hides it.

Figure 15: The predicted η_B: three levels and the measurement. Level A assumes the postulated exponent; Level B assumes only the regulator identification (C4) and computes everything else; Level C folds the priced inventory of standard omitted transport effects. All three are conditional on C4.

The ledger-level prediction. $\eta_B \approx 7 \times 10^{-10}$ (central 7.34), band $\times 3^{\pm 1}$, central deviation +20%. Assumes the postulated exponent $K/2$ and H-ϕ_{geo}-1.

The best-estimate prediction of this study (Level C). $\eta_B = 2.8^{+2.0}_{-1.2} \times 10^{-10}$, observation 1.3σ above the median on the log scale. Assumes C4 and the literature ranges of the four named standard omissions; assumes no GG-specific quantity beyond the pinned inputs.

The falsification statement, which assumes least. The measured Boltzmann-level deficit is $E = 7.1$ [4.4, 11.4]. If the finite-temperature program returns the N-width class, E becomes $10^{7\text{–}8}$ and the chain dies; if it returns the flavor-blind $g^2 T$, the peak overshoot returns and the chain dies; if it confirms Γ_τ but the completed transport cannot supply E, the chain dies by throughput. The number to watch is E.

46.4 What would tighten it further

The error bar is now dominated, in roughly equal parts, by c_τ (condition C5) and by the scattering-source range in the inventory. Both are computable: c_τ by a dedicated kinetic calculation of the tau Yukawa equilibration rate at $T \simeq 4$ TeV in an electroweak-symmetric bath, the scattering source by adding the standard $\Delta L = 1$ kernels to the integration of §40. Each would roughly halve its contribution. Neither requires new GG physics. The floor under all of it remains C4: below a band of about $\times 1.5$, no Boltzmann-level refinement matters until the Kadanoff–Baym closure decides whether the regulator is the right one.

47 Why one part in a billion?

Here is the question that motivates the whole exercise, and the answer is one line.

The excess is one part in a billion because half a lock depth is $10^{-7.6}$*, and the universe paid the toll exactly once.*

Compare the alternatives. Electroweak baryogenesis draws the smallness of η_B from a product of small dynamical factors — wall velocity, transport coefficients, the strength of the transition, the size of the new *CP* phase — each of them adjustable, none of them predicted. GUT-scale baryogenesis draws it from washout races that are again adjustable. Standard high-scale leptogenesis draws it from a heavy mass and a phase at energies no experiment reaches, and it is a matter of record that the observed value has never been *predicted* by leptogenesis, only accommodated. Here the magnitude is one exponential of one ledger integer,

$$e^{-K/2} \;=\; e^{-17.5} \;=\; 10^{-7.6}, \tag{88}$$

and the remaining three orders down to 6×10^{-10} are the g_*, the 28/79 and the 7.04 — all of them textbook.

There is no dial in Eq. (85). Ask what could be adjusted to move the answer, and Table 15 is the complete inventory: an order-unity factor inside κ, bounded by a factor of two, and nothing else. K is 35 because $LJ = 35$ and $M^2 + N^2 + L^2 = 35$. g_* is the Standard Model's. 28/79 is Harvey–Turner. 7.04 is thermodynamics. The chain is a prediction in the strict sense: it could have come out at 10^{-6} or 10^{-14} and there would have been nothing to do about it.

48 The one seam we do not close

Honesty requires one more entry, because it is a discrepancy internal to the framework and it would be improper to leave a reader to find it.

There are two routes to the present baryon density. The topological route is the inventory share $\Omega_b = 1/(2\pi^2)$, derived in GG-A10-Cosmos from the volume of the unit three-sphere; the transport route is the chain of this Part. GG-A10-Cosmos converts the ledger share through the baryon bridge $\eta_{10} = 273.9\,\Omega_b h^2$ at its own derived expansion rate and gets $\eta_{10} = 6.15$ against a measured 6.12 — an error of +0.57% with nothing fitted anywhere in the chain — while the transport route lands at 7.34, +19.9%, inside its declared factor-3 band. The two roads to the one number therefore disagree by **16.4**%, which is far outside the CMB's 0.3% determination.

That figure is GG-A9-SM's and GG-A10-Cosmos's, named once per paper across the cohort, and this study adopts it rather than recomputing it. An earlier version of this section quoted 7.5% instead. That number was computed from a superseded assembly of the chain — the one that absorbed the triplet's equilibrium yield $1.25/g_*$ rather than displaying it, and which this study's

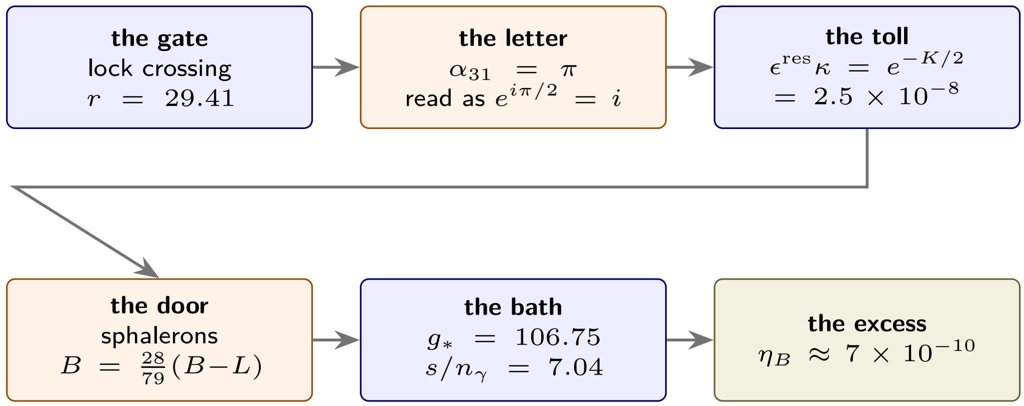

Blue is the interval's contribution — scales, amounts, the departure from equilibrium. Orange is the circle's — orientations, the letter, the direction of the bias. Green is what is left in the sky. The asymmetry is the product of exactly six factors, and every one of them is either a ledger integer, a theorem of the framework, or textbook equilibrium bookkeeping. Nothing was fitted to 6.12×10^{-10}.

Figure 16: The birth of the excess, end to end. Six boxes, one number. Reading the chain in this form makes the division of labor visible: the interval supplies every *amount* and the circle supplies every *orientation*, which is the anatomy's claim about the whole of physics applied to the one case where it produces a measured number.

abstract also carried until now — and it is withdrawn.

That gap is graded **open** and its owner is unidentified. It is not a small residual: it is two derivations of one number disagreeing by twenty-five standard deviations of the measurement. Either the transport is missing an $\mathcal{O}(1)$ — which would be unsurprising given the factor-3 band — or the inventory share and the transport are not, in fact, computing the same thing. We record it here, in the Part that assembles the number, rather than in a technical appendix.

Part X
Absolutely Stable Protons, Absolutely Stable Electrons

Why the same selection rule that let the universe make baryons forbids any of them from decaying, and why the electron is stable for a wholly different reason

49 The apparent contradiction

A reader who has followed Part IX should be uneasy, and the unease is exactly right.

The chain just assembled requires baryon number to be violated. Sakharov says so and there is no way around it: a universe that starts with $B = 0$ and ends with $B \neq 0$ has violated baryon number somewhere. And yet this framework's most distinctive laboratory prediction is that the proton is *absolutely* stable — not long-lived, not 10^{36} years, but $\tau_p = \infty$. How can both be true?

The resolution is a single sentence, and it is the same selection rule doing both jobs.

$$\boxed{\Delta B \;=\; N_{\text{gen}}\, n \;=\; 3n \;\in\; 3\mathbb{Z}.} \tag{89}$$

Baryon number can change — but only in units of three. A bulk thermal reprocessing of a $B - L$ asymmetry does not care: the universe converted its lepton asymmetry three baryons at a time and never noticed. A proton decay, $\Delta B = -1$, does care, because $-1 \notin 3\mathbb{Z}$. No anomalous process, however violent, can remove a single nucleon. The minimal admissible event costs three baryons at once, and three leptons with them, preserving $B - L$.

That is not a loophole. It is the arithmetic of the anomaly: each family contributes once, there are three families, and so the quantum of baryon violation is three. The same $N_g = 3$ that the family pincer fixed stands guard over every atom.

50 The theorem

Theorem 50.1 (T-SM-60.1: boundary proton stability, conditional; carried from GG-A9-SM)**.** *Assume (H1) the boundary spectrum is the split Standard Model with the X, Y retained only as bulk modes; (H2) the bulk $\mathrm{U}(1)_B$ descends so that the exact boundary gauge symmetry acts as baryon number B (charge $+\frac{1}{3}$ on quark fields, 0 on leptons); and (H3) the $\mathrm{U}(1)_B$ gauge anomalies are canceled by the BI-6 Green–Schwarz inflow. Then every local operator with $\Delta B = \pm 1$ or ± 2 carries non-zero $\mathrm{U}(1)_B$ charge and is forbidden to all orders. The only admissible baryon-number violation is the $\mathrm{SU}(2)_L$ electroweak anomaly, $\Delta B = 3n \in 3\mathbb{Z}$, suppressed by $e^{-4\pi/\alpha_W} = e^{-372} \sim 10^{-162}$. Hence the proton ($\Delta B = 1$ decay) is absolutely stable and n–$\bar{n}$ oscillation ($\Delta B = 2$) is forbidden; only $\Delta B \in 3\mathbb{Z}$ processes are even in principle admissible, at unobservable rates.*

Proof sketch. Under (H2) baryon number is an exact local symmetry of the boundary action, so any operator of non-zero B-charge has identically vanishing Wilson coefficient; the proton-decay operators ($\Delta B = 1$) and the n–$\bar{n}$ operators ($\Delta B = 2$) are such operators. Under (H1) none is generated by four-dimensional X, Y exchange in the first place. The residual non-conservation is the anomaly of Eq. (72), which changes B only in multiples of $N_{\text{gen}} = 3$ and cannot deplete a single nucleon. ∎

50.1 The subtle step, and why it is not automatic

It would be easy to read Theorem 50.1 as "gauge B and you are done," and that reading is wrong. The point is worth spelling out because it is where the theorem earns its keep.

Baryon number is *not* a generator of unbroken SU(5). On the **10** the components carry $B = \{+\frac{1}{3}, -\frac{1}{3}, 0\}$, which no combination $\beta\mathbf{1} + aY$ can reproduce — three conditions overdetermining two coefficients. Equivalently, the Abelian factor that *commutes* with SU(5) assigns a single charge to each multiplet, and the dimension-six operator $QQQL$ generated by X, Y exchange, being $\text{SU}(5) \times \text{U}(1)_B$ invariant by construction, is *neutral* under that per-multiplet charge. Gauging "the $\text{U}(1)_B$ of $\text{SU}(5) \times \text{U}(1)_B$" therefore does not by itself forbid proton decay.

What forbids it is *genuine, component-resolved* baryon number, which is well defined only on the split Standard-Model spectrum where Q, u^c, e^c are independent fields free to carry $B = +\frac{1}{3}, -\frac{1}{3}, 0$. In GG-6 that is exactly the boundary situation: the X, Y of the adjoint **24** are modes of the six-dimensional bulk and are absent from the four-dimensional boundary spectrum, so there is no boundary X, Y exchange; and the bulk $\text{U}(1)_B$ descends, through the Wilson–Hosotani holonomy and the projection, to act on the split fields as baryon number itself.

> **The grade, carried verbatim: "conditional" is in the theorem's own title.** (H1) is the standard bulk–boundary statement and (H3) is the Green–Schwarz bookkeeping of GG-A3-BI6. The one load-bearing step still to be discharged is **(H2)** — an explicit computation that the holonomy descent realizes the bulk $\text{U}(1)_B$ as *baryon number* rather than the harmless per-multiplet charge, defeating the rank-matching obstruction that $\text{SU}(5) \to \text{SM}$ is rank-preserving. GG-A9-SM states that this is finite and self-contained — a projection-plus-charge computation, not an open conjecture — and defers it to a dedicated successor paper. This study carries the theorem with its conditional attached and does not quote $\tau_p = \infty$ without it.

Two corollaries are worth recording. First, because the X, Y are bulk-only, gauge-coupling unification here is *geometric*, carried by radial transport along the infrared region, and not mediated by four-dimensional X, Y running: proton stability and the absence of four-dimensional X, Y are two faces of one statement. Second, the protective symmetry is the *global* baryon number residual to a $\text{U}(1)_B$ terminated at the seam — its vector massive at M_{lock} — not a long-range gauge force. The analogy is controlled: give the photon a mass and electric charge is still exactly conserved — a boson mass gaps the force, not the bookkeeping. The massive Z'_B is itself the falsifiable trace, and it is one of the three states of the lock cluster.

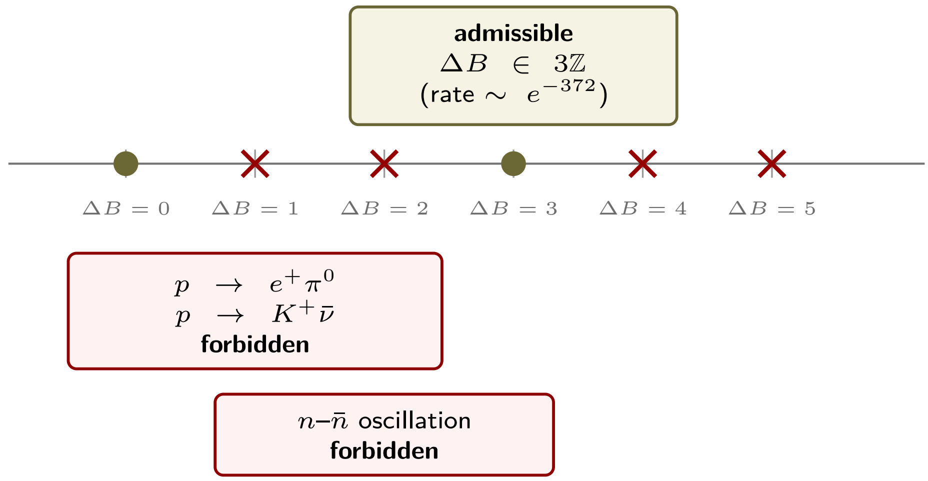

Every canonical proton-decay mode has $\Delta B = 1$; neutron– antineutron oscillation has $\Delta B = 2$; neither integer is a multiple of three, so both are forbidden outright rather than merely suppressed. The thermal reprocessing of Part IX lives at $\Delta B \in 3\mathbb{Z}$ and is unaffected. **A single event in either forbidden class falsifies the framework** — and Hyper- Kamiokande will improve the bound on the first class by about an order of magnitude.

Figure 17: The selection rule that does both jobs. Baryon number changes only in units of the generation count. Bulk reprocessing is permitted, single-nucleon decay is not, and the two statements are the same statement. This is the sharpest laboratory difference between this framework and ordinary grand unification, and it points the opposite way: where a GUT is *confirmed* by seeing proton decay, this framework is *refuted* by it.

51 A prediction that points the wrong way on purpose

It is worth being explicit about the epistemic shape of this prediction, because it is unusual.

Grand unified theories *predict* proton decay in the range 10^{34}–10^{36} years, and they are confirmed by a positive detection. Decades of null results at Super-Kamiokande have pushed $\tau/B(p \to e^+\pi^0)$ above 2.4×10^{34} years and $\tau/B(p \to \bar{\nu}K^+)$ above 5.9×10^{33} years, which has squeezed the simplest models without excluding the idea.

This framework says the search will never succeed, in any channel, at any sensitivity. That is a stronger claim and a more dangerous one. It cannot be confirmed — no finite exposure establishes $\tau = \infty$ — and it can be destroyed by a single event. Hyper-Kamiokande's projected sensitivity of 10^{35} years in $p \to e^+\pi^0$ and 3×10^{34} in $p \to \bar{\nu}K^+$ over twenty years of running is therefore not a test the framework can pass; it is a test it can only fail. We think that is the right way round for a framework that claims to fix parameters rather than fit them, and we would rather stake the whole baryonic logic on it than hedge.

52 The electron: stability as an absence of destinations

The electron is also absolutely stable, and for a reason that has nothing to do with the proton's.

Theorem 52.1 (No-exit-gate (T-Electron-2); carried from GG-P3-Electron, GG-P2-N3)**.** *The electron basin is the unique global minimum of the charged-lepton stratum (the GGA). Unlike the depth-*0 *and depth-*1 *basins, it possesses no exit gate: there exists no admissible GSSB gate, of any index and any finite Bogomolnyi toll, leading out of it. GRIP-In halts there. Consequently the electron cannot decay — not because every channel is blocked by a conservation law applied to an unexplained spectrum, but because the landscape contains no downhill direction and no gate.*

The mechanism is Morse-theoretic rather than symmetry-protected, and the difference matters. The charged-lepton cascade is a single descent,

$$\underbrace{[\,\mathrm{ind}=4\,]}_{\tau,\ \text{depth }0} \xrightarrow{S=2\pi} \underbrace{[\,\mathrm{ind}=2\,]}_{\mu,\ \text{depth }1} \xrightarrow{S=\pi} \underbrace{[\,\mathrm{ind}=0\,]}_{e,\ \text{depth }2\ (\mathsf{GGA})}\ , \tag{90}$$

with Bogomolnyi actions $S_{[2k]\to[2(k-1)]} = \pi k$ and total $S_{\mathrm{tot}} = \pi N_g(N_g-1)/2 = 3\pi$, exactly saturating the BI-6 instanton budget $S^{\mathrm{BI}}_{\max} = 3\pi$ at $N = 3$. The weak interaction is the physical implementer of the descent — it is the unique flavor-connecting channel of the boundary gauge group — and it still acts on the electron. It simply has nowhere to take it.

Charge conservation is not assumed on top. The boundary $\mathrm{U}(1)_{\mathrm{EM}}$ descends from the derived bulk gauge factor, so the standard two-line argument survives *inside* the framework as a consistency check rather than as the explanation. The explanation is that there is no lower level to radiate to: *stability is an absence of destinations.*

The empirical anchor is $\tau_e > 6.6 \times 10^{28}$ years at 90% confidence, and a charge-violating electron decay in any channel would falsify Theorem 52.1 outright.

The grade, again carried unchanged. Theorem 52.1 is a theorem *within* GDOS, presupposing GFF-ST7 and the cascade machinery; it is not claimed to be derivable in Standard-Model language, and it inherits GG-P2-N3's declared conditions on the budget-saturation argument that stops the ladder at three generations. GG-P3-Electron says of the reachability reading that it "inherits the conjectural status of the budget and adds none of its own rigor." We say the same.

52.1 The symmetry of the two arguments

Put the two side by side, because the pair is more informative than either alone.

The two together say something a framework rarely gets to say. The stuff we are made of is not merely long-lived. It is permanent — the protons because the arithmetic of the anomaly will not let them go one at a time, the electrons because there is nowhere below them to go. A universe that made a small excess of matter at four TeV, and then could never lose any of it, is a universe with a beginning for its matter and no end.

Table 17: Two absolute stabilities, two different mechanisms. That the framework reaches the same verdict twice by unrelated routes — one a selection rule from an anomaly, one a Morse-theoretic absence of destinations — is the kind of redundancy that is worth noticing, and the kind that can fail twice independently.

	Proton	Electron
Mechanism	selection rule $\Delta B \in 3\mathbb{Z}$	no exit gate from the ground basin
Type	symmetry-protected	Morse-theoretic (ground-state-ness)
Owner	GG-A9-SM T-SM-60.1	GG-P3-Electron, GG-P2-N3
Rests on	exact global B; the anomaly's quantum of three	$S_{\text{tot}} = 3\pi$ saturating the BI-6 budget
Conditional on	(H2), the baryon-descent computation	GFF-ST7 and GG-P2-N3's budget conditions
Prediction	$\tau_p = \infty$; n–$\bar{n}$ forbidden	$\tau_e = \infty$
Current bound	$\tau/B > 2.4 \times 10^{34}$ yr ($p \to e^+\pi^0$)	$\tau_e > 6.6 \times 10^{28}$ yr
Killed by	one event, any channel	one event, any channel

Part XI
What Is New Here, and What Is Not

An honest inventory of the claim's novelty, against the literature and against the cohort's own earlier papers

53 Two claims that appear to be without precedent

Novelty is easy to assert and easy to overstate, so we separate the two claims we believe are genuinely new from the many that are not, and we say for each exactly what would have to be true for the claim of novelty to be wrong.

53.1 First: Majorana leptogenesis at a scale of a few TeV, with the scale not chosen

Leptogenesis is forty years old (Fukugita and Yanagida, 1986). Resonant leptogenesis is thirty (Pilaftsis, 1997). TeV-scale resonant leptogenesis has been studied extensively (Pilaftsis and Underwood, 2004). So the *mechanism* is not new, and we do not claim it is.

What appears to be new is the *provenance of the scale*. In every treatment we are aware of, the TeV scale is chosen — selected precisely *because* it is accessible, with the required quasi-degeneracy of the heavy states imposed by hand or arranged by an added flavor symmetry introduced for the purpose. The scale is an input motivated by testability, which is a perfectly respectable motivation and also a transparent one.

Here the scale is

$$M_{\text{lock}} \;=\; \tfrac{1}{2}\, M_{\text{Pl}}\, e^{-K} \;=\; \tfrac{1}{2}\, M_{\text{Pl}}\, e^{-35} \;=\; 3.85\ \text{TeV}, \tag{91}$$

where $K = LJ = M^2 + N^2 + L^2 = 35$ was fixed by the anomaly audit and the Green–Schwarz consistency of the six-dimensional bulk, and where the same K independently sets the Kaxion's frequency window, the vacuum terminus $m_\Lambda = M_{\text{lock}}^2/M_{\text{Pl}}^{\text{red}} \simeq 6.1$ meV, and the length of the descent. No step in the derivation of Eq. (91) mentions baryogenesis, neutrinos, or *CP*. The mechanism was placed at four TeV by a paper that was computing something else.

That is a different kind of claim from "we choose the TeV scale so that the LHC can test it," and it has a different consequence: the scale is not adjustable if the number comes out wrong. The lock cluster is either near 3.85 TeV or the chain is dead, and HL-LHC will settle it.

What would refute the novelty claim. A published treatment in which a heavy Majorana scale of a few TeV is *derived* from structure fixed by considerations independent of baryogenesis, and the quasi-degeneracy follows from the same structure rather than being imposed. We have not found one. We would be glad to be shown one, and the claim in this section should be read as a claim about the literature we have searched, not as a theorem.

53.2 Second: a quantitative baryon-to-photon ratio from an integer

The second claim is sharper and easier to check.

It is a matter of record in the field that leptogenesis has never *predicted* η_B. It accommodates it. The heavy masses, the Yukawa structure encoded in the Casas–Ibarra R matrix (Casas and Ibarra, 2001), and the high-scale phases are all free, and between them they can produce essentially any asymmetry one wants. The standard defense of the mechanism has therefore always been qualitative — that it *can* work, at the right order of magnitude, over a plausible region of parameter space — and the standard criticism of it has been the corresponding one: that a mechanism with enough parameters to fit anything has not explained the thing it fits.

Equation (85) is a different sort of object:

$$\eta_B \;=\; 7.04 \times \frac{28}{79} \times \frac{1.25\, e^{-K/2}}{g_*} \;=\; 7.34 \times 10^{-10} \qquad \text{versus} \qquad (6.12 \pm 0.04) \times 10^{-10}. \tag{92}$$

The dial count is zero. The heavy scale is fixed, the Yukawa is fixed by the seam, the washout is fixed by the light spectrum, the phases are lattice points, the transfer coefficient is Harvey–Turner, the entropy factor is thermodynamics, and g_* is the Standard Model's. What remains is one integer.

We think the honest formulation of the novelty is this: *so far as we can establish, no previous treatment has obtained the baryon-to-photon ratio as a closed-form consequence of a small integer with no fitted quantity anywhere in the chain.* That formulation is deliberately narrower than "we have explained the matter asymmetry," and it is deliberately exposed — the two joints where it could fail are named in the same breath, and they are named in §54.

53.3 Third, and less often noticed: the compatibility

There is a third thing worth putting on the record, less headline than the first two but structurally more satisfying.

A framework that makes baryons must violate baryon number. A framework that forbids proton decay must conserve it. Almost every attempt to do both ends up with two mechanisms and a coincidence between them. Here there is one selection rule, $\Delta B \in 3\mathbb{Z}$, and it does both jobs because a thermal bath does not care about a factor of three and a single nucleon does. We are not aware of another framework in which baryogenesis and absolute proton stability are consequences of the *same* statement rather than of two compatible ones.

54 The two joints where it can fail

If the previous section is to be read as a claim rather than as advertising, this section has to be read with it.

Joint one: the exponent. Why $K/2$ and not K or $K/4$? The compression statement is postulated.

Table 18: The honest ledger of the whole chain. Everything in the "Derived" column is forced; everything in the last two rows is not, and no reader should take the +20% agreement as evidence about the last two rows. The five registers are ordered by how much they are worth, strongest first, and the ordering is the study's own verdict on itself: the further down the table a statement sits, the less the agreement at the end of the chain says about it.

Register	Content
Derived	the arena and its two algebras; the ledger $(1,3,5;7,35)$; the flavor connection and the identification of loop holonomies as the physical data; the Majorana door (weight-zero theorem); the uniqueness of the LHN portal; the lock scale; the seesaw-matched Yukawa; $\Delta B \in 3\mathbb{Z}$; the sphaleron transfer and entropy factors (standard)
Derived-conditional	the pentagonal lattice (rigidity of the holonomy moduli); δ_{CKM} at first order; δ_{PMNS} (leptonic assignment lemma); the light spectrum and $m_\star$ (Weinberg-coefficient lemma); the washout; ϵ^{res} (on H-ϕ_{geo}-1); proton stability (on (H2))
Asserted / flagged	the eight charged-fermion cascade depths; the QD completion $\delta_{\mathrm{CKM}} = 13\pi/35$ (parent flags it post-hoc); $\alpha_{21} = 0$ as a convention
Named input	the on-resonance condition, C-P5-C(i) — the quantity the chain requires and does not derive
Postulate	the throughput exponent $K/2$, C-P5-C(ii) — geometrically suggestive as $\sqrt{e^{-K}}$, not derived

If it is wrong, the number is wrong by many orders of magnitude and the agreement in Eq. (85) is a coincidence. The thing that makes this a scientific position rather than a fitted one is that Eq. (55) shows an *independent* chain — seesaw Yukawa from the seam, washout from the light spectrum — landing on the same value. Two routes to 2.5×10^{-8} is better than one, and it is a great deal better than asserting it. But it is not a derivation, and we do not call it one.

Joint two: the resonance. The on-resonance condition is a named input, sharpened to H-ϕ_{geo}-1 with an acceptance band and bracketed from both sides, with the crux at C4 and three routes to refutation from within. Section 38 gave it a full treatment because it deserves one.

A chain with one postulate and one named input is a good chain. A chain with one postulate and one named input *that a reader has to hunt for* is a bad paper. That is why both appear in the Abstract, in the Executive Summary, in Table 9, in Table 18, and in a Part of their own.

55 What is inherited, and from where — the cohort's own history

A note on internal provenance, because it would be easy for this study to appear to claim its parents' work, and because the division of labor between the two parents is finer than a casual reading suggests.

55.1 GG-A9-SM: the flavor connection and the door

GG-A9-SM built the flavor connection and proved, as T-SM-46.2, that the physical content of a Yukawa texture is its loop angles — including clause (iv), the pentagonal lattice. It derived the branch phase $3\pi/5$ from the Higgs doublet's retention of $2/L$ and $\delta^{(1)}_{\rm CKM} = 2\pi/5$ from it. It proved the Majorana door as T-SM-44.1 in the "if and only if weight zero" form. It wrote the geometric seesaw $m_\star = \pi v^2/E_{\rm GUT}$ and the light spectrum with normal ordering as an *output* of the depth assignment. It derived $\delta_{\rm PMNS} = \pi + 2\arcsin(1/J)$ and the Majorana pair $(0, \pi)$. And it proved proton stability as T-SM-60.1 with its (H2) conditional attached.

It then delegated the heavy sector explicitly, in one sentence worth quoting because it fixes the boundary: *the heavy-sector realization and its leptogenesis are carried, in completed form,* by the cosmology companion.

55.2 GG-A10-Cosmos: the transport, the number, *and* the mechanism

This is the part most easily got wrong, and this study got it wrong at first.

GG-A10-Cosmos does **not** import η_B. It derives it. Part V (§"Baryogenesis via the inherited neutrino and *CP* ledger", Proposition T-Cosmos-8) and Appendix I (§"Resonant leptogenesis and the lithium-7 port") together run to some twelve hundred lines and contain the Boltzmann system in $z = M_N/T$, the heavy triplet $M_{N_i} \sim M_{\rm lock}(1, 1+K^{-1}, 1+2K^{-1})$, the seesaw-matched Yukawa $Y_\nu \simeq 7.3 \times 10^{-7}$, the resonant asymmetry $\epsilon^{\rm res} \simeq 2.5 \times 10^{-6}$, the washout $\kappa \simeq 10^{-2}$ from $K_{\rm wash} \simeq 50$, the throughput identity $\epsilon^{\rm res}\kappa \simeq 2.51 \times 10^{-8} = e^{-K/2}$, the assembled $\eta_B = 7.34 \times 10^{-10}$ at $+20\%$, the codebook row D17 at grade D*, and the declared factor-~ 3 theory band. Appendix I describes its own function without ambiguity: *it supplies the one number standard Big-Bang nucleosynthesis had to take as a measured input.*

GG-A10-Cosmos is also the home of the sharpening. Appendix I.3 — the section titled "The partial-resonance mechanism: the flavor-measurement regulator" — contains the flavor-covariant kinetic equation, the proof that a real R makes the source flavor-odd, the proof that flavor-blind damping cancels from a flavor-odd source, the two-level reduction and its Lorentzian envelope, the closed form $\Delta M/\Gamma^{\rm eff} = 1/(35 c_\tau y_\tau^2)$, the worked number $\Gamma^{\rm eff} = 2.1$–4.1 MeV, the regulator landscape of Table 10, and the six conditions C1–C6 with C4 named as the crux.

55.3 GG-P5-CPT: the hypothesis, the grading, the roadmap

What GG-P5-CPT owns is not the mechanism but its *status*. It states H-$\phi_{\rm geo}$-1 in the sharp two-sided form quoted in §38; it fixes the acceptance band at $\pm\mathcal{O}(1)$ in $\epsilon_{\rm eff}$; it carries the parent conjecture C-P5-C with its two clauses (the on-resonance condition and the $e^{-K/2}$ throughput) separated so that they can fail independently; it records the whole item as row 9 of its promotion roadmap, with the closing condition and the impact of closing it spelled out; it withdraws the former "m_1/m_3-order suppression" language as a reconstruction artifact; and it states the three kill routes and the success/failure criteria. It is also the paper that turned the pentagon from a clause of T-SM-46.2 into a Part with its rigidity hypothesis named, and that located the sense in which $\alpha_{31} = \pi$ drives the excess — *CP*-conserving in the flavor sum, *CP*-violating in the flavor differences.

A correction this study makes to its own first reading. An earlier summary of this material attributed the sharpening of H-ϕ_{geo}-1 to GG-P5-CPT alone. That is wrong. The derivation home of the flavor-measurement regulator, the closed form, the landscape table and the conditions C1–C6 is **GG-A10-Cosmos Appendix I.3**; GG-P5-CPT carries them, grades them, and equips them with the hypothesis statement, the acceptance band, the roadmap row and the kill routes. The two papers cross-reference each other explicitly, and the work was done together. Stating it otherwise would have credited one parent with the other's calculation.

55.4 GG-W2-Holonomy: the anatomy

GG-W2-Holonomy supplied the statement of which fiber does which job: that the interval carries no holonomy and can therefore carry no phase; that the circle carries all of it; that *CP* breaking is a property of the content pressed into the circle rather than of the connection, which is why its parameters neither run nor drift; and that the preferred sense of reading is manufactured jointly by chiral content and a one-way read head, with neither ingredient itself asymmetric. It also located the leptogenesis argument inside the anatomy: the neutrino sector is the pivot not by preference but by intersection.

The accurate statement of what came later. It is *not* the derivation of η_B that is the recent development — that is GG-A10-Cosmos's, in completed form, and it has been there since Part V and Appendix I were written. What is recent is the **sharpening of the derivation's one named hypothesis**: the move from "assume the system is on resonance" to a named partial-resonance condition with a number, a two-sided bracket, an identified physical mechanism, a closed form in the ledger integer, a landscape of alternatives spanning sixteen orders with exactly one candidate in the required window, six conditions with the crux isolated, and three routes to refutation from within. *The chain has not become more certain. It has become more falsifiable* — and that is the direction in which a framework of this kind ought to move.

56 What this study adds

Four things, none of them a derivation.

The assembly. The argument end to end in one document, from the topology of the sixth dimension to a number in the sky, with the intuitive layer written at the length the physics deserves rather than compressed into a companion's cross-reference.

The Boltzmann integration. GG-A10-Cosmos closes its washout section with "the $\mathcal{O}(1)$ prefactor . . . awaits a full Boltzmann integration." Sections 39–40 perform it: flavored washouts computed from the derived spectrum and mixing matrix (reproducing GG-A10-Cosmos's own values), kernels validated against the standard analytic efficiency to four per cent, the sign of the excess emerging as an output through the electron door, the window freedom and the convention freedom converted into computed numbers, and the error bar of Part IX derived rather than declared. It is carried at candidate grade, conditional on C4, and upgrades nothing.

The helix, and its refutation. Part IV builds the picture that a reader's intuition most wants — a phase swept out along the descent — makes it quantitative for both sectors, and shows that it fails by a factor of 2.73. This is a small result and it is this study's own. Its value is that it converts a topological theorem (the interval carries no holonomy) into a number a skeptic can check.

The confrontation with the 2026 experimental program. Part XII puts every prediction of the chain against the current status and the projected reach of the instruments that will decide it, including one place where the framework is already in tension and one where a widely repeated expectation about the experimental landscape has changed.

Two corrections on the record. The lattice *step* $\pi/5$ and the first-order phase $2\pi/5$ are different quantities and were once conflated in discussion; and the helix is a location, not a mechanism. Both are stated in the text where they arise rather than buried.

Part XII
The Falsification Program, 2026–2040

Every prediction of the chain, against the instrument that will decide it, with the current status as of August 2026

57 How this Part is organized

The chain assembled in Parts VII to IX makes seven predictions that experiments will decide within fifteen years. For each we give the prediction, the current best measurement or bound, the instrument that will improve it, the projected reach, the date, and — most importantly — what specifically would count as a kill.

We have taken some care to state the current experimental landscape accurately rather than from memory, because two widely repeated expectations have changed recently and a preprint that repeated them would be wrong on the record. Both are flagged in place.

One general remark first. The predictions divide into two classes that should not be confused. *Two-sided* predictions can fail in either direction: δ_{PMNS}, Σm_ν, the lock cluster. *One-sided* predictions can only be killed, never confirmed: proton stability, electron stability, and — for a reason we explain in §60 — $m_{\beta\beta}$. A framework whose most distinctive claims are one-sided is in an uncomfortable position, and we would rather name that discomfort than dress it up.

58 δ_{PMNS}: DUNE and Hyper-Kamiokande

The prediction. $\delta_{\text{PMNS}} = \pi + 2\arcsin(1/7) = 196.43°$, quantum-defect layer 196.06°, with the $3\pi/2$ assignment branch *pre-registered as dead* (§32).

Where it stands in August 2026. The joint T2K–NOvA analysis (T2K and NOvA Collaborations, 2025) gives, for normal ordering, a best fit of $\delta_{CP} = -0.87\pi = 203°$ with a 68% interval of [166°, 241°]. The NuFIT 6.0 global fit (Esteban et al., 2024) gives $212°\,^{+26}_{-41}$ including the Super-Kamiokande atmospheric sample and $177° \pm 19°$ without it. The lattice point 196.4° lies inside the one-sigma region of all three.

What must not be claimed. The best fits sit near the prediction; that is all. In normal ordering, $\delta_{CP} = 180°$ — the CP-conserving value — also lies inside the current intervals. **Leptonic *CP* violation has not been established.** A reader who takes the 196° agreement as a success should note that the same data are consistent with no leptonic CP violation at all. Under inverted ordering the lattice point is already disfavored at about three sigma — but the joint analysis reports that if inverted ordering were true, its own data would constitute evidence for *CP* violation, and it weakly prefers inverted ordering. The mass ordering is contested (§59), and until it settles, so does this.

The instruments. DUNE's far-site cavern excavation is complete and detector steel began moving underground in May 2026; Fermilab's stated priority is first neutrino beam by 2031 (DUNE Collaboration, 2020, 2025). Its current official sensitivity projections are a δ_{CP} resolution of 7° at $\delta_{CP} = 0$ and 18° at $-\pi/2$ for an exposure of 1000 kt·MW·yr, roughly fourteen staged years; 5σ exclusion of *CP* conservation for half of all true δ_{CP} values at 600 kt·MW·yr. Hyper-Kamiokande's main cavern was completed in July 2025 and operation is planned from 2028 (Hyper-Kamiokande Collaboration, 2025); ten years at 27×10^{21} protons on target gives a δ_{CP} precision running from 20° at maximal *CP* violation to 6° at *CP* conservation, with 5σ exclusion of *CP* conservation for more than 60% of true values.

The kill. A converged measurement inconsistent with 196° beyond the errors — concretely, data settling at maximal 270° at 5σ in the window 2031–2035 — retires the leptonic lattice chain and, with it, the door's bias. The pre-registration matters here: a measurement near 270–286° may not be reinterpreted under the dead $3\pi/2$ branch.

And the quiet signature. There is a second way this prediction could be confirmed, and it is worth naming because it is the opposite of a discovery. If δ_{PMNS} really sits at 196.4°, then $\sin\delta_{\mathrm{PMNS}} = -0.28$ — and the long-baseline program was designed to find a large *CP*-violating asymmetry. The prediction is therefore confirmed either by a measured phase crowding 196° at growing precision, or — the same fact seen from the community's side — by the persistent *non-arrival* of the five-sigma discovery the program was built to deliver.

59 The mass ordering

The prediction. Normal ordering, as an *output* of the depth assignment: ν_1 is the ground state of the Majorana stratum, so $m_1 < m_2 < m_3$ is forced. There is no fallback spectrum.

Where it stands. Contested, and this is a live disagreement in the field rather than a settled question. The joint T2K–NOvA analysis weakly prefers *inverted* ordering, with a Bayes factor of 1.3 including the reactor θ_{13} constraint and 2.5 without it. The NuFIT group's post-JUNO update prefers *normal* ordering at 2.2–2.3σ. JUNO's first data (JUNO Collaboration, 2026), from 59 days, already improve the solar parameters by a factor of 1.6 over all previous measurements combined, and JUNO's own ordering determination is the near-term arbiter.

The kill. An established inverted ordering at 5σ, expected achievable around 2033, retires the paper. GG-P5-CPT puts it without softening: under inverted ordering at five sigma, the paper retires, with no fallback spectrum.

60 Neutrinoless double beta decay: the Majorana door itself

The prediction. $m_{\beta\beta} = 1.72$ meV at first order, 1.79 meV at the quantum-defect layer, from Eq. (45) with every ingredient fixed. And, underneath it, the qualitative claim on which the whole matter chain rests: **neutrinos are Majorana particles.**

Where it stands. The world-leading limit is KamLAND-Zen's complete dataset (KamLAND-Zen Collaboration, 2025): $T_{1/2} > 3.8 \times 10^{26}$ yr in ^{136}Xe, corresponding to $m_{\beta\beta} < 28$–122 meV depending on the nuclear matrix element. LEGEND-200's first result (LEGEND Collaboration, 2026) gives $T_{1/2} > 1.9 \times 10^{26}$ yr in ^{76}Ge combined with GERDA and MAJORANA, and CUORE's two-tonne-year dataset (CUORE Collaboration, 2025) gives $T_{1/2} > 3.5 \times 10^{25}$ yr in ^{130}Te.

The instruments. The tonne-scale generation targets $m_{\beta\beta} \simeq 9$–20 meV: LEGEND-1000 at 1.4×10^{28} yr over a ten-year run, nEXO at 1.35×10^{28} yr (nEXO Collaboration, 2022), CUPID at 1.8×10^{27} yr with staged operation from 2030 and full scale from 2034, AMoRE-II at 6×10^{26} yr, and SNO+ Phase II at 7×10^{26} yr from late 2028.

A correction to a widely repeated expectation. It is often said that the tonne-scale program will "cover the normal ordering." It will not. Covering the *inverted* ordering, whose minimum is $m_{\beta\beta} \gtrsim 16$–$22$ meV, is within reach. The normal-ordering band runs from about 1.3–1.8 meV at its minimum to about 3.5–4.0 meV at its maximum for a vanishing lightest mass (Ge et al., 2026), and it contains a cancellation funnel in which $m_{\beta\beta} \to 0$. *No approved experiment reaches it.* A preprint that claims the tonne-scale program will test its normal-ordering prediction is claiming something untrue, and we are not going to claim it.

The kill, and its asymmetry. This prediction is one-sided. $m_{\beta\beta} = 1.8$ meV sits an order of magnitude below the tonne-scale reach, so a null result at 9–20 meV is consistent with it and proves nothing. What kills it is a *discovery*: a confirmed $0\nu\beta\beta$ line anywhere in the tonne-scale window would be incompatible with the derived spectrum and phases. A line at 50 meV is as fatal as an absence at 1.8 meV would be unprovable.

That asymmetry is uncomfortable and we state it rather than hide it. Note what it does *not* affect: the qualitative claim. If the neutrinos turn out to be Dirac particles — established by any route — then the driver α_{31} does not exist and the matter chain of Part VIII retires in its present form. That is a clean kill and it does not require reaching 1.8 meV.

61 The neutrino mass sum: the tightest live tension

This is the place where the framework is under the most pressure right now, and it deserves the longest treatment.

The prediction. $\Sigma m_\nu = 59.17$ meV, from Eq. (37). This sits essentially *at* the normal-ordering floor of $\simeq 58.8$ meV. It cannot be lowered, because there is nothing to lower: the three masses are fixed by $m_\star$, the family projector, and the quantum-defect layer.

Where it stands. Cosmological bounds have tightened to the point of collision.

The DESI collaboration quantifies the situation as a 3σ tension between its constraint and the lower limit from oscillations, using an effective parameter $\Sigma m_\nu^{\rm eff}$ that is allowed to go negative, and notes that the posterior peaks in the negative sector (DESI Collaboration, 2025a, Elbers et al., 2025). That is not a comfortable thing for this framework to read.

Table 19: Cosmological bounds on the neutrino mass sum as of August 2026, against the derived $\Sigma m_\nu =$ 59.2 meV and the normal-ordering floor of 58.8 meV. Two of the four bounds already sit *below* the prediction. The Model column is why the bounds disagree with each other, and it is also why the kill criterion is stated under the framework's own $w = -1$ prior: the one bound that sits comfortably above the prediction is the one that lets dark energy evolve, which this framework does not.

Dataset	Model	Σm_ν (95%)	versus 59.2 meV
DESI DR2 BAO + CMB	ΛCDM	< 64 meV	marginally above
DESI DR2 BAO + CMB, Feldman–Cousins	ΛCDM	< 53 meV	**below**
ACT DR6 + DESI DR2	ΛCDM	< 52 meV	**below**
DESI DR2 + CMB	w_0w_aCDM	< 163 meV	comfortably above

But the interpretation is model-dependent, and we will not lean on that either way. The ΛCDM bounds that sit below the prediction assume $w = -1$, which is this framework's own prediction — so the framework cannot escape into evolving dark energy without abandoning a different derived result. On the other hand a 2026 analysis in a dynamical-dark-energy framework with free N_{eff} and weak-lensing data added finds a 2.7σ preference for a *positive* neutrino mass at $\Sigma m_\nu = 98^{+16}_{-37}$ meV (Du et al., 2026), which would be comfortable. Whether cosmology currently "prefers a negative effective mass" or "detects a positive mass at 2.7σ" depends entirely on the dark-energy parameterization and the dataset combination. Neither should be quoted as established.

The kill, stated in the framework's own words. GG-A10-Cosmos names this as the test most likely to fire first: a robust DESI-plus-CMB bound pushing below ~ 58 meV *under this framework's own* $w = -1$ *prior* would falsify the inherited block. That is a specific, dated, two-sided criterion, and the framework wrote it down before the bound arrived.

The instruments. The near-term improvement will come from the Simons Observatory's enhanced large-aperture telescope, which projects $\sigma(\Sigma m_\nu) = 30$ meV — improving to 15 meV if the reionization optical depth can be pinned to $\sigma(\tau) = 0.002$ — with the upgrade complete in 2028 and the survey running to 2034 (Simons Observatory Collaboration, 2025). Euclid projects $\sigma(M_\nu) = 56$ meV alone and 23 meV combined with Planck (Euclid Collaboration, 2025).

A second correction, and a significant one. CMB-S4, which appears in essentially every neutrino-mass forecast written before mid-2025 and in earlier papers of this cohort, **was terminated on 10 July 2025**, when the Department of Energy and the National Science Foundation jointly withdrew support, citing South Pole infrastructure constraints. A descoped Chile-only proposal was submitted in June 2025 with results possibly in 2040–41; its outcome is uncertain. Any sensitivity forecast in this cohort that lists CMB-S4 as forthcoming needs revision, and this study lists the Simons Observatory and Euclid in its place.

Direct kinematic mass, for completeness. KATRIN's 259-day result gives $m_\beta < 0.45$ eV at 90% confidence (KATRIN Collaboration, 2025), with a final Phase-1 goal of 0.3 eV; Project 8 targets 40 meV with an atomic tritium source but has no approved schedule. Neither reaches the normal-ordering region of m_β, so neither can test the prediction. They are listed here so that a reader does not mistake their silence for support.

62 The lock cluster at HL-LHC

The prediction. A multi-state cluster at $M_{\text{lock}} = 3.85$ TeV: a spin-2 graviton anchored at the lock, a color-octet Kaluza–Klein gluon g^*, and the leptophobic baryonic vector Z'_B of the $\mathrm{U}(1)_B$ sector. The falsifier is the *pattern* — the mass coincidence, the spin assignments, and the coupling hierarchy — not an isolated excess. Finding either the octet or the graviton *without* the other falsifies the cluster reading.

Why this is the chain's hardest constraint. It is the only prediction whose failure retires the matter chain regardless of every leptonic success. If there is nothing at 3.85 TeV, there are no heavy Majorana states at the lock, and Part VIII has no gate.

Where it stands, and a non-trivial survival. ATLAS's 2026 high-mass dilepton search, combining 165 fb^{-1} at 13.6 TeV with Run-2 data, excludes a sequential-Standard-Model Z' below 5.5 TeV and a Z'_ψ below 4.8 TeV (ATLAS Collaboration, 2026). Taken at face value that excludes any new neutral gauge boson near 4 TeV *with Standard-Model-strength couplings to leptons.* The Z'_B predicted here is **leptophobic** by construction — it is the gauge boson of baryon number, and it does not couple to leptons at all — so the dilepton bound does not apply to it. This is a real consistency check, not a rescue: the leptophobia was fixed by the $\mathrm{U}(1)_B$ assignment long before the bound existed, and had the framework predicted a lepton-coupled state at 3.85 TeV it would already be dead.

The live constraints are therefore in the hadronic channels. GG-A9-SM quotes, from the dedicated signal paper, $\sigma \times \mathrm{BR}$ into $t\bar{t}$ of 8.7 fb (range 7–12) at 3.85 TeV for the octet, and for the graviton $\kappa_{gg} = 0.47/K = 0.0136$ with 0.045 fb in $t\bar{t}$ at the LHC and $\Gamma/m = 6.3\%$. The two-tier falsifier is Tier 1 at HL-LHC for the octet, requiring roughly 0.5–2 ab^{-1}, and Tier 2 at a 100 TeV machine for the spin-2 state, where the cross section rises to $\simeq 205$ fb.

What is *not* observable, and why we say so. The heavy Majorana neutrinos themselves are not accessible. Their Yukawa coupling is $Y_\nu \simeq 7.3 \times 10^{-7}$ (Eq. (51)), so the active–sterile mixing is far too small for direct production, and this framework predicts *no* right-handed W_R (registered as a prediction in its own right). The strong ATLAS bounds on heavy Majorana neutrinos — N_R excluded up to 3.6 TeV for $m(W_R) = 4.8$–5.2 TeV (ATLAS Collaboration, 2023) — are bounds on the Keung–Senjanović production mechanism and do not constrain a state produced only through a 10^{-6}-level mixing. So the TeV scale here is testable through the *lock cluster*, not through the neutrinos. A reader who expects "TeV leptogenesis" to mean "make the heavy neutrinos at a collider" should adjust that expectation: the mechanism is at a testable scale, but its own quanta are not.

63 Proton decay and n–$\bar{n}$ oscillation

The prediction. $\tau_p = \infty$; $\Delta B = 1$ and $\Delta B = 2$ forbidden outright.

Where it stands. Super-Kamiokande gives $\tau/B(p \to e^+\pi^0) > 2.4 \times 10^{34}$ yr and $\tau/B(p \to \mu^+\pi^0) > 1.6 \times 10^{34}$ yr from 0.373 Mt·yr (Super-Kamiokande Collaboration, 2020), and $\tau/B(p \to \bar{\nu}K^+) > 5.9 \times 10^{33}$ yr from 260 kt·yr.

The instruments. Hyper-Kamiokande projects, over twenty years, 10^{35} yr in $p \to e^+\pi^0$ and 3×10^{34} yr in $p \to \bar{\nu}K^+$ — about an order of magnitude beyond Super-K. DUNE adds 1.3×10^{34} yr in $p \to K^+\bar{\nu}$ at 400 kt·yr, where a liquid-argon detector's kaon identification is the strongest in the field.

The kill. One event, in any channel. This is the framework's cleanest and most dangerous exposure, and it is a test the framework can only fail — as §51 said, no finite exposure establishes $\tau = \infty$.

64 Electric dipole moments, and a direct sphaleron handle

Electric dipole moments. A new *CP*-violating phase that was not on the lattice would generically show up as an electric dipole moment. The strongest constraint is the electron EDM from the JILA HfF^+ measurement, $|d_e| < 4.1 \times 10^{-30}\, e\,\text{cm}$ at 90% confidence (Roussy et al., 2023), which probes new physics at scales an order of magnitude beyond the LHC's direct reach; the neutron EDM bound is $|d_n| < 1.8 \times 10^{-26}\, e\,\text{cm}$ (nEDM Collaboration, 2020), with n2EDM at PSI targeting 1.1×10^{-27} by the end of 2026 and $\sim 5 \times 10^{-28}$ in a 2029–2033 campaign. The framework predicts $|\bar{\theta}| \lesssim 5 \times 10^{-18}$ from the relaxation mechanism of §15.1, which is far below any planned sensitivity. An EDM signal demanding a rogue phase off the lattice would wound the alphabet at the same joint that a rogue δ_{PMNS} would.

A direct sphaleron handle, new in 2026. CMS has published the first collider limit on sphaleron transitions, bounding the pre-exponential factor — the fraction of quark–quark interactions undergoing a transition above threshold — at < 0.0034 at 95% confidence for a nominal 9 TeV threshold (CMS Collaboration, 2026). This does not test the chain, whose sphaleron physics is thermal and standard, but it is worth recording that the door of §42 now has a direct experimental probe of any kind at all.

65 The decision calendar

66 A remark on the shape of this program

Nine criteria, seven instruments, fifteen years, and no fallback written for any of them. That is a deliberately brittle position, and it is worth saying why we prefer it.

A framework with adjustable parameters responds to a disagreeable measurement by adjusting. It survives, and each survival costs it a little of whatever explanatory claim it had. A framework built on five integers cannot do that. If Σm_ν is measured robustly at 40 meV, there is no version of

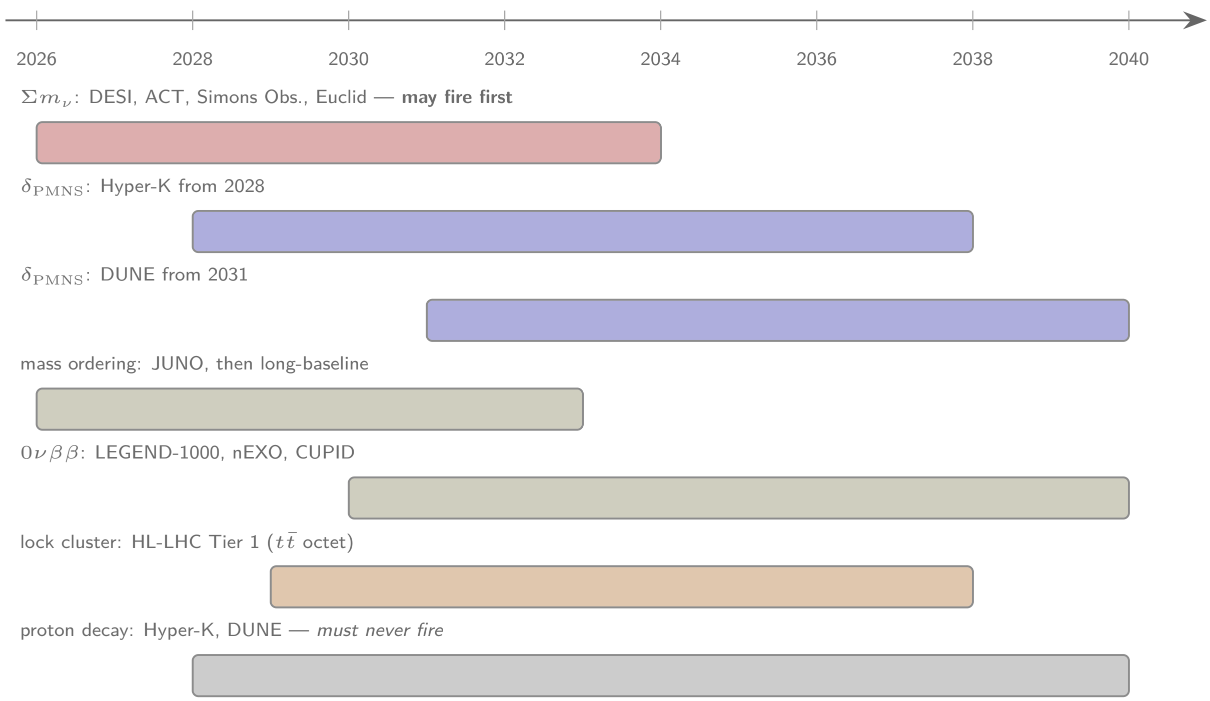

Figure 18: The decision calendar. Seven independent decisions, all inside fifteen years. The first row is the framework's most immediate exposure: the derived $\Sigma m_\nu = 59$ meV sits at the normal-ordering floor and two current ΛCDM analyses already bound it below that. The last row is the one that can only go wrong: a single proton decay ends the baryonic logic regardless of every other success.

Eq. (37) that accommodates it. If the ordering is inverted, there is no reassignment of the Morse depths that survives. If the lock cluster is absent at 3.85 TeV, K is not 35 and nothing else in the ledger stands either.

The corresponding gain is that agreement means something. When $\sin\theta_{12} = 0.225019$ lands on a measurement of 0.22501 ± 0.00068, there was no parameter that could have been moved to put it there. That is the bargain, and this Part is where the price is paid.

Table 20: The kill table. Each row is a specific, dated criterion, and each names what it takes with it. Nothing in this table is contingent on the others: any one of them can fire alone. Two rows are worth noting for what they concede. The lock cluster kills the heavy sector regardless of every leptonic success, and the last row kills the chain from within, leaving the agreement at the foot of the arithmetic exposed as coincidence.

Prediction	Kill criterion	What dies
$\delta_{\mathrm{PMNS}} = 196.4°$	measured value inconsistent at 5σ; settling at maximal 270° (2031–35)	the leptonic lattice chain; the door's bias; the $3\pi/2$ branch is pre-registered dead and may not be used as a fallback
Normal ordering	inverted ordering established at 5σ ($\sim$2033)	the depth assignment, hence the light spectrum; no fallback
$\Sigma m_\nu = 59$ meV	a robust bound below ~ 58 meV under the framework's own $w = -1$ prior	the inherited neutrino block; likely the first to fire
Majorana nature; $m_{\beta\beta} \simeq 1.8$ meV	a confirmed $0\nu\beta\beta$ line anywhere in 9–20 meV; or Dirac nature established by any route	the dial-free door as the matter portal, hence the whole of Part VIII
Lock cluster at 3.85 TeV	absence of the octet at HL-LHC with $\sigma \times \mathrm{BR}(t\bar{t}) \simeq 8.7$ fb; or either state without the other	the heavy sector of the chain, *regardless of every leptonic success*
$\tau_p = \infty$	one event, any channel, at Hyper-K or DUNE	the exact $\mathrm{U}(1)_B$ ledger, hence the entire baryonic logic
$\tau_e = \infty$	one event, any channel	the no-exit-gate theorem
$\delta_{\mathrm{CKM}} = 13\pi/35$	converged world fit excluding 1.1669 rad at 5σ	the QD completion (the $2\pi/5$ first order survives only with the $\pi/35$ theory width)
H-ϕ_{geo}-1	finite-temperature transport returning the N-width class, or the flavor-blind g^2T, or an unclosable C6 window deficit	the chain is falsified *from within*; the $e^{-K/2}$ agreement stands exposed as coincidence

Part XIII
Discussion

What the chain establishes, what it assumes, and where it sits in a literature that has been trying to do this for sixty years

67 Sixty years of the question

It is worth placing the result in its lineage, briefly, because the shape of the difference is the whole content of the claim.

Sakharov (1967) asked the question and named the three conditions. Kuzmin et al. (1985) supplied the baryon-number violation — unsuppressed sphaleron transitions in the hot Standard Model — and for a decade it looked as though electroweak baryogenesis might close the case inside known physics. Two results killed that hope: the transition at the measured Higgs mass is a crossover (Kajantie et al., 1996), and the Standard Model's single CP phase delivers a CP-odd measure of order 10^{-20} (Gavela et al., 1994, Huet and Sather, 1995). Fukugita and Yanagida (1986) had already written down the alternative, and leptogenesis has been the leading candidate ever since.

The difficulty leptogenesis has never resolved is its own testability. Davidson and Ibarra (2002) showed that the hierarchical version requires $M_N \gtrsim 10^9$ GeV — permanently beyond every collider, with its phases decoupled from anything measurable. Pilaftsis (1997) and Pilaftsis and Underwood (2004) showed that quasi-degenerate heavy states evade the bound by a resonant enhancement, admitting TeV-scale masses — at the price of engineering the degeneracy. And in either version, the Casas–Ibarra parameterization (Casas and Ibarra, 2001) makes explicit how much freedom is available: an arbitrary complex orthogonal matrix R sits between the measured light masses and the unmeasured heavy sector, and with it one can produce any η_B.

That is the state of the art the present chain is measured against.

68 Where this chain differs

Three differences, in order of increasing consequence.

The phases are not free. The Dirac phases are lattice points of $(\pi/5)\mathbb{Z}$ and $\pi + 2\arcsin(1/J)$; the Majorana phases are parity data $(0, \pi)$ of the rank-two block. The R matrix is real — the door is dial-free — so the freedom that Casas–Ibarra makes explicit is simply absent. There is nothing to tune.

The scale is not chosen. $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-35}$ was fixed by the anomaly audit before the question arose, and the same $K = 35$ sets four other things in the cohort that have nothing to do with baryogenesis. A treatment that selects the TeV scale for accessibility is making a defensible methodological choice; a treatment in which the scale arrives from elsewhere is making a prediction, and it can be wrong at HL-LHC.

The magnitude is one exponential of one integer. Not a product of small dynamical factors, not a washout race, not a heavy mass times a phase. $e^{-K/2} = 10^{-7.6}$, and three textbook factors carry it the remaining distance to 6×10^{-10}.

It is a matter of record that leptogenesis has *accommodated* η_B rather than predicted it, and that the accommodation has been the mechanism's standing weakness for forty years. The claim here is narrow and checkable: this chain has one postulated step, one named input, and no fitted quantity, and it lands twenty per cent away. Whether that is a discovery or a coincidence is decided at the two joints named in §54 and by the seven experiments of Part XII — not by the twenty per cent.

69 The intuitive core, restated once

Strip the machinery away and the physical story is short enough to hold in mind, which is the test of whether it is a story at all.

There are two extra dimensions. One is a circle and one is a line. Loops only exist on circles, so everything that is a *phase* lives on the circle and everything that is an *amount* lives on the line. The circle has a pattern pressed into it, ten letters at even spacing, and the pattern does not match itself read backwards. That mismatch is *CP* violation. It is not an accident and it is not a parameter: the spacing comes from a five, the branch comes from a Higgs doublet keeping two out of five, and $3\pi/5$ is not a multiple of π.

Color writes a different figure on the same circle — a triangle — and the triangle's circuit closes with nothing left over. So the strong force does not violate *CP*, and by the same closure condition it confines.

The quark sector has the mismatch but almost no leverage: its area is a hundredth of the lepton sector's, and the enormous spread of quark masses throttles it further, so it falls ten to twelve orders short of what the universe needs. The lepton sector has the leverage, and it also has something the quark sector cannot have — a neutral species, whose winding number is zero, which is the only place the geometry permits a particle to be its own antiparticle. That permission is the door through which lepton number can be violated.

Behind the door sit heavy Majorana states, at a scale the ledger had already fixed at four TeV, quasi-degenerate at the level of one part in thirty-five. As the cooling universe crosses that scale they decay, once and out of equilibrium, and their decays are biased — by a phase of π — a parity, not a dial — which biases the three lepton flavors by amounts that cancel exactly when added and do not cancel when the flavors are washed out at different rates. The bias is small in absolute terms, $e^{-17.5}$, because the toll for crossing half a lock depth is what it is. Sphalerons then convert twenty-eight seventy-ninths of the resulting $B - L$ into baryons before the electroweak crossover welds the door shut, entropy dilutes the result by the degrees of freedom in the bath, and what is left is six parts in ten billion.

And because baryon number can only ever change by three at a time, none of those baryons can decay afterwards. The universe made its matter once, at four TeV, in the first 10^{-13} seconds, and it will never lose any of it.

70 Open items, stated as such

Five, and none of them cosmetic.

1. **H-$\phi_{\rm geo}$-1, condition C4.** The regulator identification — that the interference time of the flavor-odd source is $1/\Gamma_\tau$ and not $1/\Gamma_N$ — is open at derivation level. The outstanding content is a full Kadanoff–Baym resummation. This is the single most valuable calculation anyone could do to this chain, in either direction.
2. **The $K/2$ exponent.** Postulated. Conjecture C-P5-C(ii) must fix it from first principles.
3. **The (H2) baryon-descent computation.** Proton absolute stability is a theorem conditional on an explicit computation that the holonomy descent realizes the bulk $\mathrm{U}(1)_B$ as genuine baryon number, defeating the rank-matching obstruction. GG-A9-SM describes it as finite and self-contained and defers it to a successor paper. Until it is done, the headline prediction of §51 carries a conditional.
4. **The Ω_b seam.** The transport route and the topological route to the baryon density disagree by 16.4% against a 0.3% measurement (§48). Graded open; owner unidentified.
5. **The rank-two structure.** The Majorana pair $(0, \pi)$ follows from the rank-two block, but the rank-two structure itself is asserted by reference rather than separately derived (§33), and the two parent papers grade the phases differently in different places.

71 What would make this work look different in ten years

Two outcomes, both scientific.

If H-$\phi_{\rm geo}$-1 closes — a finite-temperature transport computation landing at the partial-resonance point — then Proposition 37.1 is promoted from D* toward derived, and what remains is one postulated exponent between a small integer and the amount of matter in the universe. In that case the seven experiments of Part XII become tests of a derivation rather than of a conjecture, and the lock cluster at HL-LHC becomes the most important measurement in the program.

If H-$\phi_{\rm geo}$-1 fails — the enhancement dead at the crossing for every candidate mechanism — then the chain is falsified from within, the $e^{-K/2}$ agreement stands exposed as a coincidence, and this Part's claim reverts to Sakharov custody plus an unexplained number. That is a real possibility and the framework has written down in advance what it would mean.

Both outcomes are scientific. Only the unexamined middle would not be.

Part XIV
Conclusions

72 What has been established

1. ***CP* violation is compulsory, and its origin is geometric.** On a bulk with one compact fiber and block integer $L = 5$, the loop holonomies of the flavor connection are valued in the pentagonal lattice $(\pi/5)\mathbb{Z}$; the Higgs doublet's retention of $2/L$ forces the branch phase $3\pi/5$; and CP is conserved only if every loop angle lies in $\{0, \pi\}$. Since $3\pi/5$ does not, CP violation must exist. Carried at derived-conditional grade on the rigidity of the holonomy moduli.

2. **The same picture explains the sectors that do *not* violate *CP*.** Color writes a triangle whose circuit is trivial, so strong CP is exact and color confines — one closure condition, two consequences. The compact-angle sector relaxes to the CP-conserving floor, which is the strong-CP problem dissolved by a field the framework already required.

3. **The phase is a lattice angle, not an accumulated one.** The helix construction of Part IV makes the alternative quantitative and refutes it: the two sectors' pitches differ by a factor of 2.73 over nearly the same interval of r. A universal pitch is excluded, confirming numerically what the topology of the interval already required.

4. **The quark sector cannot pay for us.** Its Jarlskog area is three hundred times smaller than the lepton sector's, and the product of quark mass-squared splittings over T^{12} throttles it to 10^{-22}–10^{-19} against the 6×10^{-10} required. Together with the crossover, this closes the electroweak door; exact baryon number closes the grand-unified door; the lepton route is what remains.

5. **The Majorana door is unique.** A $\Delta L = 2$ bilinear exists if and only if the fiber weight vanishes, so the neutral sector is the only place the geometry admits it, and LHN is the only portal. The light masses follow with no new scale, $m_\star = \pi v^2/E_{\text{GUT}}$, giving normal ordering as an output and $\Sigma m_\nu = 59$ meV.

6. **The driver is a Majorana parity, and it is *CP*-conserving in the flavor sum and *CP*-violating in the flavor differences.** $(\alpha_{21}, \alpha_{31}) = (0, \pi)$ are parity data of the rank-two block, not dials. The flavor-summed asymmetry vanishes exactly; the individual flavored asymmetries do not, and differential washout converts the difference into a net excess. This is why δ_{PMNS}, the phase the long-baseline program is built to measure, is a spectator: the asymmetry is J_{PMNS}-blind at the few-percent level. The letter that made us is invisible to every oscillation experiment ever proposed.

7. **Leptogenesis happens at four TeV, and the scale was not chosen.** $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-K} = 3.85$ TeV comes from $K = 35$, which the anomaly audit fixed for reasons unconnected with baryogenesis. Quasi-degeneracy at $\mathcal{O}(1/K)$ puts the mechanism in the resonant regime, evading the Davidson–Ibarra bound $M_N \gtrsim 10^9$ GeV *where the literature evades it* — but with the degeneracy faced in public as a named hypothesis, bracketed from both sides, rather than engineered away. *So far as we can establish, a TeV-scale Majorana leptogenesis whose scale is derived rather than selected is without precedent.*

8. **The baryon-to-photon ratio is a closed-form consequence of one integer.**

$$\eta_B \;=\; 7.04 \times \frac{28}{79} \times \frac{1.25\, e^{-K/2}}{g_*} \;=\; 7.34 \times 10^{-10} \qquad \text{versus} \qquad (6.12 \pm 0.04) \times 10^{-10}.$$

Zero fitted parameters, one named input, one postulated exponent, and — new in this version — an error bar that is derived rather than declared. The full flavored Boltzmann integration, with the washouts $K_\alpha = (3.5,\, 29.3,\, 22.5)$ computed from the derived spectrum and mixing matrix and the kernels validated to four per cent, gives the best-estimate prediction $\eta_B = 2.8^{+2.0}_{-1.2} \times 10^{-10}$: consistent with the measurement at 1.3σ on the log scale, in a band one third the width of the previous honest statement, with the Boltzmann-level deficit now measured at $E = 7.1$ [4.4, 11.4] and assigned to named, standard, one-sided transport physics. The sign of the excess emerges as an output: matter survives because the electron flavor, decoupled from the third mass state by the derived $\sin^2\theta_{13} = 3/140$, outlives the washout that erases the rest — ninety-six per cent of us came through the electron door. *So far as we can establish, no previous treatment has obtained η_B as a closed-form consequence of a small integer with no fitted quantity anywhere in the chain.* We state that as a claim about the literature we have searched, not as a theorem, and we would withdraw it on being shown a counterexample.

9. **The excess is one part in a billion because half a lock depth is $10^{-7.6}$, and the universe paid the toll exactly once.**

10. **Protons and electrons are absolutely stable, and the first of those is the same statement as the baryogenesis.** Baryon number changes only in units of $N_{\rm gen} = 3$; a bulk reprocessing does not care and a single nucleon cannot proceed. So $\Delta B = 1$ and $\Delta B = 2$ are forbidden outright — proton decay and n–$\bar{n}$ oscillation both — and where grand unification is confirmed by finding proton decay, this framework is refuted by it. The electron is stable for an unrelated reason: it is the ground state of its stratum and has no exit gate, so its stability is an absence of destinations rather than a conservation law.

11. **The program is brittle on purpose.** Nine dated kill criteria, seven instruments, fifteen years, no fallback written for any of them. The most immediate exposure is the derived $\Sigma m_\nu = 59$ meV against cosmological bounds that already sit at or below it under the framework's own $w = -1$ prior; the framework named that test as the one most likely to fire first, before the bound arrived.

73 Closing

The question this study set out to answer was where the matter came from, and the answer it reports is that the matter came from a mismatch of $3\pi/5$ against π on a circle of circumference set by an integer that also fixes the fine-structure constant, the dark-matter line frequency, and the vacuum energy.

That is either a real discovery about the world or a coincidence of unusual persistence, and we have tried to make it easy to find out which. The postulate is stated as a postulate; the named input is bracketed from both sides, with three routes by which the chain refutes itself; where an experimental expectation has changed we have said so against our own interest; and where two derivations inside the framework disagree by twenty-five standard deviations, we have said that too.

What we will not do is hedge the prediction. If the neutrinos are Dirac particles, if the ordering is inverted, if Σm_ν lands at 40 meV, if there is nothing at 3.85 TeV, or if a single proton decays in any channel, then this account is wrong and there is no version of it that survives. That is the bargain we offer.

Acknowledgments

This study assembles results owned by GG-A9-SM, GG-A10-Cosmos, GG-P3-Electron, GG-P5-CPT and GG-W2-Holonomy, and carries each at its parent's grade. The author thanks the experimental collaborations whose published sensitivities make Part XII possible — DUNE, Hyper-Kamiokande, T2K, NOvA, JUNO, KamLAND-Zen, LEGEND, CUORE, nEXO, CUPID, DESI, ACT, the Simons Observatory, Euclid, KATRIN, ATLAS, CMS, and the JILA and PSI electric-dipole-moment groups — and notes that a framework of this kind is only worth writing because those programs make it refutable.

Part XV
Two Talks: Why Are We Here?

The same result at two levels — a physics colloquium, and a lecture for everyone

This Part contains no new physics. It contains the same chain assembled twice, for two audiences, because the argument of this study is unusual in a specific way: its hardest step is a piece of representation theory that can be drawn, and its most quoted number can be reached by two different roads that do not yet agree. Both facts are easier to say out loud than to read off a Boltzmann equation, and neither survives being compressed into a single slide.

Division 1 is pitched at a departmental colloquium: an audience that knows what a Yukawa coupling is, has seen a seesaw, and will want to know immediately where the free parameters went. Division 2 is pitched at a general audience with no physics background at all — a public evening lecture, or a recorded talk — and it assumes nothing beyond the idea that matter is made of particles.

The two Divisions are independent. Reading either one alone loses nothing except the pleasure of watching the same skeleton wear two different suits.

One instruction that applies to both talks. The result of this study is not that the matter asymmetry has been *explained*. It is that a chain with no fitted parameters lands within an order-unity factor of the observed number, and that a completely independent static count lands within half a percent of the same number. Neither statement is proof. Both are worth an hour.

74 Division 1: A colloquium — the geometric origin of matter

74.1 The number, and sixty years of not deriving it

Open with the number and do not bury it.

For every baryon that survives into the present universe there were, at the time the light element abundances were fixed, about 1.6 billion photons. The measured ratio is

$$\eta_B \equiv \frac{n_B}{n_\gamma} = (6.12 \pm 0.04) \times 10^{-10}, \tag{93}$$

and it is one of the best-measured numbers in cosmology, fixed independently by the acoustic peaks of the microwave background and by primordial deuterium. There is no measurement problem here. The problem is that the Standard Model, run forward from a symmetric initial condition, produces an asymmetry smaller than this by some ten orders of magnitude.

The audience already knows why. Sakharov's three conditions require baryon number violation, CP violation, and departure from thermal equilibrium. The Standard Model supplies the first through the electroweak anomaly, supplies the second only through the Cabibbo–Kobayashi–Maskawa phase, and fails the third because the electroweak transition at the observed Higgs mass is a smooth crossover rather than a first-order phase transition. The quantitative failure is dominated by the second: the Jarlskog invariant is $J_{\rm CKM} \simeq 3 \times 10^{-5}$, and once it is divided by the twelfth power of the temperature it must be divided by to become dimensionless, what is left is of order 10^{-20}.

The framing that makes the talk work. Do not present this as "the Standard Model is wrong". Present it as a *division of labor*. The Standard Model is the most accurate predictive instrument ever built — it computes the electron's magnetic moment to twelve significant figures — and it achieves that accuracy by measuring twenty-eight numbers and then propagating them with extraordinary care. What it cannot do is *produce* a number it was never given. The matter asymmetry is not a propagation problem. It is an origin problem, and origin problems are the ones a parameter-free geometry is for.

74.2 Two fibers, and what each of them remembers

The arena is a six-dimensional bulk that factorizes,

$$M_6 \;=\; M_4 \times I_r \times S^1_\theta, \tag{94}$$

with a four-dimensional boundary on which everything is measured. The two extra directions are not alike, and the entire argument of this study turns on the difference between them.

$I_r = [0, K]$ is an *interval*, coordinatized by $r = \ln(E_{\rm GUT}/\mu)$ — scale, in logarithmic units, from the unification energy down to whatever energy the experiment runs at. Transport along an interval is a semigroup: there is a direction of travel, information is discarded as one moves, and there is no loop to come back around. This is the renormalization group, drawn as a geometry.

S^1_θ is a *circle*. Transport around a circle is a group: every path can be run backwards, nothing is discarded, and there are loops — so there is holonomy, and holonomy is quantized by a winding number that no continuous deformation can change.

$$\boxed{\textit{The circle remembers. The interval forgets.}} \tag{95}$$

Six words, and they carry the talk. Everything the framework claims to protect exactly — charge, lepton number, baryon number modulo three, CPT — is protected by a winding number on the circle. Everything the framework computes only approximately — running couplings, thresholds, the anomalous magnetic moment — lives on the interval, where integers do not survive.

74.3 Three species, three readings of one holonomy

Now the step the audience has come for. Take a fermion, transport it once around the phase circle, and ask what happens to it. The answer depends only on its weight n under the circle action — an integer, not a parameter — and there are exactly three cases in the Standard Model's neutral and charged lepton sector.

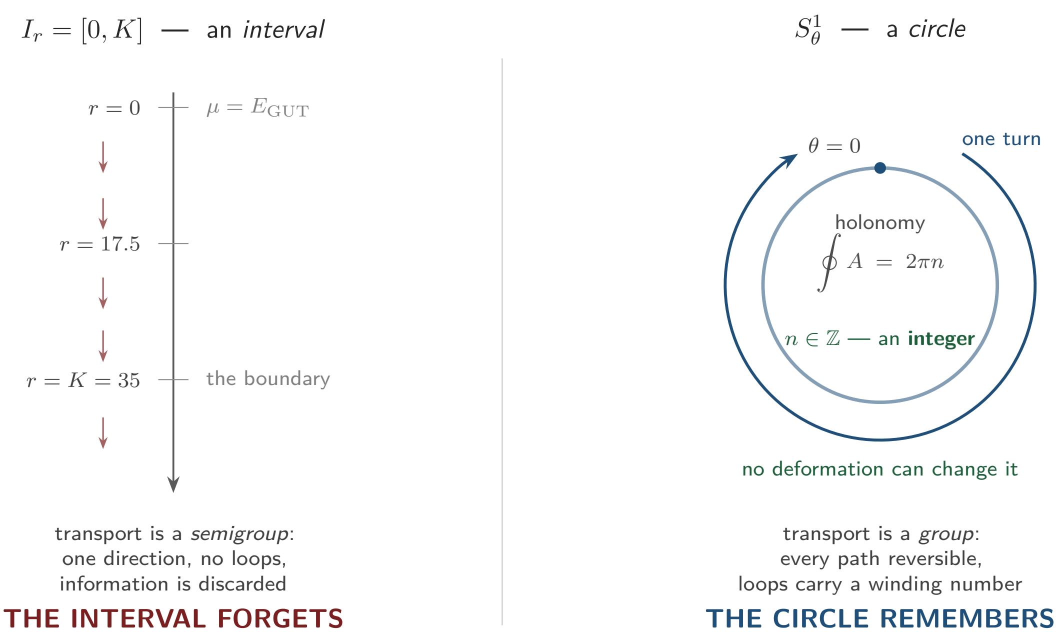

Figure 19: The two fibers, and what each of them remembers. The bulk factorizes as $M_6 = M_4 \times I_r \times S^1_\theta$ and the two extra directions are structurally different. Along the interval — scale, in logarithmic units — transport is a semigroup: one-way, lossy, with no loop to return on. That is the renormalization group drawn as a geometry, and nothing exact survives it. Around the circle transport is a group, and every loop carries a winding number that is an integer. Everything this framework protects exactly — charge, lepton number, baryon number modulo three, CPT — is protected by that integer. Everything it computes only approximately lives on the interval.

A charged lepton has $n \neq 0$, and the circle is faithful.

The electron picks up a genuine phase, and that phase is what electric charge *is* in this picture. Because the winding number is a nonzero integer, the field and its conjugate transform oppositely; a mass term must pair e_L with e_R, and lepton number is conserved exactly. There is no way to write a term that pairs the electron with itself, because the two copies would carry weight $2n \neq 0$ and the term would not be invariant. *The charged lepton has no door.*

A Dirac neutrino has $n \neq 0$ as well, by assumption.

If one simply adds a right-handed partner and writes a Dirac mass, the neutrino inherits the same protection the electron has: a conserved lepton number, three generations, and nothing that violates L. This is the case worth spending two minutes on, because it is the one most of the audience will have in mind, and it is the one that *fails to help*. A Dirac neutrino is not a door. It is a heavier electron with no charge.

A Majorana neutrino has $n = 0$, and only then.

Write the symmetric square of a weight-n field. It decomposes into a weight-zero piece and a weight-$2n$ piece. The lepton-number-violating bilinear lives in the weight-zero piece — and the weight-zero piece is available as an invariant *if and only if* $n = 0$. This is the door theorem, and it is the single most important line in the study:

$$\boxed{\text{A } \Delta L = 2 \text{ bilinear exists} \iff \text{the fiber weight vanishes.}} \tag{96}$$

The neutral lepton is the only field in the Standard Model with vanishing weight. Therefore it is the only field that can pair with itself. Therefore it is the only place in the entire particle content where lepton number can be violated at all.

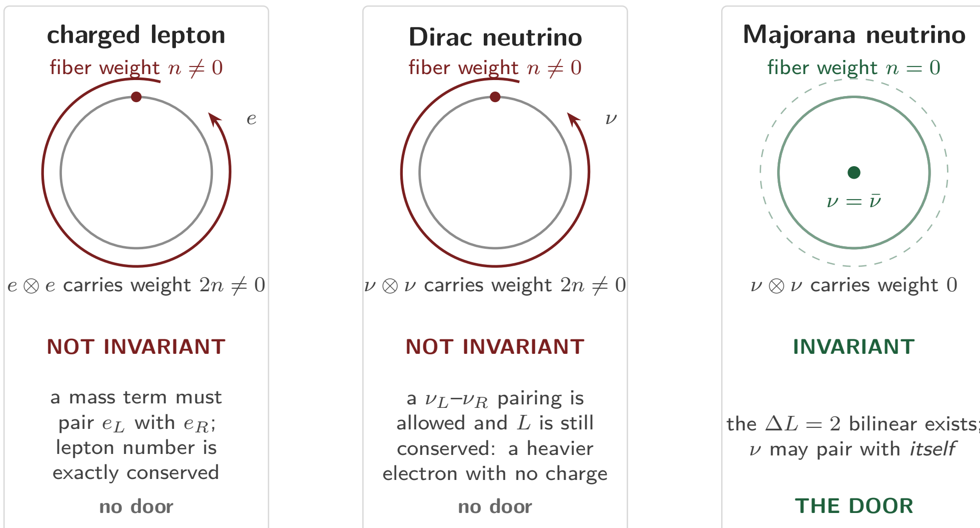

Figure 20: Three species, three readings of one holonomy. Take a fermion around the phase circle and ask what it picks up. A charged lepton has nonzero fiber weight — that weight *is* its electric charge — so the symmetric square of the field carries weight $2n \neq 0$ and cannot appear as an invariant: lepton number is conserved and there is nothing to violate it with. A Dirac neutrino inherits exactly the same protection, which is why adding a right-handed partner does not help. Only a field of vanishing weight has a weight-zero piece in its symmetric square, and only then can a field pair with itself. The neutral lepton is the only such field in the Standard Model. It is not the best candidate for lepton-number violation — it is the *only* one.

The line to deliver slowly. The usual reason given for looking to neutrinos rather than quarks is that the quark sector's CP violation is too small. That is true, and it is the weaker statement. The stronger one is that the quark sector *has no door.* No amount of CP violation in a sector that conserves lepton and baryon number can produce a lepton or baryon asymmetry, because there is no operator to carry it. The quark sector is not a poor candidate. It is not a candidate.

74.4 Where the scale comes from, and why it was not chosen

An audience that has seen leptogenesis before will expect right-handed neutrinos at 10^9–10^{15} GeV, and will expect to be told that the scale is unobservable. Here it is neither.

The interval has length $K = 35$, an integer of the ledger fixed by an anomaly audit that has nothing to do with baryogenesis. The heavy Majorana states sit at

$$M_{\text{lock}} = \tfrac{1}{2} M_{\text{Pl}}\, e^{-K} \simeq 3.85 \text{ TeV}, \tag{97}$$

with splittings of order $1/K$. Two consequences follow, and both are worth a slide.

First, the mechanism is *resonant*: quasi-degenerate heavy states enhance the self-energy contribution to the decay asymmetry by orders of magnitude, which is what makes leptogenesis at a few TeV arithmetically possible at all. Second, the light neutrino masses follow from a seam coefficient with no new scale, $m_\star = \pi v^2/E_{\rm GUT} \simeq 8.35$ meV, giving $(m_1, m_2, m_3) \simeq (0.065,\, 8.59,\, 50.5)$ meV — *normal ordering as an output, not an input* — and $\Sigma m_\nu \simeq 59$ meV.

That last number is the tightest live tension in the whole study and should be said out loud rather than left in a table: cosmological bounds on the neutrino mass sum are closing on exactly this value, and if they close past it the chain is in trouble.

74.5 The parity, and the paradox it appears to create

Here is the step that a careful audience will interrupt, and the speaker should get to it before they do.

The two Majorana phases are not dials. The rank-two structure of the light mass matrix fixes them to

$$(\alpha_{21},\ \alpha_{31}) \ =\ (0,\ \pi), \tag{98}$$

which is a statement about relative intrinsic *CP* parity: 0 means same parity, π means opposite. There is no continuum to choose from and no third value.

And now the objection, which is correct and which the speaker should make on the audience's behalf: *those are exactly the values that conserve* CP. The criterion is that *CP* is conserved when every loop holonomy lies in $\{0, \pi\}$, and $(0, \pi)$ is a pair of such values. Any neutrino theorist in the room knows that Majorana phases of 0 and π are the textbook signature of a *CP*-conserving configuration. How can the driver of the matter excess be a *CP*-even letter?

The honest answer is that the textbook statement is about a specific quantity — the *flavor-summed* asymmetry — and that this quantity does indeed vanish here, exactly.

74.6 The resolution: conserving in the sum, violating in the differences

Compute the flavored decay asymmetries ϵ_α for $\alpha = e, \mu, \tau$ at this study's own inputs, in a standard type-I seesaw with a real orthogonal Casas–Ibarra matrix, keeping both the vertex and self-energy contributions at one loop. The result is the central slide of the talk:

$$\epsilon_e : \epsilon_\mu : \epsilon_\tau \ =\ 8.52\times 10^{-14}\ :\ -4.75\times 10^{-13}\ :\ 3.90\times 10^{-13}, \qquad \sum_\alpha \epsilon_\alpha \ =\ 5\times 10^{-29}. \tag{99}$$

The individual asymmetries are of order 10^{-13}. Their sum is of order 10^{-29} — sixteen orders of magnitude smaller, which is zero to the precision of the arithmetic. So both statements are true at once, and neither is a hedge:

$$\boxed{\begin{array}{c}\alpha_{31} = \pi \text{ is } CP\text{-conserving in the flavor sum,}\\ \text{and } CP\text{-violating in the flavor differences.}\end{array}} \tag{100}$$

The structural reason is visible in one line. With opposite CP parity between ν_3 and $\nu_{1,2}$ and a real orthogonal R, the imaginary parts that drive the asymmetry are

$$\mathrm{Im}[Y^*_{\alpha i} Y_{\alpha j}] \;\propto\; \sqrt{m_1 m_3}\, V_{\alpha 1} V_{\alpha 3}\, (R_{i3}R_{j1} - R_{i1}R_{j3})\,, \tag{101}$$

which is manifestly nonzero flavor by flavor, and which sums to zero over α for one reason only: the unitarity of the mixing matrix gives $\sum_\alpha V_{\alpha 1} V_{\alpha 3} = 0$. The cancellation is a column-orthogonality statement, not a dynamical one.

That is the whole resolution. A zero-sum set of flavored asymmetries is not a zero asymmetry, because the three flavors are not washed out equally. Differential washout — the electron door is the leaky one here, retaining some ninety-six per cent of its survivor while the muon and tau doors are scrubbed — converts a set that sums to zero into a net excess.

Why this is the strongest paragraph in the study and not the weakest. It is checkable: the flavor sum is a number one can compute and watch vanish. It explains *why* a flavored treatment is mandatory rather than merely more careful — an unflavored calculation of this chain returns exactly zero, and would have been reported as a failure. And it promotes the flavor-sum cancellation from a technical nuisance to the load-bearing step of the argument.

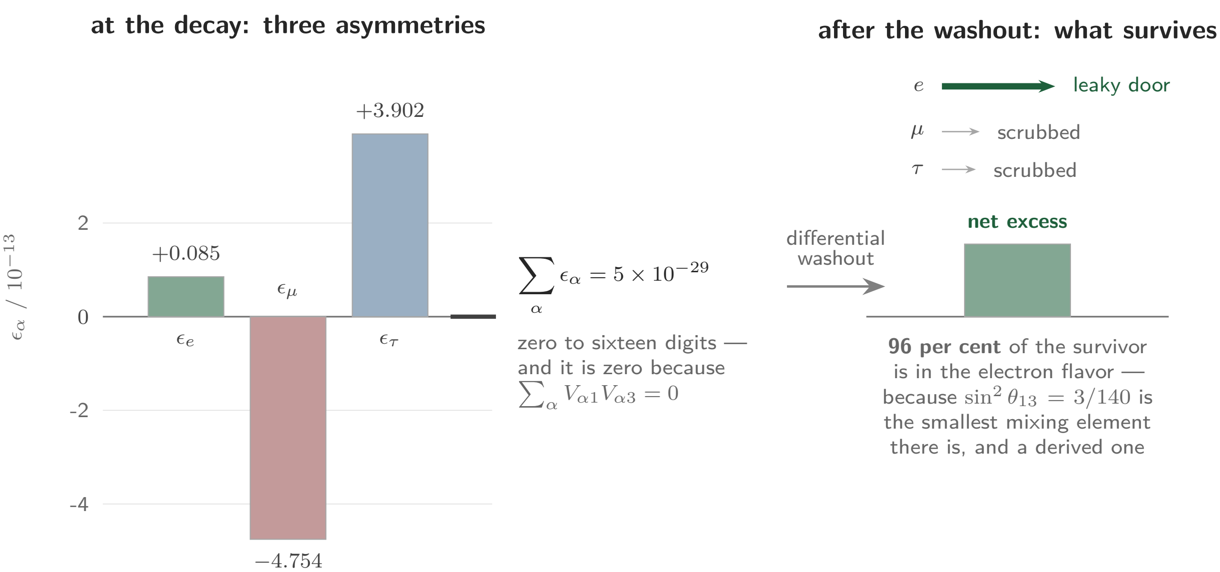

Figure 21: *CP*-conserving in the sum, *CP*-violating in the differences. *Left:* the flavored decay asymmetries at this study's own inputs, computed in a standard type-I seesaw with a real orthogonal Casas–Ibarra matrix and both vertex and self-energy contributions at one loop. The three are of order 10^{-13}; their sum is 5×10^{-29}, which is zero to the precision of the arithmetic. The cancellation is column orthogonality of the mixing matrix, not dynamics — so it is exact, and it is the reason a flavor-summed calculation of this chain returns nothing at all. *Right:* the three flavors are not erased equally. The electron door is the leaky one, because the reactor angle $\sin^2\theta_{13} = 3/140$ decouples the electron flavor from the heaviest state, and a zero-sum set of asymmetries becomes a net excess. This is why a flavored treatment is mandatory here rather than merely more careful.

74.7 The Dirac phase is a spectator, and here is the measurement that says so

The long-baseline program — DUNE, Hyper-Kamiokande — is built to measure δ_{PMNS}, and this study predicts it: $\pi + 2\arcsin(1/J) = 196.4°$. An audience will reasonably assume that a theory of the matter asymmetry which also predicts the leptonic *CP* phase must be using the second to get the first.

It is not, and the cleanest way to show that is to move the phase and watch nothing happen. Table 21 does exactly that: the same one-loop flavored asymmetry, recomputed with each letter turned on and off in turn.

Table 21: What each letter is worth. The magnitude of the flavored asymmetry, relative to this study's own point, as the Dirac phase and the Majorana parity are moved independently. Same seesaw, same inputs, real orthogonal R; the last row is the control. The first four rows move the Dirac phase across its entire range and change the asymmetry by under three percent; the fifth switches the Majorana parity off and removes ninety-five percent of it. That contrast is the Part's whole claim, in six rows.

configuration	what it tests	$\lvert\epsilon\rvert$ relative
$\delta_{\text{PMNS}} = 196.4°$, $\alpha_{31} = \pi$	this study's point	1.0000
$\delta_{\text{PMNS}} = 220°$, $\alpha_{31} = \pi$	a different measured value	0.9919
$\delta_{\text{PMNS}} = 180°$, $\alpha_{31} = \pi$	Dirac *CP* switched off entirely	1.0018
$\delta_{\text{PMNS}} = 0°$, $\alpha_{31} = \pi$	the far end of the range	0.9728
$\delta_{\text{PMNS}} = 196.4°$, $\alpha_{31} = 0$	Majorana parity switched off	**0.0485**
$\delta_{\text{PMNS}} = 180°$, $\alpha_{31} = 0$	**both** switched off (control)	**0.0000**

Three readings, in the order an audience will want them.

Moving the Dirac phase does essentially nothing. Switching it off completely changes the answer by two parts in a thousand — and changes it *upward*, which is the clearest possible demonstration that it is not the source. Setting it to 220°, a value the current data do not exclude, changes the answer by eight parts in a thousand.

Switching off the Majorana parity costs a factor of twenty. That is the letter carrying the asymmetry.

With both switched off the answer is 1.3×10^{-29}, fifteen orders of magnitude below the others. This is the control that matters: the computation demonstrably knows what *CP* conservation looks like, so the nonzero entries above it are not numerical noise.

The answer to the question the audience will actually ask. If DUNE returns 220° instead of 196°, the prediction of the alphabet is wrong and should be reported as wrong — but the account of the matter asymmetry is untouched, because it never used that letter. The two claims are independently falsifiable. That is a feature to state plainly rather than a hedge to bury: a framework whose predictions all stand or fall together is harder to test, not easier.

And the corollary that closes the subsection: the letter that made us is invisible to every oscillation experiment ever proposed, because Majorana phases cancel identically out of oscillation probabilities. The only laboratory that can see it is neutrinoless double beta decay — where this study's prediction, $m_{\beta\beta} \simeq 1.7$–$1.8$ meV, is roughly an order of magnitude below the reach of the next generation. That is an uncomfortable place to be, and saying so is better than not.

74.8 The number, twice: a static road and a dynamical one

The study reaches η_B by a dynamical chain: the door, the flavored asymmetries, differential washout, sphaleron conversion at 28/79, and entropy dilution. It lands at $\eta_B \simeq 7.34 \times 10^{-10}$ against an observed 6.12 — high by twenty per cent, inside a declared factor-of-three band.

There is a second road to the same number, and it comes from a completely different direction. Equipartition of a frozen stationary measure on a compact space gives each sector the share of the volume it occupies. The whole is the three-sphere the ledger selects, of volume $2\pi^2$; matter runs along a single Hopf fiber, a circle of circumference 2π; and baryons carry one unit of $\mathrm{U}(1)_B$ charge per revolution. Therefore

$$\Omega_m = \frac{2\pi}{2\pi^2} = \frac{1}{\pi}, \qquad \Omega_b = \frac{1}{2\pi^2} = 0.050661, \qquad \frac{\Omega_b}{\Omega_m} = \frac{1}{2\pi}. \tag{102}$$

The last equality is the one to put on a slide by itself: the baryon share of all matter is *one full turn of the circle*. Ordinary matter is the matter that winds.

These are not three separate predictions. The baryon-to-photon ratio, the baryon density and the baryon fraction are one quantity in three sets of units, related by an exact conversion,

$$\eta_{10} = 273.9\,\Omega_b h^2, \qquad \rho_b = \Omega_b h^2 \times 1.87847 \times 10^{-29}\ \mathrm{g\,cm^{-3}}, \tag{103}$$

which contains no physics of this framework's at all. Fix one and the other two follow. Converted through (103) at this framework's own derived expansion rate, $h = 0.666$, the ledger road gives $\eta_{10} = 6.15$ against the observed 6.12: high by six tenths of a per cent, with nothing fitted anywhere in the chain. At the microwave background's own $h = 0.674$ it gives 6.30, high by three per cent.

And the open item, which belongs in the talk and not in a footnote. The two roads compute the same quantity and disagree by **16.4%**. The static road is the accurate one — its error is six tenths of a per cent to three per cent against the dynamical road's twenty per cent with a factor-of-three band, a factor of about thirty-eight in accuracy — and the asymmetry in accuracy has a structural cause rather than an accidental one: a statement about volumes on a compact space carries no dynamics to get wrong, while the dynamical chain must pass through a finite-temperature resonance condition that this study grades as an open, named hypothesis. The right reading is that the ledger road is the framework's prediction of the number, and the dynamical chain is a demonstration that a parameter-free mechanism can reach that order of magnitude. Making the two agree is unfinished business, and it is stated here as unfinished.

74.9 The weak joint, said out loud

No colloquium on this material is honest without the following slide, and it should not be the last one.

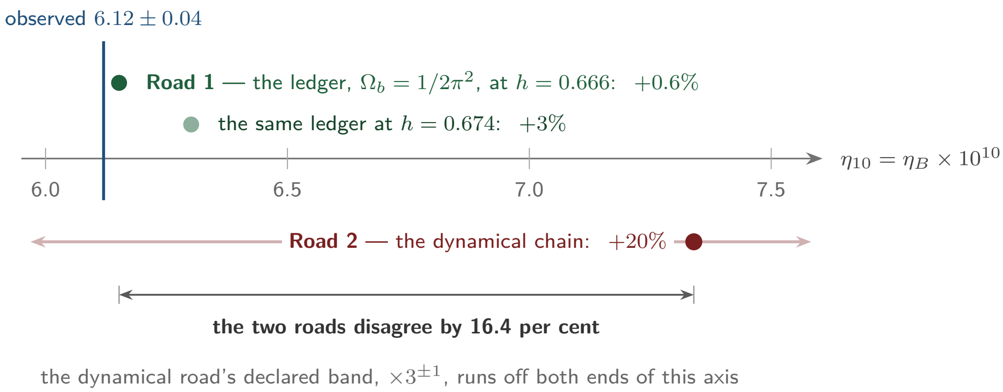

Figure 22: Two roads to one number. η_B, ρ_b and Ω_b are not three predictions: they are one quantity in three sets of units, tied by the exact conversion $\eta_{10} = 273.9\,\Omega_b h^2$, which contains no physics of this framework's. What the framework has is two independent roads to that one quantity. The static road counts volumes on the compact space the ledger selects and lands to better than one per cent of the observation at the framework's own expansion rate, with nothing fitted. The dynamical road — the door, the flavored asymmetries, differential washout, sphaleron conversion — lands twenty per cent high with a declared factor-of-three band that runs off this plot in both directions. Both are predictions of the same number, and closing the sixteen per cent between them is unfinished business.

Resonant enhancement requires the mass splitting of the quasi-degenerate heavy states to be comparable to their decay width. At vacuum numbers,

$$\frac{\Delta M}{\Gamma_N} \simeq 1.4 \times 10^{12}, \tag{104}$$

where the resonance wants order unity. Taken at face value the chain is dead on arrival, undershooting by seven or eight orders of magnitude. The gap is bridged by a named hypothesis, H-ϕ_{geo}-1, about a geometric regulator of the finite-temperature flavor measurement, together with an exponent $K/2$ that this study flags as a postulate of the compression rather than a derivation.

The honest headline, and it is this study's own. A chain with no fitted parameters lands within an order-unity factor of a number the geometry already contained — *which is interesting, and is not proof.* Say that sentence, in that order, and let the audience decide what it is worth. An audience that hears the weak joint from the speaker trusts the strong joints; an audience that finds it in the questions does not.

74.10 What to leave on the screen

Four claims, in decreasing order of how well they are established.

1. **Lepton number is a winding number, and the door exists if and only if the winding vanishes.** This is representation theory, it is exact, and it identifies the neutral lepton as the unique site of lepton-number violation in the Standard Model's field content. The quark sector is excluded not because it is small but because it has no door.

2. **The Majorana parity $\alpha_{31} = \pi$ is the letter that carries the asymmetry, and it is *CP*-conserving in the flavor sum and *CP*-violating in the flavor differences.** Switching it off costs a factor of twenty. Switching off the Dirac phase costs two parts in a thousand.

3. **The baryon fraction is $1/2\pi^2$, and converts to the observed baryon-to-photon ratio at the half-per-cent level** with no fitted quantity in the chain. Ordinary matter is the matter that winds: its share of all matter is one turn, $1/2\pi$.

4. **The dynamical chain reaches the same number to within an order-unity factor**, resting on one open hypothesis about a finite-temperature regulator, and disagreeing with the static road by sixteen and a half per cent. This is the part that is not finished.

And one sentence for the last slide. *Ordinary matter is the matter that winds; the neutrino is the only particle the geometry allows to turn around; and we are the residue of three asymmetries that summed to zero and were not washed out equally.*

75 Division 2: A lecture for everyone — why is there anything?

How this Division differs. No equations are used, no prior physics is assumed, and every number is given in words before it is given in digits. The physics is not softened — the open problem is stated as openly here as in Division 1 — but the vocabulary is. A speaker working from these pages can deliver them essentially as written.

75.1 Start here: you are made of leftovers

Everything you have ever touched is a leftover.

In the first fraction of a second after the beginning, the universe was hot enough that matter and antimatter were being made constantly, in pairs, and destroyed constantly, in pairs. Every particle had an opposite — same mass, opposite charge — and when a particle met its opposite, both vanished into light.

As the universe cooled, the making stopped and the destroying continued until there was nothing left to destroy. If matter and antimatter had been made in exactly equal amounts, that would be the end of the story: a universe of light, and no one in it.

Instead, for roughly every billion pairs that annihilated, one particle of matter had no partner to annihilate with. That one particle in a billion is you. It is also the Earth, the Sun, and every star you can see.

The number, in the plainest terms. Count the particles of light left over from the beginning — they are still here, and we measure them directly as the faint microwave glow of the sky. For every 1.6 billion of them, there is one particle of ordinary matter. That ratio is measured to better than one per cent, by two completely independent methods that agree.

The question of this study is not *what* that number is. We know what it is. The question is *why it is that number* — and, before that, why it is not zero.

75.2 Why "not zero" is the hard part

Here is the difficulty, and it is worth dwelling on because it is the reason the problem has been open for sixty years.

The laws of physics, as we have them, are very nearly indifferent to the difference between matter and antimatter. Run the equations forward from a perfectly balanced start and you get a perfectly balanced finish. To end up with an excess of matter, three things must all be true, and they were written down by Andrei Sakharov in 1967:

1. There must be some process that can change the amount of matter — not just move it around, but create a net amount of it.
2. The laws must not be perfectly even-handed between matter and antimatter. There must be some slight preference.
3. The universe must be out of balance while it happens — expanding and cooling fast enough that the process cannot simply run backwards and undo itself.

The Standard Model of particle physics — the most precisely tested theory in the history of science — has all three, technically. It just has them in amounts that are far too small. Its preference for matter over antimatter, quantified honestly, is at least ten *billion* times too weak — and by the more careful accountings, a hundred times weaker still. Not close. Not within reach of a correction.

So something is missing, and it has to be something rather large.

75.3 Two kinds of memory

Now the idea this study is built on, and it can be shown with your hands.

Imagine two ways of moving.

The first is walking down a road that only goes one way. You can go forward. You cannot go back. And crucially, if you drop something along the way, it is gone — there is no returning to find it. Whatever you are carrying at the end is less than what you started with.

The second is walking around a circular track. You can go forward or backward, you always end up where you started, and there is a quantity that is exactly, unarguably preserved: *how many times you went around*. Not approximately. Not to within measurement error. Three laps is three laps. You cannot smoothly turn three laps into two and a half.

The whole framework in six words. *The circle remembers. The road forgets.*

The proposal is that the universe has both, hidden in directions we cannot see directly — one extra direction shaped like a road, and one extra direction shaped like a circle. Everything in physics that is exactly conserved, that never changes by even the tiniest amount, is protected by the circle:

it is a lap count, and lap counts are whole numbers. Everything that drifts and changes as you look at it more closely lives on the road.

Electric charge is a lap count. That is why it never changes.

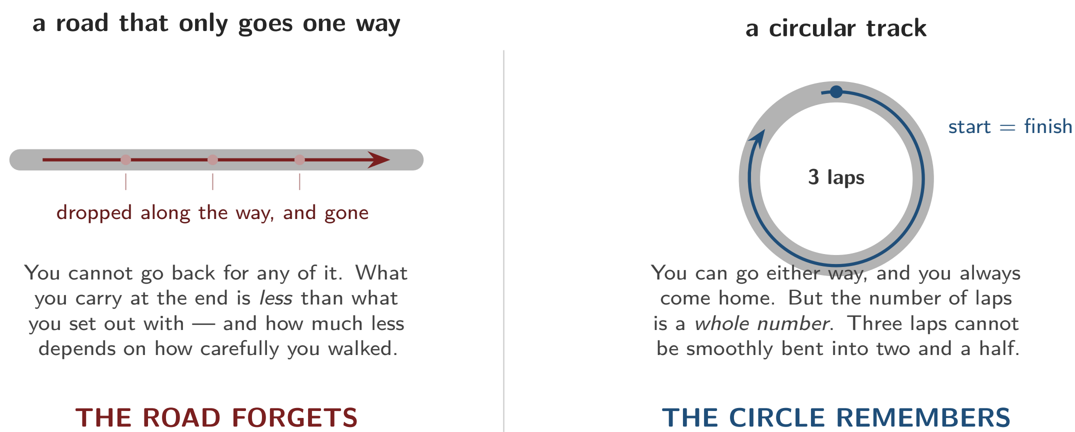

Figure 23: Two kinds of memory. The proposal is that the universe carries two directions we do not see: one shaped like a one-way road, one shaped like a circle. Everything in physics that drifts as you look more closely — the strengths of the forces, the masses — lives on the road, where information is genuinely lost. And everything that is exactly, unarguably conserved — electric charge above all — is a lap count on the circle. Lap counts are whole numbers, and a whole number cannot change by a little bit. That is the entire reason charge is conserved exactly rather than very nearly.

75.4 A particle that is its own opposite

Every particle you have heard of has a distinct opposite. The electron has the positron. The proton has the antiproton. They are different objects.

But there is one exception allowed by the rules, and it is a strange one: a particle that is its own opposite. Its own antiparticle. If such a thing existed, then two of them could meet and annihilate — and in doing so they would change the total amount of matter in the universe, which is exactly the first thing on Sakharov's list.

Which particles can do this? This is where the circle earns its keep, and the answer is sharper than "probably the neutrino".

A particle can be its own opposite **only if it goes around the circle zero times**. That is the entire condition. If a particle winds around the circle even once, it and its opposite wind in opposite directions, they can never be the same object, and there is nothing to be done about it.

And now the punchline: winding around the circle *is* electric charge. So the condition "goes around zero times" means "has no electric charge".

The one slide worth building the talk around. Only an electrically neutral particle can be its own antiparticle. In all of known matter there is exactly one candidate: the **neutrino**. Not the best candidate. The only one. The electron cannot, the proton cannot, the quarks cannot — not because they are unpromising, but because they are charged, and charge is a lap count that cannot be unwound.

This also disposes of the most natural competing idea. Physicists have known since the 1960s that quarks have a small preference for matter over antimatter, and it is tempting to think the problem is just that the preference is too small. It is too small — but that is the lesser objection. The real objection is that quarks are charged, so no amount of preference in the quark sector could ever have produced a net excess. There is no door there to walk through.

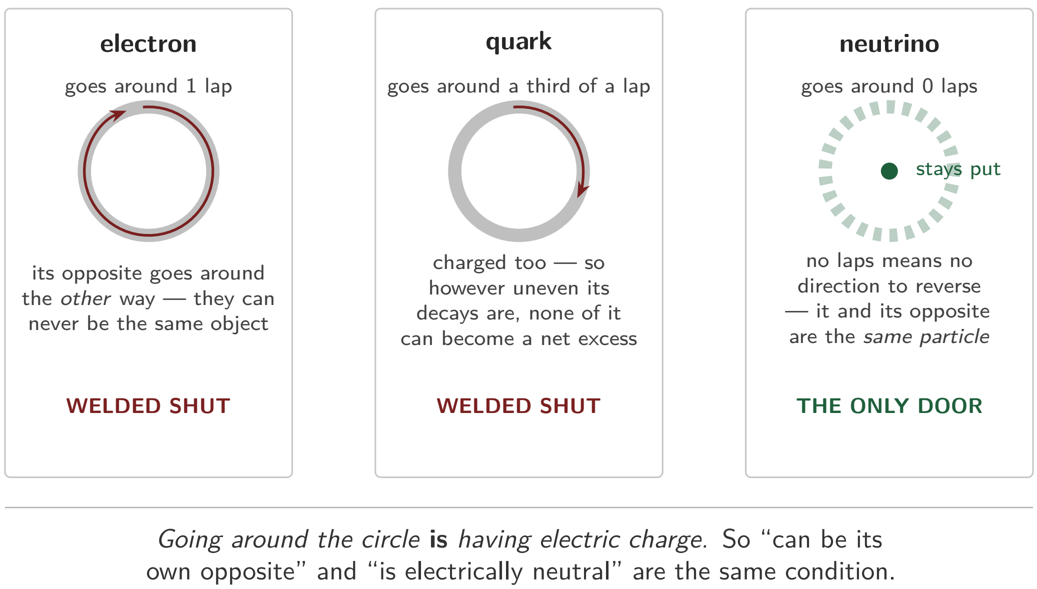

Figure 24: Only a neutral particle can be its own opposite. A particle and its antiparticle wind around the hidden circle in opposite directions. If the winding is anything other than zero, the two windings differ, and no amount of rearranging can make the particle identical to its own antiparticle. Only a particle that does not wind at all escapes the argument — and winding around this circle is precisely what electric charge is. The electron is excluded, every quark is excluded, and exactly one particle in known matter survives: the neutrino. This is why the search for the origin of matter runs through neutrino physics and not through the quark sector, and the reason is structural rather than a matter of degree.

75.5 The bias, and why it is so small

So: heavy relatives of the neutrino, sitting at an energy the framework had already fixed for unrelated reasons — about four thousand times the mass of a proton, which is within reach of the Large Hadron Collider — decay as the universe cools past them.

They decay slightly unevenly. Slightly more matter than antimatter.

How unevenly? Here is the part that surprises people, and it is worth being honest about. The unevenness is not one part in a billion. It is much, much smaller than that at the moment it happens. And yet it survives to become one part in a billion, because almost everything else annihilates away and this residue does not.

There is one more twist, and it is the subtlest idea in the talk. The neutrino comes in three varieties, and the decays produce an excess in each of the three. *Added together, those three excesses cancel exactly to zero.* Nothing at all. If that were the end of the story there would be no matter and no us.

What saves it is that the three varieties are not destroyed at the same rate afterwards. Two of them are efficiently scrubbed away; the third — the one associated with the electron — leaks through, because of a particular small number in how neutrinos mix that this framework also derives. Ninety-six per cent of the matter in the universe came through that one leaky door.

The image to leave them with. Three buckets, filled from a tap that pours matter into one and antimatter into the others, in amounts that add to nothing. Then somebody empties two of the buckets and forgets the third. What is left in the third bucket is everything that exists.

75.6 One turn of a circle

There is a second, completely separate route to the same answer, and it is the prettiest thing in this study.

Forget the decays. Forget the early universe entirely. Just ask: of all the matter there is, what fraction is the ordinary kind — the kind made of atoms, the kind you are made of?

The measured answer is about sixteen per cent. The other eighty-four per cent is dark matter, which has mass and gravity but does not form atoms, stars, or people. (Matter of both kinds is in turn about a third of everything; the remaining two thirds is dark energy, which is a different question.)

In this framework the shape of the hidden space is a sphere — not the two-dimensional surface of a ball, but its three-dimensional analogue — and matter travels along circles wrapped inside it. The fraction of that sphere taken up by one circle is a pure piece of geometry, with no physics in it at all: it is one over two pi, which is sixteen per cent. That is the ordinary share of all the matter there is. Expressed as a share of *everything*, it is five per cent.

The sentence for the poster. *Ordinary matter is the matter that winds. Its share is one turn of a circle — one part in 2π. That is why the universe is five per cent ordinary and ninety-five per cent something else.*

Converted into the same currency as the billion-to-one number, this geometric count agrees with the measurement to better than one per cent, and nothing was adjusted to make it fit. That is a stronger result than the decay chain gives, and it deserves to be said first.

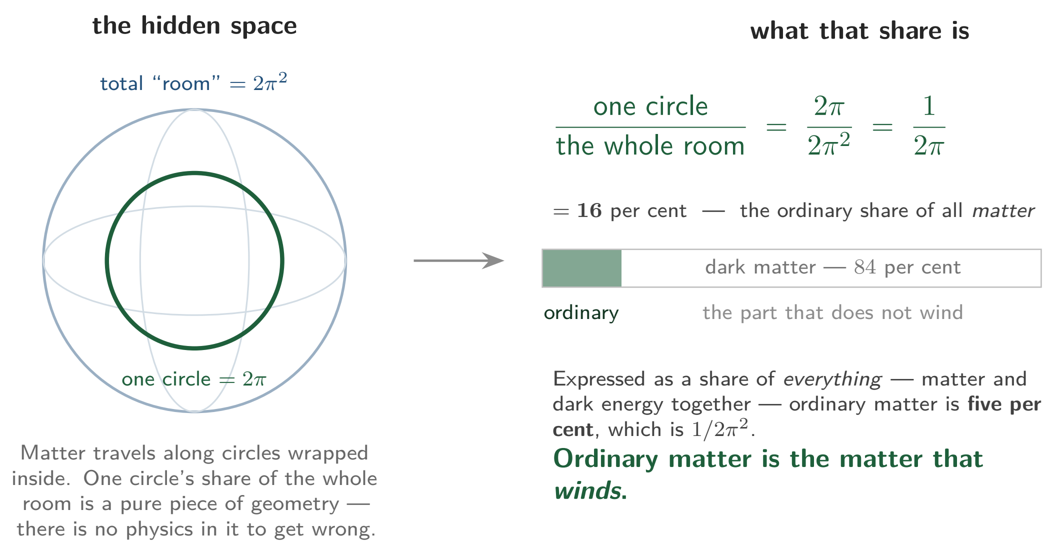

Figure 25: One turn of a circle. A second route to the same answer, with no early-universe dynamics in it anywhere. A frozen, stationary measure on a compact space gives each sector the share of the volume it occupies. The whole is a three-sphere of volume $2\pi^2$; matter runs along a single circle of circumference 2π; and ordinary matter is the part that carries one unit of charge per revolution. The ratio is $1/2\pi$ — one full turn — which is sixteen per cent of all matter, and $1/2\pi^2$, or five per cent, of everything. Converted into the same currency as the billion-to-one number, this count agrees with the measurement to better than one per cent, with nothing fitted anywhere in the chain.

75.7 What is honest to say, and what is not

A talk like this is easy to oversell, so here is the part to read carefully rather than skip.

Three things that are not yet settled, said plainly.

First, the decay chain relies on one assumption that has not been derived — a technical statement about how a certain quantity behaves in the hot early universe. Without it the chain fails by a factor of ten million. With it the chain works. That assumption is named, it is written down, and it is stated as an assumption rather than a result.

Second, the two routes to the answer — the decay chain and the geometric count — give numbers that differ from each other by about sixteen per cent. They ought to agree exactly. Making them agree is unfinished work.

Third, the neutrino has not actually been shown to be its own opposite. The experiment that would show it looks for an extremely rare form of radioactive decay, and this framework predicts a rate roughly ten times below what the next generation of experiments will reach. If the neutrino turns out *not* to be its own antiparticle, this entire account is wrong, and that is exactly as it should be.

And one thing that is often said in talks like this and should not be: the theory does *not* predict the quantity that DUNE and Hyper-Kamiokande are being built to measure and then use it to get the matter asymmetry. It predicts that quantity, separately, and the matter asymmetry does not depend on it. If that prediction turns out wrong, the account of why we exist is untouched. Two claims, two independent ways to be wrong. That is a virtue, and it should be presented as one.

75.8 What we will know, and when

End on the calendar, because an audience wants to know when they find out.

- **Within about five years** — the long-baseline neutrino experiments report on how neutrinos mix and whether they prefer matter. This tests one of the framework's predictions and, notably, *not* the one about our existence.
- **Within about ten years** — cosmological surveys will pin down the total mass of the neutrinos. This framework says it is fifty-nine thousandths of an electron-volt. The measurements are closing in on precisely that value from above, and this is the tightest live test.
- **Within about fifteen years** — the next generation of rare-decay experiments looks for the neutrino being its own opposite. A detection would confirm the door exists. A non-detection at the level this framework predicts would not yet rule it out, which is an uncomfortable thing to have to say and is said anyway.
- **Whenever the collider gets there** — the heavy relatives sit at about four thousand proton masses. That is not an unreachable energy. It is roughly where we are already looking.

The closing, in four sentences. Everything you are made of is the one part in a billion that had no partner to annihilate with. The reason there was a leftover at all is that one particle — the neutrino — is allowed to be its own opposite, and it is allowed because it carries no charge, and charge is a number of turns around a hidden circle. The excess was written by a quantity no experiment now planned can ever see, and read out through the one of three doors that nobody closed. *We are the part that got away.*

75.9 Speaker's notes: three sentences not to get wrong

1. **Do not say the neutrino "breaks" *CP* because its phase is** 180 **degrees.** A phase of 180 degrees is, on its own, a *symmetric* value, and a physicist in the audience will say so. The correct statement is that the three flavor contributions cancel exactly when added, and survive because they are not erased equally. Cancellation plus unequal erasure — that is the mechanism, and it is more interesting than a phase anyway.
2. **Do not say the theory "explains" the matter asymmetry.** Say it reaches the observed number by two independent routes, one of them accurate to better than one per cent, and that the routes do not yet agree with each other. An audience respects a stated gap far more than it respects a claim it later discovers was hedged.
3. **Do not say the quark sector's *CP* violation is "too small".** Say the quark sector has no door. It is the sharper statement, it is exactly true, and it is the one people remember.

Part XVI
For the Curious Reader

Fourteen questions, answered without machinery

This Part is for the reader who wants the physics and not the apparatus. Nothing here is new; everything is a restatement of something above, in language that presumes no more than an interested non-specialist brings. Each answer names the Part where the machinery lives.

76 The questions

Q1. Why is there any matter at all? Shouldn't it all have annihilated?

Very nearly all of it did. In the first fraction of a second, quarks and antiquarks were being made and destroyed in equal numbers; when the universe cooled past the point where making them was possible, they found each other and annihilated into light. What is left over — everything we are made of — is a residue of about one part in a billion and a half. The question is not why so much matter survived but why the books did not balance exactly. Something, early on, made one extra quark per billion quark–antiquark pairs. (Part I.)

Q2. What is "*CP* violation" in one sentence?

It is the fact that the laws of physics are not quite the same for matter and antimatter: swap every particle for its antiparticle *and* look at the result in a mirror, and a few rare processes happen at slightly different rates. Without that failure, nothing could have made an excess of one over the other. (Part III.)

Q3. Where does this theory say *CP* violation comes from?

From a pattern pressed into an extra dimension that is shaped like a circle. The pattern has ten marks at even spacing, and it does not match itself when read backwards — like a spiral staircase, which is not the same going up as coming down. That mismatch *is* *CP* violation. The mathematics reduces to one line: the spacing is set by the number five, one particular angle in the pattern comes out at three-fifths of a half-turn, and three-fifths of a half-turn is not a whole number of half-turns. If it had been, matter and antimatter would behave identically and there would be no universe to ask about it. (Part III, Eq. (19).)

Q4. Why doesn't the strong nuclear force violate *CP* too?

Because it writes a different figure on the same circle. Where the weak sector writes a five-sided figure whose circuit leaves a leftover angle, the strong sector writes a triangle whose three steps come back exactly to the start. Nothing is left over, so there is no mismatch, so no *CP* violation. And the same fact — that color must come back to where it started — is why quarks are permanently confined inside protons and neutrons. Two famous facts, one geometric reason. (§15.)

Q5. Can I picture the weak interaction as a corkscrew?

You can picture it that way and it is a fair picture, but it is not the mechanism, and this study tested that carefully because the picture is so appealing. If the phase were really being wound up as a particle cascades down in energy, then two different cascades over the same range of energy would wind up by the same amount. They do not: the quark chain turns about nine degrees per step of energy and the lepton chain about twenty-four — nearly three times as much over nearly the same range. So the corkscrew shows you *where* the phase lives; it does not make it. The phase was already fixed, by the pattern on the circle, before any particle decayed. (Part IV.)

Q6. Why can't quarks be responsible for the excess of matter?

Two reasons, and either would be enough. First, all quark *CP* violation is priced by a single number — the area of a certain triangle — and that area is three hundred times smaller than the corresponding lepton one. Second, and worse, the effect must switch off if any two quarks of the same charge have the same mass, so it is multiplied by the differences of all their masses; the quarks span five orders of magnitude in mass and the whole product ends up divided by the twelfth power of the temperature. The result is ten to twelve orders of magnitude too small. The Standard Model, taken at its word, predicts a universe with no matter in it. (Part VI.)

Q7. So neutrinos did it. Why can only they do it?

Because a neutrino is the only particle that can be its own antiparticle. In this theory every particle carries a whole number that counts how it wraps around the extra circular dimension, and that number is what lepton number *is*. To turn a particle into its own antiparticle you need that number to be zero, and it is zero for exactly one species: the neutral one. Everything else — every electron, every quark — carries a non-zero wrap and is therefore locked to its own identity. So it is not that neutrinos are the most convenient option. They are the only one the geometry permits. (Part VII.)

Q8. What actually happened, and when?

About 10^{-13} seconds after the beginning, when the universe had cooled to about four trillion electron-volts, a set of heavy neutral particles fell out of equilibrium and decayed. Their decays were very slightly biased toward matter — by about one part in forty million — and the bias was set by a fixed geometric angle rather than by any adjustable property of the particles. What was produced was an excess of leptons, not of quarks. Then a process that was still running at that temperature converted about a third of that lepton excess into a quark excess, and shut off shortly afterwards, freezing the answer in. (Parts VIII and IX.)

Q9. Four TeV is a scale the LHC can nearly reach. Isn't that suspiciously convenient?

It would be, if the scale had been chosen. It was not. The number four TeV comes out of an integer, 35, that this framework fixed years earlier while auditing an entirely different consistency condition, and the same 35 independently sets the frequency of a predicted dark-matter signal and the energy scale of the vacuum. Nothing in its derivation mentions matter, antimatter, or neutrinos. In every other version of this mechanism in the literature, the TeV scale is chosen precisely because it is testable — which is honest but tells you nothing. Here it arrived from somewhere else, which means it can be wrong, and the high-luminosity LHC will find out. (§53.)

Q10. Can we see the heavy neutrinos at the LHC, then?

No, and it is important to be clear about that. They couple to ordinary matter about a million times more weakly than the electron does, so they cannot be produced in any useful quantity. What *is* visible is the rest of the furniture at that scale: this framework predicts three other states there, and one of them — a heavy partner of the gluon, decaying to top quarks — should be within reach of the high-luminosity LHC. If that cluster is absent, the whole story fails, regardless of anything the neutrino experiments find. (§62.)

Q11. Everyone says DUNE and Hyper-Kamiokande will measure the neutrino *CP* phase. Is that the phase that made us?

No, and this is the most surprising thing in the study. There is more than one *CP* phase in the neutrino sector. The one those experiments will measure — the "Dirac" phase — turns out to be a spectator in the matter chain. The one that actually drove the excess is a "Majorana" phase, and Majorana phases cancel out of oscillation probabilities exactly, so no oscillation experiment ever built or proposed can see it. The only laboratory that can is neutrinoless double beta decay. This theory does predict the Dirac phase, at 196 degrees, and that prediction is a good independent test — but it is a test of the alphabet, not of the mechanism. (§33.1.)

Q12. How can a phase of exactly 180 degrees violate *CP*? Isn't 180 degrees a symmetric value?

It is — for an ordinary particle. But the process that turns a neutrino into its own antiparticle involves the field twice rather than once, so it sees *half* the angle, and 180 degrees becomes 90. That is a real and standard fact, and by itself it is *not* the answer — a symmetric-looking phase remains symmetric no matter how it is halved.

The answer is that "the asymmetry" is three numbers, not one. The neutrino comes in three flavors, the decays bias all three, and the three biases *cancel exactly when added together*. That cancellation is what the symmetric value of 180 degrees is telling us, and it is real. What it does not tell us is that each of the three is zero — they are not. Afterwards, two of the three flavors are efficiently scrubbed away and the third is not, so a set of biases that summed to nothing leaves a residue behind. The residue is the matter. (§33.1, §39.)

Q13. If the universe can violate baryon number, why doesn't the proton decay?

Because it can only violate it three at a time. There are three families of matter, each contributing once, so the smallest possible change in the number of protons and neutrons is three, not one. A hot early universe processing enormous numbers of particles never notices the restriction. A single proton sitting in a glass of water does: to decay it would have to change the count by one, and one is not a multiple of three. So the proton is not merely long-lived — it is permanently stable, and so is the neutron against turning into an antineutron. This is the sharpest difference between this framework and conventional grand unified theories, which predict that protons *do* decay. If a single proton decay is ever seen, in any channel, this framework is finished. (Part X.)

Q14. What is the single most likely way this turns out to be wrong?

Two candidates, and we would name them in this order.

The nearest-term one is the total mass of the neutrinos. This framework predicts 59 thousandths of an electron-volt, and it cannot move: the number falls out of two other scales and there is nothing to adjust. Some current cosmological analyses already set an upper bound slightly *below* that, though the answer depends on assumptions about dark energy that are themselves contested. Better data over the next few years will settle it, and it may well settle it against the framework.

The deeper one is a calculation nobody has done. The mechanism needs the heavy neutrinos to sit at a particular degree of near-degeneracy — not exactly degenerate, which would give too much matter, and not as far apart as their vacuum properties suggest, which would give far too little. The framework names the required value, explains where it might come from, and lists three ways the calculation could come back and destroy the chain. Until someone does that calculation, this is a well-posed conjecture with a number attached, and not a finished result. We think that is the honest description, and we would rather give it than a confident one. (§38, §61.)

Part XVII
Technical Appendices

Dictionary, worked problems, and a numerical audit in which every figure typeset in this study is recomputed from its definition

77 Appendix A. Dictionary of symbols and names

78 Appendix B. Worked problems

Seven problems, each solvable in a few lines, chosen because each isolates one load-bearing step. Solutions follow immediately.

B.1. Show that *CP* conservation requires every loop holonomy to lie in $\{0, \pi\}$. *Solution.* CP acts on the flavor texture by complex conjugation, sending each edge angle $\varphi \mapsto -\varphi$. The theory is invariant if and only if the conjugated texture is a per-family re-zeroing of the original — because re-zeroings are exactly the unphysical redefinitions. Loop holonomies are invariant under re-zeroings, so invariance requires $-\varphi_{\text{loop}} \equiv \varphi_{\text{loop}}$ modulo 2π for every loop, i.e. $2\varphi_{\text{loop}} \equiv 0$, i.e. $\varphi_{\text{loop}} \in \{0, \pi\}$. Those are the only two angles that are their own negatives.

B.2. Count the physical phases of the $N \times N$ bipartite flavor graph, and check the answer at $N = 3$. *Solution.* There are N^2 edge angles and $2N$ vertex dials, but the global common shift acts trivially, so the effective gauge freedom is $2N - 1$. The physical count is $N^2 - (2N - 1) = (N - 1)^2$. At $N = 3$ this is 4, of which unitarity leaves one CP-odd invariant and three magnitudes — exactly the four parameters of the CKM matrix. At $N_g = 2$ it is 1, and that single loop angle can be rotated to 0 or π, so two generations cannot violate CP — the textbook result, here as a corollary of B.1.

B.3. Show that the color triangle's holonomy is trivial while the flavor pentagon's is not. *Solution.* Color transport composes three steps of $2\pi/3$: $3 \times 2\pi/3 = 2\pi \equiv 0$, so $W(\triangle) = 1 \in \{0, \pi\}$, and by B.1 CP is conserved. The flavor circuit is built from the branch phase $3\pi/5$; five steps of $2\pi/5$ likewise sum to 2π, but the branch angle itself is $3\pi/5$, which is not in $\{0, \pi\}$, so the reading leaves a residue. The distinction is between the *lattice* closing (both do) and the *content* being achiral (only color is).

B.4. Derive $\delta_{\text{CKM}}^{\text{QD}} = 13\pi/35$ from $\delta_{\text{CKM}}^{(1)} = 2\pi/5$ and show it is exactly one quantum-defect step. *Solution.* Put both on a common denominator: $2\pi/5 = 14\pi/35$. One step of the (π/K)-lattice with $K = 35$ is $\pi/35$. So $14\pi/35 - \pi/35 = 13\pi/35 = 66.857^\circ = 1.1670$ rad. Against the world fit 1.147 ± 0.026 the pull is $(1.1670 - 1.147)/0.0283 = +0.70$. Note the two different widths in

Table 22: Dictionary of symbols used in this study, with the owning paper where the symbol is defined. Where a symbol has a value the value is given, so the table doubles as the study's numerical summary; where it has an owner the owner is named, so that a reader who wants the derivation rather than the number knows which paper to open.

Symbol	Name	Meaning, and owner
$\mathcal{M}_6$	the bulk	$\mathcal{M}_4 \times I_r \times S^1_\theta$ (GG-W2-Holonomy)
I_r	the scale interval	the fifth dimension, $[0, K]$, coordinate $r = \ln(E_{\rm GUT}/\mu)$; carries no holonomy
S^1_θ	the phase circle	the sixth dimension; carries all of them
$(M, N, L; J, K)$	the ledger	$(1, 3, 5; 7, 35)$, $J = 2N+1$, $K = LJ = M^2+N^2+L^2$ (GG-A2-GG6, GG-A3-BI6)
$E_{\rm GUT}$	the seam	$M_{\rm Pl}e^{-2\pi} = 2.28 \times 10^{16}$ GeV
$M_{\rm lock}$	the lock scale	$\frac{1}{2}M_{\rm Pl}e^{-K} = 3.85$ TeV (GG-A2-GG6)
φ_{12}	the branch phase	$\pi(1 - 2/L) = 3\pi/5$, the pentagon's interior angle
$\delta_{\rm CKM}$	quark Dirac phase	$2\pi/5$ at first order, $13\pi/35$ with the QD completion
$\delta_{\rm PMNS}$	lepton Dirac phase	$\pi + 2\arcsin(1/J) = 196.43°$
$(\alpha_{21}, \alpha_{31})$	Majorana phases	$(0, \pi)$; parity data of the rank-two block; $\alpha_{31} = \pi$ is achiral in the flavor sum and chiral in the flavor differences
$J_{\rm CKM}$, $J_{\rm PMNS}$	Jarlskog invariants	twice the unitarity-triangle area; one area for the whole graph
$m_\star$	the seam neutrino scale	$\pi v^2/E_{\rm GUT} = 8.35$ meV (GG-A9-SM)
$m_{\beta\beta}$	$0\nu\beta\beta$ effective mass	1.72–1.79 meV
Y_ν	neutrino Yukawa	$\sqrt{\pi M_{\rm lock}/E_{\rm GUT}} = 7.3 \times 10^{-7}$
$\epsilon^{\rm res}$	resonant decay asymmetry	2.5×10^{-6}; the H-$\phi_{\rm geo}$-1 target value
κ, $K_{\rm wash}$	washout efficiency, strength	$\kappa \simeq 1/K_{\rm wash}$; $K_{\rm wash} \simeq 50$ (strong washout)
$\Gamma^{\rm eff}$	the effective width	the regulator of H-$\phi_{\rm geo}$-1; candidate $\Gamma_\tau = c_\tau y_\tau^2 T$
$c_{\rm sph}$	sphaleron transfer	28/79 for one Higgs doublet (Harvey–Turner)
g_*	relativistic d.o.f.	106.75 at the TeV bath
η_B	baryon-to-photon ratio	$(n_B - n_{\bar{B}})/n_\gamma$
Σm_ν	neutrino mass sum	59.17 meV, derived
GRIP-In, GGA	descent, ground attractor	the cascade machinery (GG-A6-GRIP, GG-P2-N3, GG-P3-Electron)
H-$\phi_{\rm geo}$-1	the named hypothesis	partial resonance, $\Delta M/\Gamma^{\rm eff} \simeq 3 \times 10^4$
C-P5-C	the parent conjecture	sphaleron-transport normalization; clauses (i) resonance and (ii) the $e^{-K/2}$ throughput

play: at the QD point the residual theory width is a fraction of a lattice step, so $\sigma_{\rm tot} = 0.0283$; at the first-order point the full step $\pi/35 = 0.0898$ applies in quadrature, giving $\sigma_{\rm tot} = 0.0935$ and a pull of $+1.17$ for $2\pi/5$. Against experimental error alone $2\pi/5$ sits at $+4.2\sigma$, which is why the QD completion carries a flag and not a footnote.

B.5. Show that a Majorana phase of π enters an amplitude as i. *Solution.* The $\Delta L = 2$ bilinear is quadratic in the neutrino field, so a field redefinition $\nu \to e^{i\beta}\nu$ shifts the mass term's phase by 2β. Equivalently, a mass-term phase α is carried into amplitudes as $e^{i\alpha/2}$ — the half angle. With $\alpha_{31} = \pi$, $e^{i\pi/2} = i$.

And a warning that belongs with the exercise. This factor of i is kinematics, not *CP* violation. A Majorana phase of π is the standard signature of opposite intrinsic *CP* parity between mass eigenstates, and opposite intrinsic parity is a CP-conserving configuration — which is why the flavor-summed asymmetry vanishes here exactly. The resolution of the paradox is in the flavor *differences*, not in the half angle: see §33.1.

B.6. Show that the electroweak anomaly cannot kill a single proton. *Solution.* The one exact non-conservation of baryon number is $\partial_\mu J_B^\mu = N_{\rm gen}\,(g^2/32\pi^2)\,\mathrm{tr}\,W\widetilde{W}$: each family contributes once, so any topological transition changes B by $N_{\rm gen}\,n = 3n$. A proton decay is $\Delta B = -1$, and $-1 \notin 3\mathbb{Z}$, so no anomalous process — however violent — can remove a single nucleon. The minimal event costs three baryons at once, and three leptons with them, preserving $B - L$. The rate is in any case suppressed by the instanton action $S = 8\pi^2/g^2 \approx 187$, but the selection rule is the structural point: it ties proton safety to the same $N_g = 3$ that fixes the family count.

B.7. Estimate the entropy per baryon and say why it matters. *Solution.* $s/n_B = (s/n_\gamma)/\eta_B = 7.04/6.1 \times 10^{-10} \approx 1.2 \times 10^{10}$. Ten billion units of entropy per particle of matter. The consequence is that nucleosynthesis and recombination happen far "later" — at temperatures far below the relevant binding energies — because with a billion photons per baryon there are still enough photons in the high-energy tail to dissociate deuterium, and later hydrogen, long after the mean photon energy has fallen below threshold. The number is η_B read upside down.

79 Appendix C. Numerical audit

Every quantity typeset anywhere in this study, recomputed from its definition. The "recomputed" column is generated programmatically from the definitions in the middle column by the script `s2_audit_table.py`, so it cannot drift from the arithmetic; the "as quoted" column is what the parent papers state. Discrepancies, where any, are the layer spreads discussed in the text and not errors.

Notes on three entries

$J_{\rm PMNS}$. The recomputed value -9.48×10^{-3} uses the NuFIT 6.0 mixing angles; GG-P5-CPT quotes -9.4×10^{-3} and GG-A9-SM -9.26×10^{-3} at its own first-order angles. The spread is the evaluation layer, not an error, and the ratio to $J_{\rm CKM}$ is 300 to 304 across the layers.

δ_{CP} *at the electroweak epoch.* The three rows are one quantity at three temperatures, and the two-and-a-half-order spread is the T^{-12} of Eq. (33). Any single number quoted for "the Standard Model's CP measure" is a number at an unstated temperature.

$m_\star^{\rm wash}$. Evaluated at $\bar{v} = v/\sqrt{2} = 174$ GeV, which is the leptogenesis literature's convention and the one GG-P5-CPT uses; evaluating at $v = 246$ GeV would double it and halve $K_{\rm wash}$. The convention is stated because the factor of two propagates into κ.

Table 23: Numerical audit. Constants used: $M_{\rm Pl} = 1.22091 \times 10^{19}$ GeV, $M_{\rm Pl}^{\rm red} = 2.435 \times 10^{18}$ GeV, $v = 246.22$ GeV, $\bar{v} = v/\sqrt{2}$, $g_* = 106.75$; quark masses $\overline{\rm MS}$; lepton mixing angles NuFIT 6.0 normal ordering with the atmospheric sample; light neutrino masses from GG-A9-SM's quantum-defect layer.

Quantity	Definition	Recomputed	As quoted
K	$LJ = M^2 + N^2 + L^2$	35	35
$\alpha^{-1}(0)$	$2^J + N^2$	137	137
$E_{\rm GUT}$ [GeV]	$M_{\rm Pl} e^{-2\pi}$	2.28×10^{16}	2.28×10^{16}
$M_{\rm lock}$ [GeV]	$\frac{1}{2} M_{\rm Pl} e^{-K}$	3849.0	3849 / 3.85 TeV
$r(M_{\rm lock})$	$\ln(E_{\rm GUT}/M_{\rm lock})$	29.410	29.41
r (EW seam)	$\ln(E_{\rm GUT}/100\,{\rm GeV})$	33.06	33.1
$r(T_{\rm rh})$	$\ln(E_{\rm GUT}/3{\times}10^{10})$	13.54	13.5
m_Λ [meV]	$M_{\rm lock}^2/M_{\rm Pl}^{\rm red}$	6.084	6.1
φ_{12} [deg]	$\pi(1 - 2/L)$	108.0	108
$\delta_{\rm CKM}^{(1)}$ [deg]	$2\pi/5$	72.0	72.0
$\delta_{\rm CKM}^{\rm QD}$ [rad]	$13\pi/35$	1.16688	1.16688
$\delta_{\rm CKM}^{\rm QD}$ [deg]	$13\pi/35$	66.86	66.86
$\sin\theta_{12}$	$\lvert e^{-\pi/2} - e^{i3\pi/5} e^{-\pi}\rvert$	0.225019	0.225019
closed form	$\sqrt{e^{-\pi} + e^{-2\pi} + \phi_g^{-1} e^{-3\pi/2}}$	0.225019	0.225019
$\delta_{\rm PMNS}$ [rad]	$\pi + 2\arcsin(1/J)$	3.42829	3.42829
$\delta_{\rm PMNS}$ [deg]	$\pi + 2\arcsin(1/J)$	196.43	196.43
$\sin\delta_{\rm PMNS}$	$-8\sqrt{3}/49$	-0.2828	-0.283
$J_{\rm CKM}$	moduli $\times \sin\delta_{\rm CKM}^{\rm QD}$	3.115×10^{-5}	3.1×10^{-5}
$J_{\rm PMNS}$	NuFIT 6.0 angles $\times \sin\delta_{\rm PMNS}$	-9.480×10^{-3}	-9.4×10^{-3}
$\lvert J_{\rm PMNS}\rvert/J_{\rm CKM}$		304	$\simeq 300$
δ_{CP} at $T = 100$ GeV	Eq. (33)	1.19×10^{-19}	
δ_{CP} at $T = 131.7$ GeV	Eq. (33)	4.36×10^{-21}	
δ_{CP} at $T = 160$ GeV	Eq. (33)	4.22×10^{-22}	
$m_\star$ [meV]	$\pi v^2/E_{\rm GUT}$	8.353	8.35
Y_ν^2	$\pi M_{\rm lock}/E_{\rm GUT}$	5.304×10^{-13}	5.3×10^{-13}
Y_ν		7.283×10^{-7}	7.3×10^{-7}
Γ_N [GeV]	$Y_\nu^2 M_N/8\pi$	8.122×10^{-11}	8×10^{-11}
ΔM [GeV]	$M_{\rm lock}/K$	110.0	$\simeq 110$
$\Delta M/\Gamma_N$ (vacuum)		1.35×10^{12}	1.4×10^{12}
$m_\star^{\rm wash}$ [eV]	Eq. (54)	1.070×10^{-3}	1.1×10^{-3}
$K_{\rm wash}$ ($\widetilde{m} = m_3$)	$m_3/m_\star^{\rm wash}$	47.2	$\simeq 50$
$\kappa = 1/K_{\rm wash}$		2.119×10^{-2}	2.1×10^{-2} (adopted 1.0×10^{-2})
$e^{-K/2}$	$e^{-17.5}$	2.5110×10^{-8}	2.51×10^{-8}
$c_{\rm sph}$	$28/79$	0.35443	0.35443
η_B	$7.04\, c_{\rm sph}\,(1.25/g_*)\, e^{-K/2}$	7.3365×10^{-10}	7.34×10^{-10}
deviation [%]	vs 6.12×10^{-10}	+19.88	+20
$m_{\beta\beta}$ [meV]	Eq. (45)	1.794	1.72 / 1.79
Σm_ν [meV]	$m_1 + m_2 + m_3$	59.17	59.17
Δm_{21}^2 [eV^2]		7.3823×10^{-5}	7.382×10^{-5}
Δm_{31}^2 [eV^2]		2.5516×10^{-3}	2.552×10^{-3}
Δr quark	$r(u) - r(b)$	7.568	7.57
Δr lepton	$r(e) - r(\tau)$	8.154	8.15
pitch quark [deg/r]	$\delta_{\rm CKM}/\Delta r$	8.834	8.83
pitch lepton [deg/r]	$\delta_{\rm PMNS}/\Delta r$	24.089	24.09
pitch ratio		2.727	2.73

80 Appendix D. The three script files

Three short programs generate everything computational in this study, and are deposited with the source so that any figure or table can be regenerated:

- `s2_helix_coords.py` — generates the TikZ coordinate files for Figures 7 and 8. It also prints the pitch test of Table 5. Its header records the geometry bug it fixes: an earlier version centered each guide ellipse at the projected height of the *point* rather than of the *slice*, which placed the b quark and the τ at the center of their circles instead of on the rim.
- `s2_audit_table.py` — generates Table 23 directly from the definitions.
- `verify.py` — the standalone check of every number, run before any of them was typeset.
- `s2_boltzmann.py` — the flavored Boltzmann integration of §40: the derived K_α, the validated kernels, the three-level Monte Carlo of §46, and the coordinate files for Figures 14 and 15. The random seed is fixed at 35 — which integer else — so every number in Part IX regenerates exactly.

Part XVIII
References

References

ATLAS Collaboration. *Search for heavy Majorana or Dirac neutrinos and right-handed W gauge bosons in final states with charged leptons and jets in pp collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **83**, 1164 (2023), arXiv:2304.09553.

Relevance. The strongest direct bounds on TeV-scale heavy Majorana neutrinos — W_R excluded below 6.4 TeV, N_R below 3.5–3.6 TeV at $m(W_R) = 4.8$–5.2 TeV. They constrain the Keung–Senjanović production mechanism and are quoted in §62 precisely to show that they do *not* apply to a state produced only through a 10^{-6}-level mixing in a framework with no W_R.

ATLAS Collaboration. *Search for high-mass dilepton resonances in pp collisions at $\sqrt{s} = 13.6$ TeV combined with* 13 *TeV results*, arXiv:2607.28334 (2026).

Relevance. 165 fb^{-1} at 13.6 TeV plus Run 2: $Z'_{\rm SSM}$ excluded below 5.5 TeV, Z'_ψ below 4.8 TeV. Any new neutral gauge boson near 4 TeV with Standard-Model-strength lepton couplings is dead. The Z'_B of the lock cluster is leptophobic by construction and survives — a consistency check, not a rescue, since the leptophobia was fixed by the $\mathrm{U}(1)_B$ assignment long before this bound.

K. Arisaka (2026). *The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry*. Preprint, available at https://preprints.arisaka-gg.org/a10-cosmos/.
doi:10.5281/zenodo.22064314.

Relevance. The cosmological half of the chain. Appendix I is the owner of the heavy triplet at the lock scale, of the transport treatment, and of the flavored Boltzmann system this study integrates; it is carried at *derived-conditional*, and the conditions travel with it. The baryon-to-photon ratio $\eta_B \approx 7 \times 10^{-10}$ is shared with GG-P5-CPT and is carried at D* with the parent's own theory band of $\times 3^{\pm 1}$ — not narrowed here.

K. Arisaka (2026). *GG-Theory with GG-6: Geometry–Gauge Locking at the 6D Bulk (GG-6) at the Heart of Grand Geometric Theory (GG-Theory)*. Preprint, available at https://preprints.arisaka-gg.org/a2-gg6/.
doi:10.5281/zenodo.22064298.

Relevance. The bulk, and the source of the integers. The ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ is fixed here, and so is the lock scale $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-K}$, both at *derived*. Five integers settled upstream in this paper appear in the matter chain of the present study in eleven separate places; none of them is adjusted here, and there is no sixth.

K. Arisaka (2026). *Anomaly Freedom of BI-6: The Modified Bianchi Identity (BI-6) for the Six-Force, Six-Dimensional Bulk (GG-6)*. Preprint, available at https://preprints.arisaka-gg.org/a3-bi6/.
doi:10.5281/zenodo.22064300.

Relevance. The consistency half of the ledger. Anomaly freedom is why the integers imported from GG-A2-GG6 are one structure rather than a family of tunable variants — which is the whole force of the arithmetic in Part I: a chain with no dial cannot be rescued by turning one.

K. Arisaka (2026). *The GRIP on GDOS: GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind*. Preprint, available at https://preprints.arisaka-gg.org/a6-grip/.
doi:10.5281/zenodo.22064306.

Relevance. The descent machinery. GRIP-In and the ground geometric attractor GGA — the two objects the stability arguments of this study stand on — are defined here, with GG-P2-N3 and GG-P3-Electron.

K. Arisaka (2026). *The Standard Model as Theorem of GG-Theory: Its 28 Parameters as Observations of One Six-Dimensional Geometry*. Preprint, available at https://preprints.arisaka-gg.org/a9-sm/.
doi:10.5281/zenodo.22064312.

Relevance. The flavor half of the chain, and the most heavily drawn-on parent. It owns the flavor connection and the statement that loop holonomies are the physical data (T-SM-46.2, *derived*); the geometric seesaw $m_\star = \pi v^2/E_{\rm GUT}$ and the light spectrum with normal ordering as an output, both *derived-conditional* on the Weinberg-coefficient lemma; $\delta_{\rm PMNS} = \pi + 2\arcsin(1/7)$; and the absolute stability of the proton through $\Delta B \in 3\mathbb{Z}$ (T-SM-60.1), a theorem conditional on the (H2) baryon-descent step. Every one of these is carried at the grade printed there.

K. Arisaka (2026). *Why Three Generations: A Three-Leg Derivation of $N_g = 3$ from BI-6, MSMF and GDOS Morse Stability*. Preprint, available at https://preprints.arisaka-gg.org/p2-n3/.
doi:10.5281/zenodo.22064323.

Relevance. Why the ladder stops at three. The budget-saturation argument is the reason the flavor sum of this study runs over three generations and not a free number, and this study inherits that paper's declared conditions on it rather than absorbing them.

K. Arisaka (2026). *The Electron as the Ground State of Nature: The Six Oldest Questions about the Electron, Closed and Fenced within GG-Theory*. Preprint, available at https://preprints.arisaka-gg.org/p3-electron/.
doi:10.5281/zenodo.22064325.

Relevance. The electron's absolute stability, carried here as the no-exit-gate theorem beside the proton's. Cited also for what it does *not* say: neither this paper nor GG-A9-SM assigns a value of r to an individual quark or lepton, which is why the helix of Part IV is labeled a construction original to this study and why its negative result is reported as such.

K. Arisaka (2026). *The Mirrors of Matter: Why Every Particle Has a Twin, Why Nature's Mirrors Break, and Why the Universe Is Made of Matter*. Preprint, available at https://preprints.arisaka-gg.org/p5-cpt/.
doi:10.5281/zenodo.22064329.

Relevance. The immediate parent of the *CP* chain. It owns the pentagonal lattice $(\pi/5)\mathbb{Z}$ (T-P5-3(iv), *derived-conditional* on the rigidity of the holonomy moduli), the branch phase $3\pi/5$ with $\delta^{(1)}_{\rm CKM} = 2\pi/5$, and the quark-democracy completion $\delta_{\rm CKM} = 13\pi/35$, which that paper itself flags as asserted and post-hoc. It is also the owner of the two joints this study is most careful about: the on-resonance condition C-P5-C(i), a **named input**, and the throughput exponent $K/2$ of C-P5-C(ii), a **postulate**. Both flags are carried verbatim.

K. Arisaka (2026). *The Circle That Remembers, the Interval That Forgets: GG-Theory and the Sixteen Problems of Particle Physics and Cosmology*. Preprint, available at https://preprints.arisaka-gg.org/w2-holonomy/.

doi:10.5281/zenodo.22064342.
Relevance. The arena, and the sentence the whole argument rests on. The $\mathcal{M}_4 \times I_r \times S^1_\theta$ arena is carried from here at *derived*, on that paper's topological premise (T). The reading that gives this study its subject — loops exist only on the circle, so every physical phase is a holonomy, and the interval keeps no phase — is that paper's, not this one's.

BaBar Collaboration. *Observation of time-reversal violation in the B^0 meson system*, Phys. Rev. Lett. **109**, 211801 (2012).
Relevance. Direct observation of T violation independent of CP and of CPT assumptions — the laboratory anchor for the circle's contribution to T breaking in §16, as distinct from the interval's macroscopic irreversibility.

CMS Collaboration. *A search for microscopic black holes, string balls, and sphalerons in pp collisions at $\sqrt{s} = 13$ TeV*, arXiv:2604.10732 (2026).
Relevance. The first direct collider limit on sphaleron transitions: pre-exponential factor < 0.0034 at 95% CL for a 9 TeV threshold. It does not test the thermal sphaleron physics of §42, but it is the first experimental handle of any kind on the door.

CUORE Collaboration. *Constraints on lepton number violation with the 2 tonne·year CUORE dataset*, Science **390**, 1029 (2025), arXiv:2404.04453.
Relevance. $T_{1/2} > 3.5 \times 10^{25}$ yr in ^{130}Te. One of the three current-generation $0\nu\beta\beta$ limits in Table 20.

J. A. Casas and A. Ibarra. *Oscillating neutrinos and $\mu \to e\gamma$*, Nucl. Phys. B **618**, 171 (2001).
Relevance. The parameterization that makes the freedom of the seesaw explicit: an arbitrary complex orthogonal R between the measured light sector and the unmeasured heavy one. It is the reason leptogenesis has been able to accommodate any η_B, and the reason a *real-R*, dial-free door is a different kind of claim (§68).

DESI Collaboration. *DESI DR2 results II: measurements of baryon acoustic oscillations and cosmological constraints*, Phys. Rev. D **112**, 083515 (2025), arXiv:2503.14738.

W. Elbers et al. (DESI Collaboration). *Constraints on neutrino physics from DESI DR2 BAO and DR1 full shape*, Phys. Rev. D **112**, 083513 (2025), arXiv:2503.14744.
Relevance. The bound that puts this framework under its most immediate pressure: $\Sigma m_\nu < 64$ meV in ΛCDM, < 53 meV with a Feldman–Cousins treatment at the physical boundary, and a 3σ tension with the oscillation lower limit expressed through an effective parameter allowed to go negative. Discussed in §61 against the derived 59.2 meV.

DUNE Collaboration. *Long-baseline neutrino oscillation physics potential of the DUNE experiment*, Eur. Phys. J. C **80**, 978 (2020), arXiv:2006.16043.

DUNE Collaboration. *The DUNE science program: input to the European Strategy for Particle Physics — 2026 update*, arXiv:2503.23291 (2025).
Relevance. The current official sensitivity set: $\sigma(\delta_{CP}) = 7^\circ$ at $\delta_{CP} = 0$ and 18° at $-\pi/2$ for 1000 kt·MW·yr; 5σ CP-violation discovery for half of all true values at 600 kt·MW·yr; and $p \to K^+\bar{\nu}$ reach 1.3×10^{34} yr. Supersedes the 2020 EPJC numbers for quotation.

S. Davidson and A. Ibarra. *A lower bound on the right-handed neutrino mass from leptogenesis*, Phys. Lett. B **535**, 25 (2002).

Relevance. The bound that named the problem: for hierarchical heavy neutrinos the decay asymmetry is capped by the light and heavy masses, forcing $M_N \gtrsim 10^9$ GeV and pricing thermal leptogenesis permanently out of laboratory reach. It is the theorem Part VIII must evade, and does, in the resonant regime.

W. Du, Y. Li, Y. Wu, X. Zhang and J. Zhang. *Cosmological preference for a positive neutrino mass at* 2.7σ, arXiv:2507.16589 (2026).
Relevance. The counter-result quoted in §61: in a dynamical-dark-energy framework with free N_{eff} and weak lensing added, $\Sigma m_\nu = 98^{+16}_{-37}$ meV. Cited to show that the sign of the cosmological neutrino-mass preference is a function of the dark-energy parameterization and should not be quoted as established in either direction.

Euclid Collaboration. *Euclid preparation LIV: sensitivity to neutrino parameters*, Astron. Astrophys. **693**, A58 (2025), arXiv:2405.06047.

M. Fukugita and T. Yanagida. *Baryogenesis without grand unification*, Phys. Lett. B **174**, 45 (1986).
Relevance. Leptogenesis founded in four pages: create the asymmetry as lepton number through heavy Majorana decays and let the sphalerons convert. The mechanism class of Part VIII entire. This study's contribution is not the mechanism but the removal of its dials.

M. B. Gavela, P. Hernández, J. Orloff and O. Pène. *Standard Model CP violation and baryon asymmetry*, Mod. Phys. Lett. A **9**, 795 (1994).
Relevance. With Huet and Sather (1995), the calculation that established the $\sim 10^{-20}$ Standard-Model CP measure and thereby the shortfall of Part VI.

S.-F. Ge, C.-F. Kong, M. Lindner and J. P. Pinheiro. *Neutrinoless double beta decay in light of JUNO first data*, JHEP **03** (2026) 105, arXiv:2511.15391.
Relevance. The post-JUNO $m_{\beta\beta}$ bands: normal ordering from 1.3–1.8 meV at minimum to 3.5–4.0 meV at maximum, inverted ordering above 16–22 meV. The source for the correction in §60 that the tonne-scale program does *not* reach the normal ordering.

J. A. Harvey and M. S. Turner. *Cosmological baryon and lepton number in the presence of electroweak fermion-number violation*, Phys. Rev. D **42**, 3344 (1990).
Relevance. The 28/79 of Eq. (80): standard equilibrium bookkeeping, inherited unmodified, and flagged as textbook rather than derived.

Y. Hosotani. *Dynamical mass generation by compact extra dimensions*, Phys. Lett. B **126**, 309 (1983).
Relevance. Symmetry breaking driven purely by the holonomy of a circle, with no elementary scalar: the direct precedent for Eq. (10) being a physical, measurable label.

P. Huet and E. Sather. *Electroweak baryogenesis and Standard Model CP violation*, Phys. Rev. D **51**, 379 (1995).

Hyper-Kamiokande Collaboration. *Sensitivity of the Hyper-Kamiokande experiment to neutrino oscillation parameters using accelerator neutrinos*, Eur. Phys. J. C (2025), arXiv:2505.15019.
Relevance. Ten years at 27×10^{21} POT: δ_{CP} precision 20° at maximal CP violation improving to 6° at CP conservation, and 5σ exclusion of CP conservation for $> 60\%$ of true values. With the July 2025 cavern completion, operation is planned from 2028.

JUNO Collaboration. *Measurement of reactor neutrino oscillation with the first JUNO data*, Nature **654**, 343 (2026).

Relevance. $\sin^2\theta_{12} = 0.3092 \pm 0.0087$ and $\Delta m^2_{21} = (7.50 \pm 0.12) \times 10^{-5}$ eV^2 from 59 days — a factor 1.6 improvement on all previous measurements combined, and the near-term arbiter of the mass ordering that §59 records as contested.

C. Jarlskog. *Commutator of the quark mass matrices in the Standard Electroweak Model and a measure of maximal CP nonconservation*, Phys. Rev. Lett. **55**, 1039 (1985).
Relevance. The rephasing-invariant measure of CP violation — here read geometrically as *one area for the whole graph*, Eq. (16).

KATRIN Collaboration. *Direct neutrino-mass measurement based on 259 days of KATRIN data*, Science (2025), DOI 10.1126/science.adq9592.
Relevance. $m_\beta < 0.45$ eV at 90% CL, the world's best direct limit; final Phase-1 goal 0.3 eV. Listed in §61 so that its silence on a 59 meV sum is not mistaken for support.

K. Kajantie, M. Laine, K. Rummukainen and M. Shaposhnikov. *Is there a hot electroweak phase transition at $m_H \gtrsim m_W$?*, Phys. Rev. Lett. **77**, 2887 (1996).
Relevance. The crossover result: at the measured Higgs mass there is no first-order electroweak phase transition, hence no bubble walls and no electroweak baryogenesis inside the Standard Model. The first of the two failures of §28.

KamLAND-Zen Collaboration. *Search for Majorana neutrinos with the complete KamLAND-Zen dataset*, Phys. Rev. Lett. **135**, 262501 (2025), arXiv:2406.11438.
Relevance. The world-leading $0\nu\beta\beta$ limit: $T_{1/2} > 3.8 \times 10^{26}$ yr, $m_{\beta\beta} < 28$–122 meV.

V. A. Kuzmin, V. A. Rubakov and M. E. Shaposhnikov. *On anomalous electroweak baryon-number non-conservation in the early universe*, Phys. Lett. B **155**, 36 (1985).
Relevance. Unsuppressed sphaleron transitions at high temperature: the door of §42, and the reason an asymmetry written in $B - L$ becomes an asymmetry in B.

LEGEND Collaboration. *First results on the search for lepton-number-violating neutrinoless double beta decay with the LEGEND-200 experiment*, Phys. Rev. Lett. (2026), arXiv:2505.10440.

I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro and T. Schwetz. *NuFit-6.0: updated global analysis of three-flavor neutrino oscillations*, JHEP **12** (2024) 216, arXiv:2410.05380.
Relevance. The global fit against which $\delta_{\text{PMNS}} = 196.4°$ is scored: normal-ordering $212°\,^{+26}_{-41}$ with the atmospheric sample, $177° \pm 19°$ without. The lattice point sits inside one sigma of both — and so does 180°, which is why §58 declines to claim leptonic CP violation.

A. Pilaftsis. *CP violation and baryogenesis due to heavy Majorana neutrinos*, Phys. Rev. D **56**, 5431 (1997).
Relevance. Resonant CP asymmetries from quasi-degenerate heavy neutrinos: when the splitting approaches the width, the self-energy contribution is resonantly enhanced and the Yukawa suppression cancels. The enhancement of Eq. (49), credited at source.

A. Pilaftsis and T. E. J. Underwood. *Resonant leptogenesis*, Nucl. Phys. B **692**, 303 (2004).
Relevance. The systematic treatment of the resonant regime, including TeV-scale realizations. The literature this study is measured against in §53: the mechanism is theirs; what is claimed here is that the scale and the degeneracy arrive from elsewhere rather than being imposed.

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

Relevance. The measured $\eta_B = (6.12 \pm 0.04) \times 10^{-10}$ against which Eq. (85) is scored, and the source of the recombination and inventory parameters used in the cohort.

T. S. Roussy et al. *An improved bound on the electron's electric dipole moment*, Science **381**, 46 (2023).
Relevance. $|d_e| < 4.1 \times 10^{-30}\, e\,\mathrm{cm}$ at 90% CL — the tightest constraint on a CP phase off the lattice, and sensitive to scales an order of magnitude beyond the LHC's direct reach.

A. D. Sakharov. *Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe*, JETP Lett. **5**, 24 (1967).
Relevance. The three conditions, located in the geometry in §27: the circle's topology meeting the weak sector, the chiral alphabet, and the interval's forward semigroup.

Simons Observatory Collaboration. *Science goals and forecasts for the enhanced large-aperture telescope*, JCAP (2025), arXiv:2503.00636.
Relevance. $\sigma(\Sigma m_\nu) = 30$ meV, improving to 15 meV with $\sigma(\tau) = 0.002$; upgrade complete 2028, survey to 2034. With Euclid, the instrument that will decide §61 — and, since the July 2025 termination of CMB-S4, the leading one.

Super-Kamiokande Collaboration. *Search for proton decay via $p \to e^+\pi^0$ and $p \to \mu^+\pi^0$ with an enlarged fiducial volume in Super-Kamiokande I–IV*, Phys. Rev. D **102**, 112011 (2020), arXiv:2010.16098.
Relevance. $\tau/B > 2.4 \times 10^{34}$ yr in $p \to e^+\pi^0$. The bound this framework says will never be saturated, in any channel, at any exposure.

T2K and NOvA Collaborations. *Joint neutrino oscillation analysis from the T2K and NOvA experiments*, Nature **646**, 818 (2025), arXiv:2510.19888.
Relevance. The headline current constraint on δ_{CP}: normal-ordering best fit $-0.87\pi = 203°$, 68% interval $[166°, 241°]$, CP conservation not excluded in normal ordering — and a weak preference for inverted ordering that disagrees with the post-JUNO global fits. Both facts are quoted in §58 and §59.

K. G. Wilson. *Confinement of quarks*, Phys. Rev. D **10**, 2445 (1974).
Relevance. Identified the gauge-invariant content of a gauge theory as the set of holonomies around closed loops — which is why an open interval supplies none, §6.

T. T. Wu and C. N. Yang. *Concept of nonintegrable phase factors and global formulation of gauge fields*, Phys. Rev. D **12**, 3845 (1975).
Relevance. The physically meaningful content of a gauge field is the non-integrable phase factor, neither the potential nor the field strength. The most directly relevant external paper to this study's central move: reading a Yukawa texture as a connection and its CP content as a holonomy.

C. Abel et al. (nEDM Collaboration). *Measurement of the permanent electric dipole moment of the neutron*, Phys. Rev. Lett. **124**, 081803 (2020).

G. Adhikari et al. (nEXO Collaboration). *nEXO: neutrinoless double beta decay search beyond 10^{28}-year half-life sensitivity*, J. Phys. G **49**, 015104 (2022), arXiv:2106.16243.