UCLA/Arisaka-GG-P2-N3
v20.2 — August 22, 2026 — *Particle Branch*

Why Three Generations

A Three-Leg Derivation of $N_g = 3$ from BI-6, MSMF, and GDOS Morse Stability

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

Everything ordinary in the universe is built from the first family of matter particles: the electron, the up quark, the down quark, and their neutrino. Yet nature manufactured two heavier copies of that family — and then stopped. Why three? For half a century the Standard Model has carried the number 3 as an unexplained input: nothing in its equations prefers three families to two or to seventeen.

This paper derives it within Grand Geometric Theory (GG-Theory). Three independent directions of constraint close on one value. The observed violation of matter–antimatter symmetry requires at least three families. The relaxation capacity of the family stratum — proved here to be π of instanton budget per family, from the quantum geometry that the boundary's own adiabatic theorem forces on the family moduli — admits at most three: relaxation to the ground family runs through the weak interaction, and a fourth family's road would cost more than the stratum can pay. And the same integer, fed into the theory's mass formulas, reproduces all twelve charged-fermion masses to four significant figures. Every fence is declared: two empirical inputs, premises shared with the Born rule, and the coupling clause — which is no longer assumed but enforced, by a pushforward of the bulk's anomaly data onto the boundary's own geometry that also discharges the paper's one topological residue and closes both of its scope refinements at class level. The integrality statement is derived, its lattice forced by the bulk's spin structure.

After all, a number that nature has held fixed through every experiment ever performed deserves better than to be an input; it deserves to be a theorem.

Executive Summary

The user-level question behind this note is direct: why are there three generations of elementary fermions, and not one, two, or four? Within the Standard Model the question has no answer. $N_g = N_{\text{gen}}$ is an empirical fact without an internal mechanism. Within GG-Theory the question has a closed answer. This paper assembles that answer, in a single self-contained proof, out of three legs that are independent in their directions of constraint (§8.2).

The first leg is the Kobayashi–Maskawa observation, recast under the MSMF postulate of GG-Theory. The number of physical CP-violating phases in an N_g-family unitary mixing matrix is $(N_g - 1)(N_g - 2)/2$. CP violation is empirically real, so this number must be at least one, which forces $N_g \geq 3$. This gives the lower bound (L1).

The second leg is genuinely new and sets the present paper apart from the historical Kobayashi–Maskawa argument. It uses the open-system framework of GDOS together with Morse theory on the family-projector moduli space CP^{N_g-1}. The Geometric Free Functional admits a canonical Morse function whose critical points encode the cascade of geometric attractors, and descending the full family ladder costs the triangle number of gate tolls, $S_{\text{tot}}(N_g) = \pi N_g(N_g - 1)/2$. What pays those tolls is the central question of the paper. The boundary runtime's own adiabatic theorem — proved here at exactly the Born rule's grade of conditionality — forces the Berry connection and the Fubini–Study geometry on the family moduli (Theorems T-AD-1–T-AD-3, Appendix F); two obstruction theorems (Appendix G) prove that neither the bulk alone nor the boundary alone can set the inflow coupling; the integrality theorem L-INT is derived at physics rigor from Dirac quantization of the Berry line bundle and the bulk's own spin structure, leaving a single named topological residue (R-INT); and the capacity theorem T-N3-6 then delivers $\text{Cap}(N_g) = \pi N_g$ — exactly π of relaxation capacity per family — from which $N_g \leq 3$ follows by arithmetic, with $N_g = 2$ excluded in the same stroke (its stratum has no four-cycle to pay the single decay step) and exact saturation at $N_g = 3$. This is leg (L2). Its physical reading, adopted from the ground-state refinement of the cohort (Arisaka, 2026, A6-GRIP; Arisaka, 2026, A9-SM), is dynamical: all three generations are genuine geometric attractors — which is why m_τ and m_μ are sharply measurable — but only the first generation is the ground attractor (GGA); the $\tau \to \mu \to e$ weak-decay chain is a single GRIP-In descent whose gate tolls exhaust the stratum's instanton budget exactly at $N_g = 3$; and a fourth family's ground state would be dynamically unreachable within that budget — the deeper reason the bound is $N_g < 4$ (§11.6).

The third leg is the constructive cross-check. GG-A9-SM gives a minimal first-order codebook for fermion Yukawas in which every charged-lepton and quark mass is built out of the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. With $J = 7$ and $L = 5$ fixed by BI-6 closure, the entire codebook is a function of N_g, and the empirical mass ratios $m_\tau/m_\mu = 16.817$ and $m_\mu/m_e = 206.768$ uniquely select $N_g = 3$ at the four-significant-figure level. This is leg (L3).

The combination is significantly stronger than any of the three legs taken alone. Leg (L1) by itself gives only $N_g \geq 3$; leg (L2) by itself gives only $N_g \leq 3$; leg (L3) by itself does not exclude solutions in which the codebook would be re-parameterized. Together they force $N_g = 3$ exactly — conditional on the declared empirical inputs, on premises shared with the Constitutional Born Theorem, and on the integrality theorem L-INT, derived with one named topological residue (R-INT, discharged on the benchmark, with the warped and relative scope refinements closed at class level by Theorems T-INF-7 and T-INF-8) — and they validate that value with twelve quantitative mass predictions agreeing with observation at the four-significant-figure level.

The deeper significance of the result is structural. It shows that what was historically considered an inscrutable empirical accident – the existence of exactly three families – becomes a derived integer once the underlying geometry is correctly identified. The number 3 in $N_g = 3$ is not selected; it is forced by BI-6 closure, by Morse stability, and by the empirical non-vanishing of the Jarlskog invariant. This is the GG-Theory pattern: an empirical accident reduces, under sufficient geometric structure, to a number-theoretic identity.

A further structural ingredient is the complete dissolution of the assumptions that the natural axiomatization of leg (L2) would require. The budget cap quoted across the cohort's cross-reference ledgers is delivered by the capacity theorem; the two would-be technical hypotheses H_1 and H_2 are replaced, respectively, by a proved obstruction (the Rigidity Lemma: every functorial descent of the bulk curvature to the family moduli vanishes identically, so the coupling was never spacetime's to supply) and by the integrality theorem L-INT (the honest successor of non-degeneracy, now derived with the single topological residue R-INT); and strict Morse-positivity leaves the critical path entirely, surviving as an independent consistency check that the admissible value passes. The cleanest way to state the new structural content is this: quantum mechanics itself equips the family moduli with Berry curvature and Fubini–Study geometry (the adiabatic theorem, at Born-grade); the only natural capacities built on that geometry that yield a consistent and decisive theory give $\mathrm{Cap}(N_g) = \pi N_g$ or the equivalent constant form — both saturating at 3π — and πN_g funds exactly three tolls of the triangle-number cost $\pi N_g(N_g - 1)/2$. Three families spend the budget exactly; a fourth could not pay for its own road to the ground state.

A closing methodological discussion (Section 14) records the road not taken. Before settling on the three-leg proof of the present paper, we conducted an extensive but unsuccessful trial to derive $N_g = 3$ within closed-system GG-6/BI-6 alone. Four proposals were attempted – a Chern–Simons multiplicity axiom, a magnetized-torus index, a $\mathbb{Z}_3$ orbifold of the compact direction, and partial GG-6 anomaly saturation – and each produced only necessary conditions on N_g but no upper bound. The structural diagnosis is that no closed-system framework, no matter how elaborate, can produce an upper bound on integer multiplicities like N_g: closed-system methods give necessary admissibility conditions (anomaly cancellation, integrality, index theorems) but never sufficient conditions, because all unitary states are equally admissible by symmetry. Upper bounds require open-system dynamics: irreversible projection, finite instanton-action budgets, cascade closure on a moduli of attractors. This is exactly what GDOS supplies, and only GDOS supplies. We argue, in the new Section 14, that this is the missing link that has held $N_g = 3$ open for the century since the foundation of quantum mechanics in 1925. The individual ingredients of the open-system / holographic perspective have been understood for decades – Lindblad master equations (1976), decoherence theory (Zeh 1970, Zurek 1981), the holographic principle ('t Hooft 1993, Susskind 1995), AdS/CFT (Maldacena 1997), Hopkins–Singer differential cohomology (2005) – but they have never been assembled into a foundational framework. GDOS is this assembly, and the present derivation of $N_g = 3$ is the first concrete demonstration that the open-system shift has the foundational power to derive structural integers that closed-system methods cannot reach.

The remainder of the paper is organized as follows. Section 3 recalls the relevant pieces of the GG-Theory framework. Sections 5–7 prove the three legs. Section 8 combines them into the main theorem. Section 9 interprets the angular cascade $e^{-n\pi}$ as instanton suppression along the GFF Morse flow. Sections 10–11 discuss falsifiability and prior approaches. Sections 12–13 record open problems and conclusions. Seven appendices contain the detailed calculations; Appendix E carries the parity consistency route and the record of the withdrawn hypotheses, Appendix F proves the adiabatic theorems, and Appendix G proves the two obstruction theorems, the capacity classification, and the integrality theorems.

After all, the difference between two families and three is the difference between a universe without CP violation and the one we inhabit; a theory that advertises no dials must get that integer right — and must say plainly under which declared assumptions it does so.

Keywords: three generations; family number; $N_g = 3$; MSMF; GDOS; Morse theory; CP^{N_g-1}; family-projector moduli; BI-6; OGN-5; fermion mass codebook; Kobayashi–Maskawa; Jarlskog invariant; GFF cascade; GRIP; instanton suppression; GRIP-In; GGA; reachability; capacity theorem; integrality theorem; adiabatic theorem; Berry connection; quantum geometric tensor; GG-Theory

Three Families Spend the Budget Exactly; A Fourth Could Not Pay for Its Own Road

Why not two, and why not four. Descending the family ladder costs the triangle number of gate tolls; the family stratum has exactly π of relaxation capacity per family to pay them with, and none at all below three families, where the four-dimensional homology generator does not exist. Those two lines are the whole upper bound. Evaluated at every candidate count they cross exactly once: $N_g = 2$ fails for want of a four-cycle to pay its single decay step, $N_g = 3$ pays its way to the penny, and from four the cost has outrun the capacity and the gap widens at every step thereafter. The capacity theorem stands at the Born rule's own grade of conditionality, and the paper declares that before a reader asks. In full: Figure 4, p. 36.

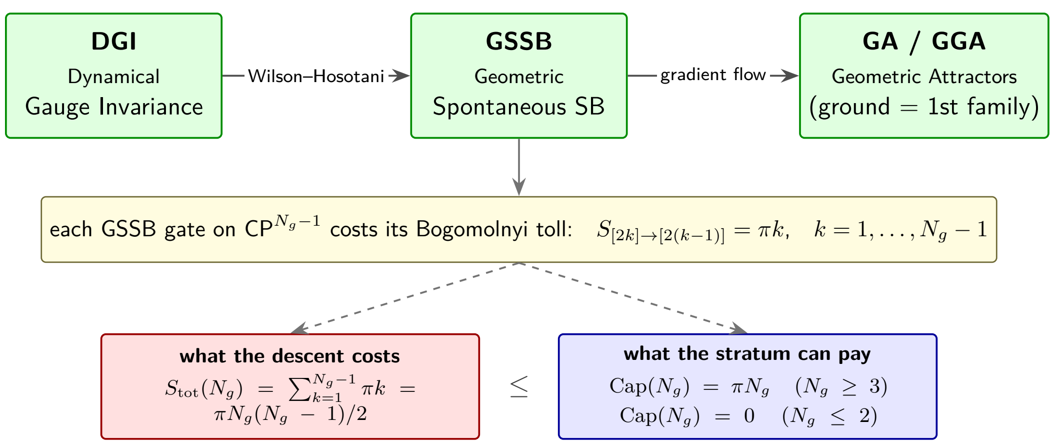

closure $S_{\mathrm{tot}}(N_g) \leq \mathrm{Cap}(N_g)$, evaluated at every candidate

N_g	1	2	3	4	5	6
Cap/π	0	0	3	4	5	6
S_{tot}/π	0	1	3	6	10	15
verdict	pass	fail	saturates	fail	fail	fail

$N_g = 2$ fails because $H_4(\mathrm{CP}^1) = 0$ leaves the single decay step with no payer; every $N_g \geq 4$ fails the budget. $N_g = 1$ passes here with nothing on either side and is excluded by leg (L1). Exact saturation: $S_{\mathrm{tot}}(3) = \mathrm{Cap}(3) = 3\pi$.

Overview of Contents

List of Figures

List of Tables

Contents

Preface

This paper answers one question — why three fermion generations — and it is written to be audited rather than admired. Three conventions therefore govern every page.

First, every fence is printed. Where a step is a theorem, it says so; where it rests on a measurement, the measurement is named; where a residue remains, it is named too; the chain carries exactly one (R-INT, the topological residue of the integrality theorem L-INT of §7.6), and Appendix H discharges it on the benchmark and closes its two scope refinements at class level. Two obstruction theorems are included precisely because they say what *cannot* be derived, and from where — a negative result is the honest form of progress on a question that has resisted proof for fifty years.

Second, the physics is stated beside the mathematics. Each technical movement closes with a paragraph in plain terms: what the landscape is, what the tolls are, who pays them, and why the books balance only through three families. A reader who wants the argument without the machinery can read those paragraphs alone and lose nothing essential.

Third, the arithmetic is reproducible. Every number in the upper bound is an integer multiple of π obtained by pairing a characteristic class against a four-cycle, and the worked computation is carried out in full at $N_g = 2$, 3 and 4 (§7.9). Nothing is asserted that a reader cannot recompute in an afternoon.

The result, stated once and plainly: the family stratum earns π of relaxation budget per family and the descent charges the triangle number of tolls, so the ledger balances only up to three — and at three it balances exactly.

Part I
The Question, the Claim, and the Framework

The three-generation puzzle and its half-century of resistance; the central claim as one deductive funnel; the inherited machinery — BI-6, OGN-5, GDOS, GRIP, and the GFF Morse functional; and the twenty-two-paper cohort the paper stands in, with the roadmap of the eight Parts

1 Introduction and statement of the problem

1.1 The three-generation puzzle

The Standard Model of particle physics organizes the known elementary fermions into three families with identical gauge quantum numbers and progressively heavier masses (Particle Data Group, 2024). This three-fold replication is among the most robust empirical facts in physics: the LEP measurement of the invisible Z width (ALEPH et al., 2006) fixes the number of light, weakly interacting neutrino species at

$$N_\nu = 2.984 \pm 0.008, \tag{1}$$

consistent with three and inconsistent with both two and four at the multi-sigma level. Cosmological probes from primordial nucleosynthesis and the Cosmic (Arisaka, 2026, A10-Cosmos) Microwave Background give an independent and equally tight bound (Planck Collaboration, 2020), $N_{\text{eff}} = 2.99 \pm 0.17$. Yet within the Standard Model itself the family number is left as a free input: nothing in $\mathrm{SU}(3) \times \mathrm{SU}(2) \times \mathrm{U}(1)$ requires three replications rather than one or seventeen. Generations of theoretical attempts have sought to derive $N_g = 3$ from deeper principles, and have produced partial successes but no closed proof.

The leading prior approaches and their respective limitations may be summarized as follows.

(a) **Heterotic string compactifications** (Candelas et al., 1985; Green, Schwarz & Witten, 1987). The Calabi–Yau manifold of compactification has Hodge numbers $(h^{1,1}, h^{2,1})$, with $|\chi/2| = |h^{1,1} - h^{2,1}|$ counting net chirality. Three generations corresponds to $|\chi/2| = 3$. This is realized by certain CICY threefolds and quotients but the choice of manifold is not forced; the family number is selected, not derived.

(b) **Magnetized toroidal compactifications** (Blumenhagen et al., 2007; Ibáñez & Uranga, 2012). The index theorem gives the chiral asymmetry as the magnetic flux quantum on each torus factor, $\sum_a m_a = N_g$. This explains how three could be generated, but again leaves the magnetic flux quanta as free integers.

(c) $\mathbb{Z}_3$ **orbifolds and discrete symmetries** (Ibáñez et al., 1987; Kikuchi et al., 2021). A $\mathbb{Z}_3$

point group trivially produces three twisted sectors. The choice of $\mathbb{Z}_3$ rather than $\mathbb{Z}_2$ or $\mathbb{Z}_4$ is a model-building input.

(d) **Anomaly-cancellation arguments** (Dobrescu & Poppitz, 2001; Witten, 1985). These constrain combinations such as the mixed gauge-gravitational anomaly but do not by themselves fix the number of replicas; in the special setting of two universal extra dimensions, global anomaly cancellation does force $N_g = 3$ outright, at the price of the extra-dimensional assumption (Dobrescu & Poppitz, 2001).

(e) **Neutrino-physics-driven CP arguments** (Kobayashi & Maskawa, 1973). Kobayashi and Maskawa famously observed that $N_g = 3$ is the smallest number of generations admitting a physical CP-violating phase. This is widely accepted as a consistency argument but has the same shortcoming as (a)–(d): it shows $N_g \geq 3$, not $N_g = 3$.

In every case the residual gap is the same: prior approaches close one direction of the inequality but not the other. Either they explain why N_g cannot be too small, or they explain how a particular geometric model can produce three, but they do not provide an upper bound that excludes $N_g \geq 4$.

1.2 The GG-Theory hypothesis

In GG-Theory the family number is not introduced as an input; it is derived. The derivation runs alongside, and not through, the integer ledger that fixes the bulk, so it is worth setting the two out together before either is used. The ledger is

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

whose uniqueness as the minimal Diophantine solution is proved in Arisaka (2026, A2-GG6) and whose BI-6 closure and integrality are proved in Arisaka (2026, A3-BI6). The five integers are forced, not fitted, and each carries a definite meaning: $M = 1$ the single tensor seam, $N = N_c = 3$ the color index, $L = 5$ the SU(5) block, $J = 7$ the odd closure index, and $K = 35 = LJ$ the throat depth. The color reading of N is settled without reference to any generation count: it follows from the dual lock together with the block embedding $N < L$, Theorem T-BI6-10.2a of Arisaka (2026, A3-BI6). The bulk's two non-spatial dimensions — a radial scale and a compact phase — fix the only two transcendentals the theory admits: e, the base of every scale-exponential, and π, the angle of the gauge circle. The fermion family number is a *different quantity*, written N_g throughout this paper. It is not a component of (2) and cannot be read off it: the ledger fixes the color index, and no step of the ledger's Diophantine cascade bears on how many generations the boundary carries. What the present paper proves is that $N_g = 3$, from premises of its own — the CP-violation combinatorics of the boundary mixing matrix and the relaxation capacity of the family stratum. That N_c and N_g both equal 3 is then a result of two independent derivations meeting at one integer, and it is discussed as such in §16.2 rather than assumed anywhere in what follows.

The purpose of the present short paper is to assemble, in a single self-contained proof, the three independent legs that force $N_g = 3$:

(L1) the MSMF lower bound from CP-violation combinatorics;

(L2) the GDOS Morse upper bound from boundary-projector stability on CP^{N_g-1};

(L3) the constructive codebook check from GG-A9-SM.

The combination of (L1), (L2), and (L3) closes the derivation: $N_g \geq 3$ and $N_g \leq 3$ and the resulting twelve fermion masses agree with observation to four significant figures. No leg is by itself sufficient; together they are. All conditioning — the empirical inputs, the Born-grade runtime premises, the coupling clause, and the topological residue R-INT of the integrality theorem L-INT — is declared explicitly in Theorem 9.1 and Appendices C, F and G. Two framings deserve emphasis at the outset. First, the genuinely new and fully rigorous content of the upper bound is a topological identity: the alternating Morse sum on the family-projector moduli equals the Hirzebruch signature, $a^+(N_g) - a^-(N_g) = \sigma(\mathsf{CP}^{N_g-1})$, which is 1 for odd N_g and 0 for even N_g, so *even N_g is excluded by the signature alone* (Appendix E, Steps 1–2), independently of the capacity theorem — which excludes $N_g = 2$ by a different route (its stratum has no four-cycle) and every $N_g \geq 4$ by the budget. Second, the integer derived here is the *generation* number N_g, and it is not the color index. That the OGN-5 ledger's second component is also 3 is a separate and *unconditional* result — the dual lock with the block embedding, Theorem T-BI6-10.2a of Arisaka (2026, A3-BI6), checked against the BI-6 trace ratio $(L - N)/L = 2/5$ — reached from premises that mention no generation count at all. Two derivations with nothing in common therefore land on the same integer. GG-Theory reads that as a structural fact about the bulk rather than an accident, and §16.2 sets out what would have to be shown for the reading to become a theorem; nothing in the present derivation rests on it.

1.3 Plan of the paper

Section 2 states the central claim as one deductive funnel, shared in structure with the companion papers of the cohort. Section 3 recalls the relevant pieces of the GG-Theory framework: BI-6, OGN-5, GDOS, GRIP, and the GFF Morse functional, in just enough detail to make the present paper self-contained. Section 4 places the paper in the twenty-two-paper cohort and lays out the roadmap of the eight Parts. Section 5 proves the MSMF lower bound $N_g \geq 3$. Section 6 proves the GDOS upper bound $N_g \leq 3$ through the capacity architecture: the toll ledger of the cascade, the two obstruction theorems that locate the axiom frontier, the adiabatic theorems that force the Berry/Fubini–Study geometry at Born-grade, the reduced coupling clause with its integrality theorem (derived; one named topological residue), and the capacity theorem with its unit classification and form-robustness result. Section 7 reviews the GG-A9-SM codebook and shows that it is consistent with $N_g = 3$ and inconsistent with any other value. Section 8 combines the three legs into the closure theorem. Section 9 interprets the angular cascade $1, e^{-\pi}, \ldots, e^{-4\pi}$ as instanton suppression along the GFF Morse flow. Section 10 discusses falsifiability. Section 11 compares with prior approaches. Section 12 catalogs open problems. Section 13 concludes. Section 14 closes with the methodological reflection on why $N_g = 3$ required open-system thinking. Seven appendices contain the detailed calculations.

2 The central claim: the family number $N_g = 3$ as an open-system observable

It is worth stating, in one place, the recognition this paper shares with the whole GG-Theory cohort. No physical system is closed — so the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection *is* observation itself; its dynamics is the open-system law GDOS (GG-A5-GDOS), forced by the geometric Gi-4 (GG-A6-GRIP) and quantum-information Qi-4 (GG-A7-Qi4) pillars. It follows that the family number $N_g = 3$ — that there are exactly three fermion generations — is forced as a *registered observable* of one projection, not chosen — which is why this paper can derive it, and why a single clean measurement can refute it. The one clause that the derivation could not at first do better than assume, the coupling between the bulk's inflow and the stratum's quantum geometry, is enforced in Appendix H by a computation that consumes no quantum-mechanical input and arrives at the boundary's own answer. The funnel below is shared, in structure, with the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs.

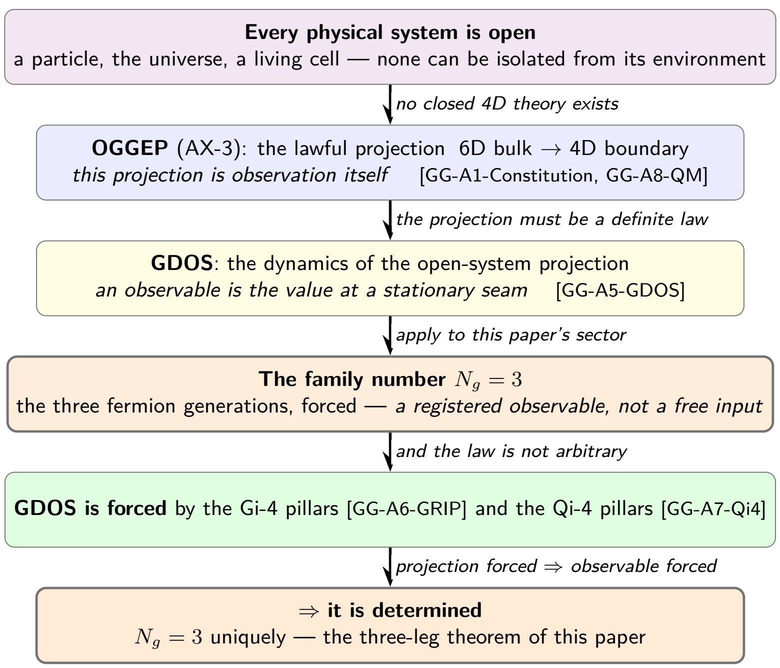

Figure 1: The central claim, as one deductive funnel, shared across the GG-Theory cohort. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection is observation itself; its dynamics is GDOS, forced by the Gi-4 and Qi-4 pillars. The same funnel opens the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs — here, the family number $N_g = 3$.

3 Framework: BI-6, OGN-5, GDOS, and the GFF Morse functional

We present the minimum machinery needed for the proof, with citations to the foundation papers.

3.1 The 6D bulk and Spin(1,5)

GG-Theory takes the fundamental arena to be a six-dimensional Lorentzian manifold $\mathcal{M}_6$ with isometry group containing $\mathrm{Spin}(1,5)$, the quaternionic Lorentz double cover (Arisaka, 2026, A1-Constitution; Arisaka, 2026, A2-GG6). The accidental isomorphism

$$\mathrm{Spin}(1,5) \cong \mathrm{SL}(2,\mathbb{H}) \tag{3}$$

explains why $d = 6$ is the unique dimension in the magic-dimension ladder $\mathbb{R} \to \mathbb{C} \to \mathbb{H} \to \mathbb{O}$ (corresponding to $d = 3, 4, 6, 10$) at which the holonomy algebra splits cleanly into a quaternionic doublet of complex doublets without either over-constraining (smaller d) or under-constraining (larger d) the gauge content. The boundary $\partial\mathcal{M}_6 = \mathcal{M}_4$ is our observed four-dimensional spacetime.

3.2 BI-6: Green–Schwarz-modified Bianchi identity

The fundamental closure condition on $\mathcal{M}_6$ is the Green–Schwarz-modified Bianchi identity (Arisaka, 2026, A3-BI6; Green & Schwarz, 1984)

$$dH_3 = X_4 = \mathrm{tr}\, R \wedge R - \mathrm{Tr}\, F_5 \wedge F_5 - \tfrac{2}{5}\, F_B \wedge F_B - \tfrac{2}{5}\, F_{Qi} \wedge F_{Qi}, \tag{4}$$

where H_3 is the Kalb–Ramond three-form, R the curvature two-form of the $\mathrm{Spin}(1,5)$ frame bundle, F_5 the field strength of $\mathrm{SU}(5)$, F_B that of $\mathrm{U}(1)_B$ (gravity-baryon), and F_{Qi} that of $\mathrm{U}(1)_{Qi}$ (quantum information). The numerical coefficients $-2/5$ are not fitted: they are forced by the inner-product normalizations on H_2 and by the requirement that X_4 be an integral cohomology class, i.e. the de Rham representative of an element of $H^4(\mathcal{M}_6;\mathbb{Z})$.

The integrality conditions on (4) admit a unique minimal solution (Diophantine uniqueness: Arisaka, 2026, A2-GG6; closure and integrality: Arisaka, 2026, A3-BI6), namely the OGN-5 ledger

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

The five integers satisfy two algebraic identities,

$$K = M^2 + N^2 + L^2 = 1 + 9 + 25 = 35, \qquad K = LJ = 5 \cdot 7 = 35. \tag{6}$$

Their physical readings are: $M = 1$ the primitive tensor slot, $N_g = 3$ the color-triplicity index (its three-generation face a downstream consequence), $L = 5$ the defining block size of $\mathrm{SU}(5)$, $J = 7$ the minimal symmetric closure cardinality of the charge packet $\{0, \pm1, \pm2, \pm3\}$, and $K = 35$ the total lock index. The identities (6) are equivalent to BI-6 closure on the index 35 (Arisaka, 2026, A3-BI6).

3.3 GDOS: Geometric Dynamics of Open Systems

Boundary observation is open-system dynamics in the sense of GDOS (Arisaka, 2026, A5-GDOS). Given a bulk state $x \in \mathcal{M}_6$ and a boundary projector $\Pi : \mathcal{M}_6 \to \mathcal{M}_4$ acting via bulk-to-boundary restriction, the boundary observable is

$$O_4(\Pi[x]) = \int_{\mathcal{M}_6} d^6X\, K_{\text{hol}}(\Pi; X, x)\, O_6(x), \tag{7}$$

where K_{hol} is the bulk-to-boundary transport kernel. The dynamics of boundary attractors – Geometric Attractors (GA) in the GRIP language – is governed by the Geometric Free Functional (GFF; developed as a Gi-4 pillar in Arisaka, 2026, A6-GRIP)

$$F_{\text{GFF}}[\Pi] = F_{\text{GEP}}[\Pi] + F_{\text{CEP}}[\Pi], \tag{8}$$

with F_{GEP} the geometric energetic part (curvature, gauge action, matter overlaps) and F_{CEP} the curvature-entropic part (Riemannian volume, configurational entropy of the projector class).

3.4 GRIP: the cascade of unstable critical points

The GRIP cascade (Arisaka, 2026, A6-GRIP)

$$\mathsf{DGI} \to \mathsf{GSSB} \to \mathsf{GA} \tag{9}$$

encodes how a deeply geometrically invariant bulk descends, by spontaneous symmetry breaking on the boundary, into a finite catalog of stable boundary attractors. Critical points of F_{GFF} are stationary projector configurations. Stable critical points are GA's; saddle points correspond to unstable boundary phases that decay into GA's by instantons.

Two regime terms, defined at runtime level in Arisaka (2026, A6-GRIP) and used throughout Arisaka (2026, A9-SM), are needed in Section 9. GRIP-In is the descent on a *fixed* landscape: gradient flow within basins, instanton steps across the GSSB gates between them, halting uniquely at the stratum's Ground Geometric Attractor (GGA). GRIP-Out is the complementary regime: a transient out-of-equilibrium event that reorganizes the landscape itself and supplies the initial datum from which a new GRIP-In descent begins. For the family sector the division is sharp: the electroweak (Wilson–Hosotani) branch selection was the boundary's one GRIP-Out event — it built the Yukawa landscape — and every flavor transition since, including the $\tau \to \mu \to e$ chain, is GRIP-In on that fixed landscape.

3.5 The MSMF postulate

Postulate 3.1 (MSMF ledger-minimality corollary; MSMF = Maximum Symmetry, Minimum Freedom, Arisaka, 2026, A1-Constitution)**.** Among all ledger choices that close BI-6 (4) and that admit a non-empty boundary observable spectrum compatible with empirical CP violation, GG-Theory selects the ledger of minimum total complexity, defined as the sum of integer absolute values $|M| + |N| + |L| + |J| + |K|$ subject to the closure constraints (6).

MSMF is the constitutional axiom AX-4 of Arisaka (2026, A1-Constitution) ("Maximum Symmetry, Minimum Freedom": no hidden continuous knobs); Postulate 3.1 is its ledger-minimality corollary as applied to BI-6 closure, the same usage as the MSMF minimality criterion of Arisaka (2026, A6-GRIP). The corollary is not a tautology: there exist solutions to (6) with larger entries (e.g. $K = 35 \cdot k^2$ for any $k \in \mathbb{Z}_{>0}$) and smaller entries that fail integrality. MSMF picks the unique smallest valid solution. We will use a direct corollary of MSMF below, applied to the boundary spectrum rather than to the ledger: the family number N_g realized on the boundary is the smallest integer compatible with the empirically observed flavor structure there.

4 The cohort inheritance and the roadmap of the eight Parts

Before the three legs of the derivation are walked (Part II onward), it is useful to lay out two reader-orientation maps on the same opening: a panorama of the twenty-two-paper cohort in which this paper stands (§4.1, Figure 2), and a panorama of the eight Parts of the present paper (§4.2, Figure 3). A reader of any later Part should be able to flip back to this Section and orient in both directions: which cohort companion carries which imported result, and where the present argument stands in the eight-Part arc.

4.1 The cohort inheritance

Figure 2 places the present paper among the twenty-two papers of the GG-Theory cohort. The reader meeting the program for the first time can read its architecture off the figure top to bottom — the constitutional root, the bulk trio, the open-system boundary law and its pillars, the three textbook readouts, the capstone, the two phenomenology branches, and the Foundations Branch on the spine's left flank — before continuing with the present paper's own thread.

Within that frame the present paper is the middle member of the Predictions Branch. It consumes the OGN-5 integer ledger of GG-A2-GG6, the BI-6 closure of GG-A3-BI6, the open-system boundary law of GG-A5-GDOS with the GRIP/GFF runtime of GG-A6-GRIP, and the minimal mass codebook of GG-A9-SM — exactly the imports recalled in Section 3 — and it supplies the family number $N_g = 3$ that its branch sibling GG-P3-Electron (Arisaka, 2026, P3-Electron) imports verbatim: the internal arrow drawn inside the Predictions frame of the map.

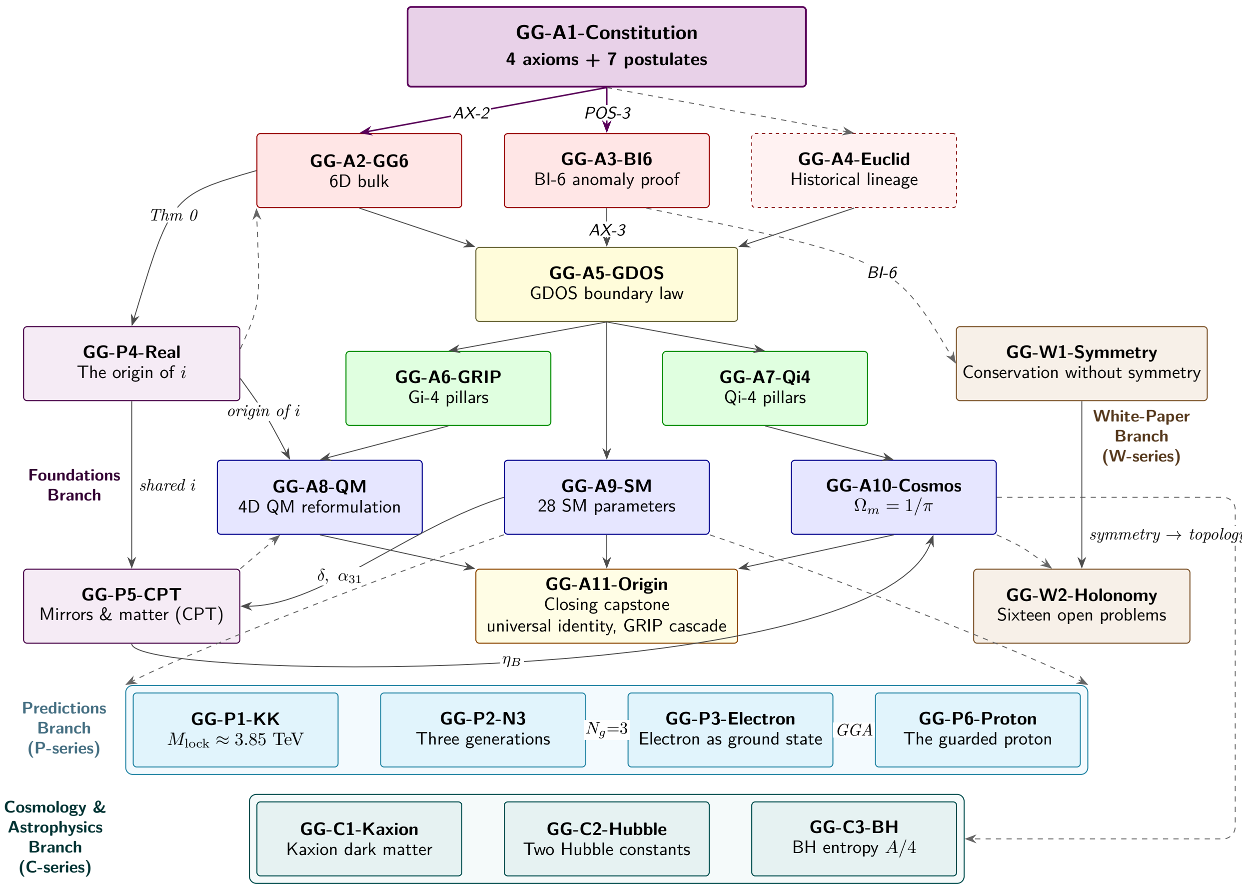

Figure 2: The twenty-two-paper cohort map of GG-Theory. The constitutional root GG-A1-Constitution (violet, top) supplies the four axioms and seven postulates from which the rest of the cohort descends: the GG6 Trio (GG-A2-GG6, GG-A3-BI6, GG-A4-Euclid, red) realizes the bulk geometry; GG-A5-GDOS (yellow) the boundary law; the GDOS Trio pillars (GG-A6-GRIP, GG-A7-Qi4, green) the geometric and quantum-information runtime; the 4D-Physics Trio (GG-A8-QM, GG-A9-SM, GG-A10-Cosmos, blue) the textbook readouts; and the closing capstone GG-A11-Origin (gold) the quarks-to-consciousness identity cascade. Two framed branches carry the phenomenology — the **Predictions Branch** (cyan), whose four papers hand their results along the row by the internal arrows (GG-P2-N3's N_g=3 into GG-P3-Electron, and that paper's ground-attractor method, GGA, on into GG-P6-Proton), and the **Cosmology & Astrophysics Branch** (teal). The spine is flanked by two columns of two, which do opposite jobs. On the left the **Foundations Branch** (violet) feeds it: GG-P4-Real, fed by GG-A2-GG6's Theorem 0 arena, manufactures the i that GG-A8-QM consumes, and GG-P5-CPT is the GG-A9-SM–GG-A10-Cosmos interface, fed by the ledger phases and feeding the η_B chain across the width of the map, with the shared-i edge descending solid between them. On the right the **White-Paper Branch** (brown) reads it: GG-W1-Symmetry walks the century from Noether to Hopkins–Singer, and GG-W2-Holonomy dissects the structure that arc arrives at — protection by topology rather than by symmetry — across sixteen open problems. The white papers derive nothing new, carrying every result at its parent's declared grade; the single dashed feed drawn into each names its heaviest source, not its only one. *Reading conventions:* solid arrows mark axiom-forced descent and carry an edge label naming the driving axiom or postulate where most illustrative; dashed arrows mark the deployment of derived results, or historical links inherited without being constitutionally forced. The branches this paper neither feeds nor draws on are on the map all the same, and are cited here at roster level, which is all they can honestly claim from it: Arisaka, 2026, A4-Euclid; Arisaka, 2026, A7-Qi4; Arisaka, 2026, A11-Origin; Arisaka, 2026, C2-Hubble; Arisaka, 2026, C3-BH; Arisaka, 2026, P6-Proton; Arisaka, 2026, W1-Symmetry; Arisaka, 2026, W2-Holonomy.

The Foundations Branch. The map's left flank carries the cohort's third branch — two member papers, a single column beside the spine, whose subject is the boundary's complex structure and its mirrors. GG-P4-Real (Arisaka, 2026, P4-Real) is the origin-of-i member: it develops at full length the real-bulk account of the boundary's imaginary unit, derived from two real structures of the GG-A2-GG6 arena — the compact phase circle and the causal clock — the same complex structure that the unitary mixing matrices of leg (L1) presuppose at boundary level. GG-P5-CPT (Arisaka, 2026, P5-CPT) is the CPT member, and it is the Foundations member nearest to the present paper: it consumes the three-generation structure derived here. Three flavors are the arena its chiral-pentagon reading of CP violation runs on, and its flavor-holonomy calculus ties the CP-odd quartets of the mixing matrix to the single surviving Jarlskog phase precisely at $N_g = 3$ — the same empirical non-vanishing that leg (L1) turns into the lower bound. Both carry P-series designators; their place at the spine's flank, not the phenomenology rows, marks their foundational role.

4.2 Roadmap of the eight Parts

The paper proceeds in eight Parts. The present Part — the question, the claim, and the framework — is nearly done: the three-generation puzzle and its half-century of resistance, the central claim stated as one deductive funnel, the inherited machinery (BI-6, OGN-5, GDOS, GRIP, and the GFF Morse functional) recalled in just enough detail to make the paper self-contained, and — on this opening — the cohort the paper stands in. Everything after this page is the derivation and its assessment, and the itinerary below says where each piece lives.

Part II is the heart, and it is highlighted in the roadmap because it is what the title claims. It walks the three legs in their natural order and closes the pincer. Leg (L1) turns the observed violation of CP symmetry, under the MSMF postulate, into the lower bound $N_g \geq 3$. Leg (L2) is the genuinely open-system leg: the GFF Morse cascade on the family-projector moduli space CP^{N_g-1}, priced against the capacity $\mathrm{Cap}(N_g) = \pi N_g$ that the stratum's own quantum geometry funds — the Berry/Fubini–Study structure forced at Born-grade by the runtime's adiabatic theorem — yields the upper bound $N_g \leq 3$, with two proved obstruction theorems locating the axiom frontier, on which the integrality theorem L-INT is then derived — its residue R-INT discharged on the benchmark background in Appendix H, and the coupling clause with it. Leg (L3) closes the circle constructively: with the ledger integers fixed, the GG-A9-SM codebook reproduces the twelve charged-fermion masses at four significant figures for $N_g = 3$ and for no other value. The closure section combines the three into the paper's theorem — $N_g = 3$ exactly, under the declared conditioning.

Part III steps back from the proof to its physical reading and its exposure. The GA-cascade section reads the angular suppression $1, e^{-\pi}, \ldots$ as instanton descent along the Morse flow — why the heavy generations are sharply measurable attractors and why only the first is the ground state — and the predictions section converts the derivation into declared falsifiers: the kill criteria a single clean measurement could trip.

Part IV is the paper's self-critical stratum: the comparison with the prior approaches to the family number — from Calabi–Yau Euler characteristics to anomaly cancellation to the anthropic route — and the open problems that remain, stated as open. Part V concludes twice: once technically, gathering what has been established and under which declared fences, and once methodologically, in the reflection on why $N_g = 3$ required open-system thinking at all — the lesson the paper regards as its deepest export.

Part VI collects the eight technical appendices — the CKM phase count, the Morse background on CP^{N_g-1}, the two capacity formulations and their completion, the numerical mass comparison, the parity route with the record of the withdrawn hypotheses, the adiabatic theorems of the boundary runtime, the obstruction theorems with the capacity classification, and the inflow bridge that enforces the coupling clause (Appendix H) — while Part VII answers twenty-one objections and Part VIII carries the references, each with a relevance note. Figure 3 renders the whole development as a single flowchart, so the big picture is visible at a glance.

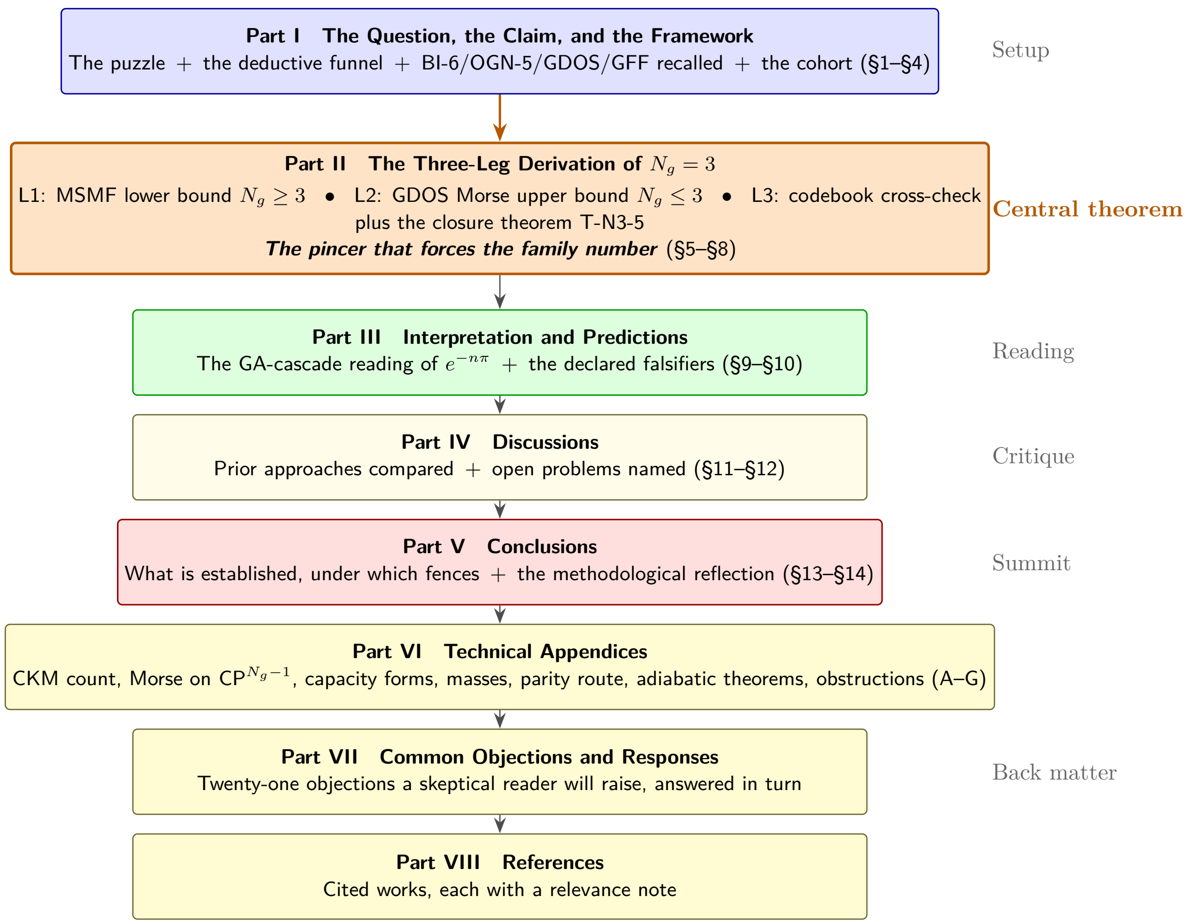

Figure 3: Roadmap of the eight Parts of GG-P2-N3. **Part II** is highlighted because it carries the paper's central claim — the three-leg derivation of the family number: leg (L1) forces $N_g \geq 3$ from observed CP violation under MSMF; leg (L2) forces $N_g \leq 3$ from the GFF Morse cascade on the family-projector moduli priced against the BI-6 instanton budget; leg (L3) validates the closure constructively through the GG-A9-SM mass codebook; and the closure section combines them into T-N3-5: $N_g = 3$ exactly, under the declared conditioning. Every later Part interprets, exposes, or documents that result.

Part I has now done its work: the question is posed, the claim is stated, the machinery is recalled, and the paper's place in the twenty-two-paper program is drawn. What remains is the derivation itself. The reader who wants the theorem should walk straight into Part II; the reader who wants the meaning first may skip ahead to Part III and return. Either way, the number at the end is the same: three.

5 Part I in plain English: a number that has resisted for fifty years

Ordinary matter needs one family of quarks and leptons. Nature built three, identical in every charge and differing only in mass, and then stopped. The count is not in doubt: the invisible width of the Z fixes the number of light neutrino species at 2.984 ± 0.008, which excludes two and four by many standard deviations, and primordial nucleosynthesis and the microwave background agree from an entirely different direction.

What is missing is a reason. The Standard Model accommodates any number of families without strain — write down a fourth and nothing breaks — so the observed three is an input, and fifty years of attempts to derive it have produced constructions that permit three rather than require it.

This Part sets up the problem and says what this paper will claim: that three is forced, under stated conditions, by three arguments that constrain from different directions. It also fixes the framework the argument runs in and marks what is imported from companion papers rather than proved here. Nothing is derived in this Part. Its job is to make the claim precise enough to be wrong.

Part II
The Three-Leg Derivation of $N_g = 3$

The pincer that closes: the MSMF lower bound $N_g \geq 3$ from CP violation, the GDOS Morse upper bound $N_g \leq 3$ from the GFF cascade against the BI-6 instanton budget, the constructive codebook cross-check from GG-A9-SM, and the closure theorem $N_g = 3$

6 Leg 1 (L1): MSMF lower bound $N_g \geq 3$ from CP violation

6.1 Statement of the result

Theorem 6.1 (MSMF lower bound (T-N3-1))**.** *Let N_g denote the number of fermion families carried by the boundary. Suppose that the boundary observable spectrum admits a non-zero Jarlskog CP-violating invariant $J_{\mathrm{CKM}} \neq 0$ in the quark sector. Then under MSMF*

$$N_g \geq 3. \tag{10}$$

Proof. The proof has three steps.

Step 1. Boundary chiral fermions of GG-Theory descend from a 6D quaternionic Dirac spinor by chiral projection (Arisaka, 2026, A1-Constitution; Arisaka, 2026, A9-SM). After Higgs descent the up-type and down-type quark mass matrices M_u, M_d are general complex $N_g \times N_g$ matrices. The CKM matrix is

$$V_{\mathrm{CKM}} = U_u^L U_d^{L\,\dagger}, \qquad U_{u,d}^L \in U(N_g). \tag{11}$$

In the standard parameterization V_{CKM} has N_g^2 phases, $N_g(N_g-1)/2$ mixing angles, and $N_g(N_g+1)/2$ moduli; subtracting the $2N_g - 1$ unphysical rephasings of u- and d-flavor eigenstates leaves

$$n_{\mathrm{phase}}(N_g) = \frac{(N_g-1)(N_g-2)}{2} \tag{12}$$

physical CP-violating phases.

Step 2. By definition, a Jarlskog invariant exists if and only if there exists a basis-independent imaginary part of a quartic product of CKM matrix elements,

$$J_{\mathrm{CKM}} = \mathrm{Im}\big(V_{ij} V_{kl} V_{il}^* V_{kj}^*\big), \qquad i \neq k,\ j \neq l, \tag{13}$$

which is non-zero. This requires at least one physical phase, i.e. $n_{\mathrm{phase}}(N_g) \geq 1$.

Step 3. Solving (12) for $n_{\mathrm{phase}} \geq 1$ gives

$$\frac{(N_g-1)(N_g-2)}{2} \geq 1 \iff N_g \geq 3. \tag{14}$$

The boundary case $N_g = 2$ has $n_{\text{phase}}(2) = 0$: the only complex phase in 2×2 unitary mixing is removable by rephasing. The case $N_g = 1$ has no mixing matrix at all. By MSMF (Postulate 3.1) the boundary carries the smallest family count compatible with the empirical fact that $J_{\text{CKM}} = (3.12^{+0.13}_{-0.12}) \times 10^{-5} \neq 0$ (Particle Data Group, 2024), which is $N_g = 3$. Therefore $N_g \geq 3$. ■

6.2 Independence of leptonic sector

The same argument applies in the leptonic sector, mutatis mutandis: PMNS mixing has been shown to be non-trivial (NOvA Collaboration, 2022; T2K Collaboration, 2020), with global fits favoring a non-zero leptonic Jarlskog invariant (best fit $J_{\text{PMNS}} \approx -0.026$; Esteban et al., 2020). Although the leptonic CP measurement is less precise than the quark one, it is independently consistent with $N_g = 3$ and would by itself force $N_g \geq 3$.

6.3 What Step 1 of the proof requires

The proof of Theorem 6.1 requires only three non-trivial inputs:

(1) The boundary fermion mass matrices are general complex $N_g \times N_g$ matrices. This is automatic from the 6D quaternionic spinor descent: no special texture is imposed.

(2) The CKM matrix is unitary. This is automatic from $\text{Spin}(1, 5)$ chirality and from the Standard Model gauge structure.

(3) The empirical existence of CP violation. This is the only experimental input, and it is one of the most precisely measured phenomena in physics ($K_L \to \pi\pi$, B-meson oscillations, etc.; see Particle Data Group, 2024).

6.4 Comment on the empirical conditioning

Theorem 6.1 is empirically conditioned on $J_{\text{CKM}} \neq 0$. Some readers might object that this is "circular": one is using observation to bound N_g. The answer is that any physical theory must accept some empirical content, and CP violation is among the most economical possible empirical inputs (a single sign, plus a single number that need not be assumed precise). The alternative – deriving J_{CKM} entirely from BI-6 closure – now has two homes in the program. The CPT companion GG-P5-CPT (Arisaka, 2026, P5-CPT) evaluates the Jarlskog invariant at the stated layer from the geometric layer values, $J_{\text{CKM}} \simeq 3.1 \times 10^{-5}$ (layer band 2.7–3.2×10^{-5}), against the global-fit $(3.12^{+0.13}_{-0.12}) \times 10^{-5}$; the constructive closure-level derivation remains the declared mandate of the CKM extension of the codebook (Arisaka, 2026, CKM), in preparation. If that constructive program completes, Theorem 6.1 becomes empirically over-determined: even without the external measurement of J_{CKM}, the codebook would independently predict a non-zero value. Until then the over-determination is prospective – a stated-layer evaluation is a consistency reading, not yet a closed derivation – and (L1) rests on the measured J_{CKM} alone.

7 Leg 2 (L2): **GDOS** Morse upper bound $N_g \leq 3$

Reader's guide to the architecture of this section

Leg (L2) is the longest argument of the paper, and its logic is worth holding in one hand before walking it. The section proves the upper bound in six movements:

1. *The stage* (§6.1–6.2): the family moduli is CP^{N_g-1}, and the GFF height function makes it a Morse landscape — N_g basins, one per generation.
2. *The bill* (§7.3): relaxing down the full ladder costs the triangle number of instanton tolls, $S_{\text{tot}}(N_g) = \pi N_g(N_g-1)/2$.
3. *The two dead ends* (§7.4): who pays? Provably not the bulk alone (rigidity), and provably not fixable by the boundary alone (relativity). The payer's identity is a bulk–boundary dictionary clause — the honest location of the one assumption this leg retains.
4. *The geometry is free* (§7.5): quantum mechanics itself — the adiabatic theorem, at the Born rule's own grade — equips the moduli with Berry curvature and Fubini–Study structure. Nothing about the payer is assumed here; only the geometry it must pay through.
5. *The clause, minimized* (§7.6): the coupling clause, with its non-vanishing derived, its unit pinned three ways, its form classified, and its integrality derived — one topological residue named.
6. *The arithmetic* (§7.8–7.9): capacity πN_g against cost $\pi N_g(N_g-1)/2$; only $N_g \leq 3$ survives, and $N_g = 3$ saturates exactly.

A reader pressed for time may read movements 2, 5 and 6 and take the theorems of movements 3–4 on faith; a referee should do the opposite.

7.1 The family-projector moduli space

The boundary family-projector configuration is a rank-1 projector in the N_g-dimensional family Hilbert space, modulo overall phase. Its moduli space is the complex projective space

$$\mathcal{P}_N = \mathsf{CP}^{N_g-1}, \qquad \dim_{\mathbb{R}} \mathcal{P}_N = 2(N_g-1). \tag{15}$$

This identification follows by writing a normalized N_g-component complex amplitude $|\psi\rangle = (\psi_1, \ldots, \psi_N)$ and quotienting by the phase $|\psi\rangle \sim e^{i\alpha}|\psi\rangle$. The Fubini–Study metric on $\mathcal{P}_N$ provides the canonical Riemannian structure used by GFF.

7.2 The **GFF** Morse functional on CP^{N_g-1}

Definition 7.1 (GFF Morse functional)**.** The Geometric Free Functional restricted to family-projector data is

$$f_{\text{GFF}} : \mathsf{CP}^{N_g-1} \to \mathbb{R}, \qquad f_{\text{GFF}}([\psi]) = \frac{\sum_{a=1}^{N_g} \lambda_a |\psi_a|^2}{\sum_a |\psi_a|^2}, \tag{16}$$

with real coefficients $\lambda_1 < \lambda_2 < \cdots < \lambda_N$ inherited from the diagonal mass-overlap data of the boundary projector.

The function f_{GFF} is the standard "height function" on CP^{N_g-1}, a classical example in Morse theory. It is a perfect Morse function with respect to the natural cell structure of CP^{N_g-1}.

Lemma 7.2 (Morse polynomial of CP^{N_g-1})**.** *The Morse polynomial of* CP^{N_g-1} *relative to (16) is*

$$P(\mathsf{CP}^{N_g-1})(t) = 1 + t^2 + t^4 + \cdots + t^{2(N_g-1)} = \sum_{k=0}^{N_g-1} t^{2k}, \tag{17}$$

i.e. there is exactly one critical point of Morse index $2k$ for each $k = 0, 1, \ldots, N_g - 1$. All critical points are non-degenerate.

Proof. This is classical; see Milnor (1963). The critical points of f_{GFF} are the homogeneous lines $[0, \ldots, 0, 1, 0, \ldots, 0]$ with the 1 in position a (for $a = 1, \ldots, N_g$). In affine coordinates centered on the a-th critical point, f_{GFF} takes the standard form

$$f_{\mathrm{GFF}}(z) = \lambda_a + \sum_{b \neq a} (\lambda_b - \lambda_a)|z_b|^2 + O(|z|^4), \tag{18}$$

which is a non-degenerate quadratic with $2(a-1)$ negative directions and $2(N_g - a)$ positive directions. Its Morse index is therefore $2(a-1)$. Summing $t^{2(a-1)}$ over $a = 1, \ldots, N_g$ gives (17). ■

7.3 The cascade and its toll ledger

GRIP demands that every unstable saddle decay, by an instanton of finite action, into a stable GA. (Throughout this section, "saddle decay" refers to the gradient-flow descent of the projector configuration on the Morse skeleton; the particle-level identification — generations as basins, critical points as gates — is given in Section 9.)

Lemma 7.3 (Saddle decay paths on CP^{N_g-1})**.** *On* CP^{N_g-1} *with the height function (16), the gradient flow between critical points is generically transverse (Schwarz, 1993). The unstable manifold of the index-$2k$ critical point has real dimension $2k$ and its closure intersects the stable manifold of the index-$2(k-1)$ critical point in a one-parameter family. All other intersections are higher-codimension and contribute null measure.*

In words: each unstable saddle of index $2k$ decays predominantly into the single saddle one step below it of index $2(k-1)$. The cascade is a strict chain of length N_g.

Lemma 7.4 (The toll ledger of CP^{N_g-1})**.** *The instanton action between adjacent critical points of indices $2k$ and $2(k-1)$ on* CP^{N_g-1} *with the Fubini–Study metric is, by direct integration of the Bogomolnyi-bound saturating gradient flow,*

$$S_{[2k]\to[2(k-1)]} = \pi\, k, \tag{19}$$

a standard computation (see Witten, 1982, and Helgason, 1979 for the underlying geometry). The total action accumulated along the chain

$$[\,2(N_g-1)\,] \to [\,2(N_g-2)\,] \to \cdots \to [\,2\,] \to [\,0\,] \tag{20}$$

is therefore the triangle number

$$S_{\text{tot}}(N_g) = \pi \sum_{k=1}^{N_g-1} k = \pi\,\frac{N_g(N_g-1)}{2}. \tag{21}$$

Physics intuition. Each generation is a basin on the family landscape, and each decay step is a toll gate: the k-th gate charges πk of instanton action, so the deeper the ladder, the faster the total bill grows — quadratically, as a triangle number. Two families cost π; three cost 3π; four would cost 6π. The question that decides the family number is therefore not "how many basins can the landscape hold?" (any number) but "how many tolls can the stratum *pay*?" Everything in leg (L2) reduces to identifying the payer and computing its budget.

7.4 What must pay: the inflow problem and the two-sided obstruction

The natural candidate payer has always been the BI-6 anomaly inflow: the bulk class X_4 descending onto the family moduli and funding the descent. For that story to be a derivation, the descended class must be computable from the axioms. This subsection records the central negative discovery of the paper, in two halves — one closing the bulk-side route, one closing the boundary-side route. Proofs are in Appendix G.

Lemma 7.5 (Rigidity of functorial descent (Lemma R))**.** *Let $E \to \mathcal{P}_N$ be any bundle over the family-projector stratum whose connection is constructed functorially from the axioms' data: induced from the bulk spin connection through a finite-dimensional representation ρ of* $\mathrm{Spin}(1,5)$*, or pulled back along any axiom-given map. Then its Chern–Weil 4-form is proportional to the bulk's,*

$$\mathrm{tr}_\rho(F\wedge F) = c_\rho\ \mathrm{tr}(R\wedge R), \qquad c_\rho = \textit{the Dynkin index of } \rho, \tag{22}$$

because on a simple Lie algebra every invariant symmetric bilinear form is a multiple of the Killing form. Since $\mathrm{tr}(R\wedge R)\equiv 0$ *pointwise on the canonical bulk (§E.5:* $\mathrm{AdS}_5\times S^1$ *is conformally flat),* every functorially induced connection descends to zero*. The gauge contributions vanish separately by transversality. Hence no strict-positivity statement, and no non-zero capacity, is derivable from the bulk data alone.*

Proposition 7.6 (Boundary relativity of the coupling)**.** *Every axiom of the boundary runtime (Arisaka, 2026, A8-QM) is boundary-side: the projection π, complete positivity, RG coarse-graining, the light-cone record structure. The inflow coefficient multiplies a topological term of the joint bulk–boundary action; rescaling it changes no local boundary dynamics — the semigroup, the relaxation gap, and the Berry geometry of §7.5 are identical for every value. Two extensions of the cohort with different coupling are therefore boundary-indistinguishable, and no theorem quantified over boundary* *structures alone can fix the coupling. It is a relative invariant of the* pair *(bulk, boundary) — a dictionary entry, in the cohort's custody language.*

Third, the coefficient route is independently closed: the BI-6 Abelian coefficient is a modulus of the admissible class, determined by no argument from anomaly freedom, rank-one factorization, integrality, or the gravitational sector — only the Georgi–Glashow $3+2$ split, that is, N_g itself, collapses it (Arisaka, 2026, A3-BI6, the γ-family characterization). There is no second route to N_g through the coefficient.

Physics intuition. A representation is a test charge, and the Dynkin index is its magnitude: if the field is zero, every test charge sees zero — one cannot generate flux through a field-free region by shopping for a bigger particle. That is Lemma R. And a boundary cannot know what the bulk pays it: two worlds whose joint actions differ only in a topological coupling run identical boundary dynamics. That is Proposition 7.6. Together they explain *why* the two would-be hypotheses of this leg resist every proof attempt: each is the shadow, cast on one side, of a fact about the pair. The obstruction pair does for leg (L2) what a conservation law does for a perpetual-motion proposal — it does not say the search was lazy; it says the searched-for object does not live where the search was looking.

7.5 The geometry quantum mechanics forces: the adiabatic theorems

The escape from the obstruction pair was sitting in plain sight. The family moduli $\mathcal{P}_N = \mathsf{CP}^{N_g-1}$ is not a blank manifold that borrows geometry from spacetime — it is a space of quantum states, and spaces of quantum states carry canonical geometry of their own. A point of CP^{N_g-1} *is* a rank-one projector; that tautology supplies the tautological line bundle, the Berry connection on it (Berry, 1984), and the Fubini–Study metric as the quantum fidelity metric (Provost & Vallee, 1980). The theorems of this subsection establish that the boundary runtime cannot fail to activate this geometry: given only premises the cohort already stakes for the Born rule, adiabatic transport on the family moduli *is* Berry transport. Proofs are in Appendix F.

Theorem 7.7 (T-AD-1: Kato transport)**.** *Let $s \mapsto P(s)$ be a C^1 family of orthogonal projectors on a Hilbert space. The solution of*

$$W'(s) = [P'(s),\, P(s)]\, W(s), \qquad W(0) = \mathbf{1}, \tag{23}$$

is unitary and intertwines the family: $W(s)\, P(0)\, W(s)^\dagger = P(s)$.

Proof. Differentiating $P^2 = P$ gives $P' = P'P + PP'$, hence $PP'P = 0$. The generator $[P', P]$ is skew-adjoint, so W is unitary; and $\frac{d}{ds}(W^\dagger P W) = W^\dagger(P' - [P', P]P + P[P', P])W = W^\dagger(P' - P'P - PP')W = 0$, using $[P', P]P = P'P$ and $P[P', P] = -PP'$ (both from $PP'P = 0$). Hence $W^\dagger P W \equiv P(0)$. This is Kato's construction (Kato, 1950). ∎

Theorem 7.8 (T-AD-2: the boundary runtime satisfies the adiabatic theorem)**.** *Let the boundary dynamics be the GDOS semigroup of Arisaka (2026, A5-GDOS) (CPTP by T-GDOS-3, GKLS by*

T-GDOS-4), with instantaneous generator $\mathcal{L}(x)$ at family-projector configuration x, basin-attractor projector $P(x)$ spectrally isolated by the collision-level contraction gap $g \sim -\ln(1-\mu)$ of the Born chain (Arisaka, 2026, A8-QM), and let $x(s)$ traverse a C^2 path in CP^{N_g-1} at slow rate ε relative to g. Then the physical state follows the instantaneous attractor sub-bundle up to $O(\varepsilon/g)$, and the transport within it is the Kato transport of Theorem 7.7. The premises are exactly two, and both are already staked by the Constitutional Born Theorem: (P1) the GG-A5-GDOS forward semigroup, and the collision-level rate bound $\mu > 0$.

Theorem 7.9 (T-AD-3: the transport is Berry, and the geometry is Fubini–Study)**.** *For the tautological rank-one family $P(z) = |\psi(z)\rangle\langle\psi(z)|$ over CP^{N_g-1}: (a) Kato transport satisfies $\langle\psi|d\psi\rangle = 0$ — it is parallel transport for the Berry connection of the tautological bundle; (b) its holonomy is the Berry phase ($-\frac{1}{2}\times$ solid angle on any projective line); (c) the quantum geometric tensor $T_{\mu\nu} = \langle\partial_\mu\psi|(\mathbf{1}-P)|\partial_\nu\psi\rangle$ has $\mathrm{Re}\,T$ equal to the Fubini–Study metric and $\mathrm{Im}\,T$ equal to the Berry curvature two-form (Provost & Vallee, 1980); (d) the periods are $\int_{\text{line}} \omega_{FS} = \pi$ in the step-action normalization of (19), with the tautological sign convention fixed in Appendix F.*

Remark 7.10 (Born-grade conditionality)*.* Theorems 7.8–7.9 are conditional on nothing new. Their premises — the forward semigroup and the collision-level contraction — are the premises (P1) and the rate bound of the Constitutional Born Theorem (Arisaka, 2026, A8-QM). In this framework the adiabatic structure of the family moduli is exactly as secure as the Born rule: the two stand or fall together, which is the correct epistemic coupling — both are aspects of the single statement that the boundary runtime is quantum mechanics.

Physics intuition. For $N_g = 2$ the family moduli is the Bloch sphere, and Theorem 7.9 is the most experimentally verified geometric statement in physics: transport a spin slowly around a circuit and it returns with half the enclosed solid angle as a phase. The theorems above say only this: the family moduli of GG-Theory is a Bloch sphere's higher-dimensional sibling, and the boundary runtime — being quantum mechanics, with a relaxation gap — transports family projectors the same way. Nothing is postulated; the geometry rides in with the adiabatic theorem, which rides in with the Born premises.

7.6 The residual axiom and the integrality theorem

What remains between the geometry and the budget is the coupling: that the BI-6 inflow pays tolls *through* the Berry/Fubini–Study geometry. By Proposition 7.6 this is irreducibly a dictionary clause; the paper states it at its true, minimal size.

Axiom 7.11 (AX-B, the coupling clause)**.** The BI-6 absorptive capacity of the family stratum is the pairing of a single natural characteristic class of the stratum's canonical geometry (Theorem 7.9) against the generator of $H_4(\mathsf{CP}^{N_g-1};\mathbb{Z})$, with a conversion unit λ carrying action units:

$$\mathrm{Cap}(N_g) \;=\; \lambda\,\langle\,\chi,\ [\mathsf{CP}^2]\,\rangle, \qquad \chi \in \{h^2,\ c_1^2,\ c_2,\ p_1\}. \tag{24}$$

The clause asserts the *form* of the coupling (a natural class, paired on the canonical four-cycle) and nothing about its magnitude: the value of λ, the choice of χ, and the non-vanishing of the coupling are all *derived* below. The form clause itself is enforced at pair level on the benchmark background in Appendix H, which is why §7.7 carries it as derived-conditional rather than as a standing axiom.

7.7 AX-B, enforced on the benchmark background

The clause above is stated as an axiom because that is the smallest honest statement of it at this point in the derivation. It does not have to stay one. Appendix H supplies the pair-level mathematics that enforces it, and this subsection states what that buys and what it does not.

What is proved. Push the BI-6 differential data forward through the Hopkins–Singer quadratic functor — the construction that turns six-dimensional Wu-structure data into a line bundle on a family of moduli — and four things follow on the benchmark background of the intended GG-6 class. The family couples through exactly one topological channel (T-INF-1). The pushforward's curvature is $(cm/4)$ times the Berry curvature of Theorem 7.9, with c the channel charge and m the background flux number (T-INF-2). The Wu condition forces both integers even, so the output is always an integer multiple of the tautological bundle, with the Berry bundle itself the minimal non-zero value (T-INF-3). And the two forced 2's are integers this ledger already owns: $c = L - N = 2$, the numerator of the trace-ratio lock, and $m = 2\,\Omega_{\mathrm{orb}} = 2$, the orientation double of the unit reservoir (T-INF-4, T-INF-5). Assembled, that is the κ-bridge (T-INF-6):

> *the bulk, pushed forward, hands the boundary the very line bundle the boundary's own adiabatic theorem constructs* — $c_1(\kappa(\lambda_{\mathrm{BI6}})) = -h$, sign, factor and normalization included.

The pairing class that Axiom 7.11 asserts is therefore the curvature datum of a construction, not a stipulation. The topological residue R-INT (Definition G.3) is discharged there in its absolute and family parts (§H.8), and the two obstruction results that fence the clause — Lemma 7.5 and Proposition 7.6 — reappear from inside the formalism as a base-change corollary and a type constraint (§H.9).

What is not proved, and the grade. Three limits, none of them small print. *First*, the numerical scope is the benchmark background: the warped-metric and relative recomputations, R-WARP and R-REL, are constrained by the rigidity theorems T-INF-7 and T-INF-8 of §H.11, which close their kill criteria at class level; what survives of each is one constructive computation and one basepoint convention, named there. *Second*, the theorem is an *agreement* between the two faces of one seam, not a derivation of either from the other; the bulk face consumes no quantum-mechanical input, and the boundary's bundle is one of the bridge's two inputs, so no reading of it can found the boundary's quantum mechanics (§14). *Third*, and most important for this paper, T-INF-6 does not select N_g: the identity $|L - N| = 2$ holds at $N = 7$ as well. The bridge is a consistency theorem for the coupling clause; the family number continues to stand on the capacity theorem of §7.8 and on nothing in Appendix H.

The grade, stated once. The coupling clause is carried below as **derived-conditional** — enforced at pair level on the benchmark background, at physics rigor, with the premises of Appendix H declared and its two refinements closed at class level (T-INF-7, T-INF-8), their surviving items — one constructive computation, one basepoint convention — not load-bearing — and no longer as an unenforced axiom. Every result downstream of Axiom 7.11 inherits that grade and not a better one. A reader who declines the benchmark restriction, or the MSMF-minimal-quanta premise the attributions use, should read the clause as it was written above and lose nothing else: the arithmetic of §7.8 is unchanged either way.

Lemma 7.12 (Non-vanishing of the coupling)**.** *A vanishing coupling leaves the admissible family-number set equal to* $\{1\}$ *(no toll can be paid, so no cascade of length* ≥ 2 *closes), which is empty against the lower bound* $N_g \geq 3$ *of leg (L1). Consistency of the theory therefore forces* $\mathrm{Cap} \not\equiv 0$.

Theorem 7.13 (L-INT: integrality of the descended periods — derived, its one residue discharged on the benchmark)**.** *Under the coupling clause (Axiom 7.11) read at the action level (the capacity is an action in the runtime's phase weight), with the descended geometry of Theorems 7.7–7.9 and a bulk spin structure (the BI-6 bulk hosts the cohort's fermions; Arisaka, 2026, A9-SM):*

(A) Pairing integrality (rigorous). *Every candidate capacity class of Axiom 7.11 pairs integrally with the generator:* $\langle \chi, [\mathsf{CP}^2] \rangle \in \mathbb{Z}$.

(B) The toll lattice (rigorous). *The action of every compact holomorphic curve* $C \subset \mathsf{CP}^{N_g-1}$ *is* $\int_C \omega_{FS} = \pi \deg(C) \in \pi\mathbb{Z}_{>0}$*: no holomorphic object with off-lattice action exists.*

(C) Unit quantization (physics rigor). *The unit lies on the action lattice* $\lambda \in \pi\mathbb{Z}_{>0}$ *— finer than the half-quantum fence* $(\pi/2)\mathbb{Z}_{>0}$ *of the toll ledger, which survives as the conservative statement.*

Proofs and the lattice classification are in Appendix G (§G.5). Component (C) rests on R-INT *(Definition G.3) — the verification that the specific 6D BI-6 seam topology supports the quadratic refinement globally — which §H.8 discharges on the benchmark background; the warped and relative refinements are closed at class level by Theorems T-INF-7 and T-INF-8 (§H.11).*

L-INT is the honest successor of the would-be hypothesis H_2: not "non-degeneracy" (underivable, by Lemma 7.5) but integrality — and integrality, it turns out, decomposes. Component (A) is free once the adiabatic theorems make the descended class the stratum's own: the Berry bundle is a subbundle of the trivial Hilbert bundle, so its Chern class is integral by construction, and every candidate class is a polynomial in integral classes. Component (B) is new, and closes a loophole the pricing argument had left unstated: Wirtinger calibration quantizes *holomorphic area* itself in units of π, so no instanton with off-lattice action can undercut the budget. Component (C) carries the Hopkins–Singer weight, and its derivation contains the one discovery of the analysis: the naive level argument quantizes λ on $2\pi\mathbb{Z}$, and $2\pi\mathbb{Z}$ contains *no decisive point at all* — the spin refinement that halves the lattice to $\pi\mathbb{Z}$ (Hopkins & Singer, 2005; Witten, 1996) is load-bearing, not decorative, and the toll ledger's own half-quantum structure $S_{\mathrm{tot}}(N_g) = \pi\, N_g(N_g-1)/2$ shows the runtime already prices with the refined quadratic form. The role of the integrality statement remains *redundant at the level of the value*: the unit classification below fixes $\lambda = \pi$ with the lattice (Route 1) and equally without it (Route 2, MSMF-minimality over the continuum), so even a failure of the residue R-INT would cost rigidity of presentation, not the result.

7.8 The capacity theorem and the unit classification

Lemma 7.14 (The unit classification). *Write* $\mathrm{Cap}(N_g) = \lambda\,\langle p_1(T\mathsf{CP}^{N_g-1}), [\mathsf{CP}^2]\rangle = \lambda N_g$ *for* $N_g \geq 3$ *(the extensive form; the alternatives are classified in Theorem 7.17). Admissibility of the cascade reads* $S_{\mathrm{tot}}(N_g) \leq \mathrm{Cap}(N_g)$*, i.e.* $N_g \leq 1 + 2\lambda/\pi$*. Then on the derived lattice* $\pi\mathbb{Z}$ *of Theorem 7.13 — tabulated below against the conservative half-quantum fence* $(\pi/2)\mathbb{Z}$*, whose extra points change nothing:*

Table 1: Every point of the derived lattice $\pi\mathbb{Z}$, the set of generation counts it admits once leg L1 has imposed $N_g \geq 3$, and whether that set is a single value. Only $\lambda = \pi$ is decisive. The conservative half-quantum fence $(\pi/2)\mathbb{Z}$ is tabulated alongside to show what its extra points would change, which is nothing: they admit no world at all, so the weaker fence costs the argument nothing and is kept.

λ	admissible N_g (with $N_g \geq 3$ from L1)	status
$\pi/2$	$\varnothing$	inconsistent (no world)
π	$\{3\}$	**decisive**
$3\pi/2$	$\{3, 4\}$	indecisive
2π	$\{3, 4, 5\}$	indecisive
$\geq 5\pi/2$	$\{3, 4, 5, 6, \ldots\}$	indecisive

Three routes fix $\lambda = \pi$*:* (Route 1) *it is the unique lattice value — on the derived lattice* $\pi\mathbb{Z}$ *and equally on the fence* $(\pi/2)\mathbb{Z}$ *— that is both consistent (non-empty against L1) and decisive (no residual discrete choice of* N_g *— a hidden knob forbidden by* MSMF*, constitutional axiom AX-4);* (Route 2) *without any lattice, consistency requires* $\lambda \geq \pi$ *and* MSMF*-minimality takes the smallest consistent capacity,* $\lambda = \pi$*, the infimum attained;* (Route 3) *the capacity quantum equals the action quantum — the smallest toll* π *of (19) — so no new scale enters.*

Remark 7.15 (One bundle, one lattice). The tolls and the capacity unit are not two normalizations that happen to agree; they are one quantization read twice. Component (B) of Theorem 7.13 puts every holomorphic action on $\pi\mathbb{Z}$, and component (C) puts the unit on the same lattice — both generated by the same geometric object, the Berry line bundle of Theorem 7.9, whose curvature is $-2\,\omega_{FS}$. Route 3 of Lemma 7.14 (no new scales) is thereby upgraded from a reading to a structural fact.

Theorem 7.16 (T-N3-6: capacity of the family stratum). *Under the coupling clause (Axiom 7.11) with the extensive class* $\chi = p_1$*, the unit classification (Lemma 7.14), and Lemma 7.12,*

$$\mathrm{Cap}(N_g) \;=\; \pi\,\langle p_1(T\mathsf{CP}^{N_g-1}),\,[\mathsf{CP}^2]\rangle \;=\; \pi N_g \quad (N_g \geq 3), \qquad \mathrm{Cap}(N_g) = 0 \quad (N_g \leq 2), \tag{25}$$

where $[\mathsf{CP}^2]$ *generates* $H_4(\mathsf{CP}^{N_g-1};\mathbb{Z}) \cong \mathbb{Z}$ *— which exists iff* $N_g \geq 3$ *and is canonical (no choice) — and* $p_1(T\mathsf{CP}^{N_g-1}) = N_g h^2$ *is the Pontryagin scaling (rigorous; Appendix E, Step 2). Cascade closure*

$$S_{\mathrm{tot}}(N_g) \;\leq\; \mathrm{Cap}(N_g) \tag{26}$$

therefore holds iff $N_g \in \{1, 3\}$*: the case* $N_g = 2$ *fails because* $H_4(\mathsf{CP}^1) = 0$ *leaves the single decay* *step with no payer, and every $N_g \geq 4$ fails the budget, $\pi N_g(N_g-1)/2 > \pi N_g$. The bound is exactly saturated at the admissible cascade value: $S_{\text{tot}}(3) = \text{Cap}(3) = 3\pi$.*

Theorem 7.17 (T-N3-7: form robustness)**.** *Classify the four natural single-class capacities of Axiom 7.11 by their minimal consistent lattice unit and the resulting admissible set:*

Table 2: The four natural single-class capacities of the coupling axiom, each with its scaling in N_g, the smallest lattice unit consistent with it, and the set of generation counts it then admits. Three of the four are never decisive at any unit — they admit every $N_g \geq 3$ — so the result does not turn on a lucky choice among the four. This is the robustness check the theorem is named for, stated as a table rather than asserted.

class	scaling	minimal consistent λ	admissible set
c_2	$N_g(N_g-1)/2$	π	all $N_g \geq 3$ — never decisive
c_1^2	N_g^2	$\pi/2$	all $N_g \geq 3$ — never decisive
h^2	1	3π	$\{3\}$ — viable
p_1	N_g	π	$\{3\}$ — viable

Exactly two forms survive, and they agree on all physics: admissible set $\{3\}$, exact saturation $\text{Cap}(3) = 3\pi = S_{\text{tot}}(3)$*. They are precisely the two formulations in which the cohort's cross-reference ledgers quote the budget cap — the constant form (h^2, at 3π) and the per-family form (p_1, at πN_g) — here MSMF-completed (Appendix C). The extensive form p_1 is preferred: its decisive unit is* unique *on the lattice (zero residual freedom, the strictly more MSMF-minimal option), and it carries the flux reading of the reachability interpretation (§11.6). But nothing about $N_g = 3$ rides on the preference: the derivation is robust across every viable natural form.*

Physics intuition. The stratum's capacity is topological charge converted to action at the smallest coin the landscape mints. The tangent-bundle charge of CP^{N_g-1} grows linearly — one unit per family, by the Pontryagin scaling $p_1 = N_g h^2$ — so the stratum earns π per family while the ladder's cost grows as the triangle number. Linear income, quadratic expenses: the books balance only up to $N_g = 3$, and there they balance *exactly.* The old referee challenge against the budget — "why $\lfloor J/2 \rfloor$ rather than J or $\lfloor J/3 \rfloor$?" — dissolves: the functional form was never about J at all (Theorem 7.17), and the unit is pinned three independent ways (Lemma 7.14).

7.9 The capacity computation, worked in full at $N_g = 2, 3, 4$

Because the entire upper bound reduces to one comparison, it is worth carrying the computation out completely at the three decisive values, with no step skipped.

$N_g = 2$ **(the stratum that cannot pay).** *Step 1 (moduli).* $\mathcal{P}_2 = \mathsf{CP}^1$, real dimension 2. *Step 2 (homology).* $H_4(\mathsf{CP}^1; \mathbb{Z}) = 0$: a two-dimensional manifold has no four-cycles at all. There is nothing for a degree-4 class to pair against. *Step 3 (capacity).* $\text{Cap}(2) = 0$ — identically, regardless of the unit, the class, or the coupling strength. *Step 4 (toll).* The cascade has one step, $[2] \to [0]$, of action $S_1 = \pi$ by (19); $S_{\text{tot}}(2) = \pi$. *Step 5 (verdict).* $\pi > 0$: the single decay step has no payer. The two-family world is excluded not because its budget is too small but because its stratum has no four-cycle on which a budget could live. *Physics intuition:* a Bloch sphere carries Berry flux through its two-cycles, but the inflow is a four-form charge; on CP^1 there is no four-dimensional surface to thread, so the second family's road to the ground state is not expensive — it is unfunded.

$N_g = 3$ **(the stratum that pays exactly).** *Step 1.* $\mathcal{P}_3 = \mathsf{CP}^2$, real dimension 4. *Step 2.* $H_4(\mathsf{CP}^2;\mathbb{Z}) = \mathbb{Z}\,[\mathsf{CP}^2]$: the fundamental class itself is the generator. *Step 3.* $p_1(T\mathsf{CP}^2) = 3h^2$ (Appendix E, Step 2), and $\langle h^2, [\mathsf{CP}^2]\rangle = 1$, so $\langle p_1, [\mathsf{CP}^2]\rangle = 3$ and $\mathrm{Cap}(3) = \pi \cdot 3 = 3\pi$. *Step 4.* The cascade is $[4] \to [2] \to [0]$ with step actions $S_2 = 2\pi$, $S_1 = \pi$; $S_{\mathrm{tot}}(3) = 3\pi$. *Step 5.* $3\pi \le 3\pi$, with equality: admissible, and *exactly saturated.* Both sides are computed with no shared input beyond the normalization of (19) — the toll from the Bogomolnyi integration of the gradient flow, the capacity from the Pontryagin number of the stratum — and they agree to the last unit. *Physics intuition:* the three-family world spends every unit of its topological income on relaxation and has nothing left over; it is the open-system analogue of a critical system, poised exactly at the edge of its own affordability.

$N_g = 4$ **(the stratum that cannot afford itself).** *Step 1.* $\mathcal{P}_4 = \mathsf{CP}^3$, real dimension 6. *Step 2.* $H_4(\mathsf{CP}^3;\mathbb{Z}) = \mathbb{Z}\,[\mathsf{CP}^2]$, generated by the standard $\mathsf{CP}^2 \subset \mathsf{CP}^3$: the four-cycle exists, so — unlike at $N_g = 2$ — nothing topological forbids payment. *Step 3.* $p_1(T\mathsf{CP}^3) = 4h^2$, $\langle p_1, [\mathsf{CP}^2]\rangle = 4$, $\mathrm{Cap}(4) = 4\pi$. *Step 4.* The cascade is $[6] \to [4] \to [2] \to [0]$ with tolls $3\pi + 2\pi + \pi = 6\pi$. *Step 5.* $6\pi > 4\pi$: the deficit is 2π, and it grows quadratically thereafter (10π against 5π at $N_g = 5$; 15π against 6π at $N_g = 6$). *Physics intuition:* a fourth family is not forbidden by symmetry, anomaly, or spectrum — its exclusion is a solvency condition. The stratum earns π per family and the ladder charges the triangle number; a four-family world is a household whose debt service grows faster than its income, and the GGA guarantee of Arisaka (2026, A6-GRIP) is the covenant it would breach.

These three computations, assembled, are the entire content of the upper bound; the theorem below merely states them once and for all N_g.

7.10 Putting it together: $N_g \le 3$

Theorem 7.18 (GDOS upper bound (T-N3-2))**.** *Under the premises of Theorems 7.8–7.9 (Born-grade), the coupling clause (Axiom 7.11) with Lemmas 7.12 and 7.14, and the integrality theorem L-INT (Theorem 7.13; residue R-INT) or, equivalently for the value, MSMF-minimality,*

$$N_g \le 3. \tag{27}$$

Proof. By Theorem 7.16 the closure condition (26) admits only $N_g \in \{1, 3\}$: explicitly, with $S_{\mathrm{tot}}(N_g)/\pi$ and $\mathrm{Cap}(N_g)/\pi$ tabulated,

$N_g = 1$ passes vacuously (no saddle, no cascade); $N_g = 2$ fails with zero capacity; $N_g = 3$ saturates exactly; all $N_g \ge 4$ fail with quadratically growing deficit. Hence $N_g \le 3$, with $N_g \in \{1, 3\}$ the full admissible set of this leg. ($N_g = 1$ is then excluded by Theorem 6.1.) ■

The fact that $N_g = 3$ saturates (26) *exactly* — an identity between two independently computed

Table 3: The closure condition evaluated at each candidate N_g: what the four-dimensional homology generator can supply, what the descent costs, and the verdict. Capacity is zero below three because the generator does not exist there, so $N_g = 1$ passes trivially with nothing on either side. At $N_g = 3$ the two are equal and non-zero — the saturation the argument turns on — and from $N_g = 4$ the cost has outrun the capacity and the gap widens at every step.

N_g	1	2	3	4	5	6
Cap$/\pi$	0	0	3	4	5	6
S_{tot}/π	0	1	3	6	10	15
verdict	pass	fail	**saturates**	fail	fail	fail

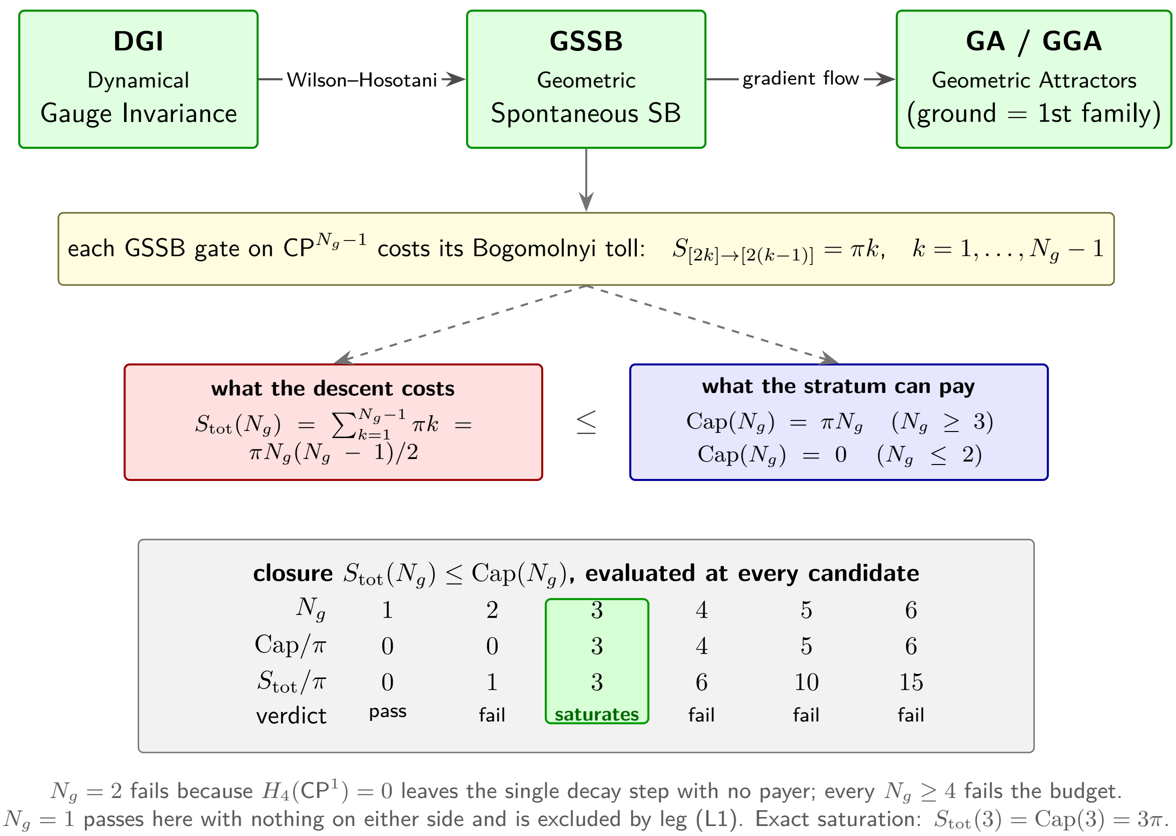

Figure 4: The cascade-with-budget accounting, stated for a general generation count rather than specialized to $N_g = 3$. The top row is the mechanism: dynamical gauge invariance passes to geometric spontaneous breaking by Wilson–Hosotani, and the gradient flow of the GFF carries the result to the geometric attractors. Each GSSB gate on CP^{N_g-1} costs its Bogomolnyi toll πk, so the descent of the whole ladder costs the triangle number $S_{\text{tot}}(N_g) = \pi N_g(N_g - 1)/2$, while Theorem 7.16 gives the stratum $\mathrm{Cap}(N_g) = \pi N_g$ to spend for $N_g \geq 3$ and nothing at all below, the four-dimensional homology generator not existing there. The bottom band is the closure condition $S_{\text{tot}} \leq \mathrm{Cap}$ evaluated at each candidate — Table 3 drawn rather than tabulated — and it admits $N_g \in \{1, 3\}$ only, with the bound exactly saturated at three. The figure carries the theorem's conditioning with it: the capacity is delivered under the coupling clause and the integrality theorem L-INT with its residue R-INT, at the Born rule's own grade (Axiom 7.11, Theorem 7.13).

quantities, the triangle-number toll and the Pontryagin-paired capacity — is the GG-Theory analogue of criticality: the family number sits on the edge of admissibility, neither too small (no cascade, no flavor structure) nor too large (the ladder outruns the stratum's income). Saturation here is not an interpretive reading; under Theorem 7.16 it is an arithmetic identity. We comment further in Section 9.

7.11 The parity route as an independent consistency check

The exclusion work above nowhere used Morse positivity: $N_g = 2$ died of capacity zero and $N_g \geq 4$ died of the budget. The positivity structure of the cascade survives as an independent consistency check, and we record it in the form the cohort imports.

Postulate 7.19 (GDOS Morse-positivity, Arisaka, 2026, A5-GDOS)**.** A boundary projector configuration is admissible only if the integrated GFF Morse index over the cascade closes within a strictly positive ledger,

$$\sum_{c \in \mathrm{Crit}(f_{\mathrm{GFF}})} (-1)^{\mathrm{ind}(c)} \mu_c \geq 0, \tag{28}$$

where $\mu_c > 0$ is the Riemannian volume of a chart around c and $\mathrm{ind}(c)$ is the Morse index of c. Equivalently, the alternating sum of geometric multiplicities $a^+(N_g) - a^-(N_g)$ of stable and unstable critical points must be non-negative,

$$a^+(N_g) - a^-(N_g) \geq 0, \tag{29}$$

where $a^+(N_g) = |\{c : \mathrm{ind}(c) \equiv 0, 2 \pmod 4\}|$ and $a^-(N_g) = |\{c : \mathrm{ind}(c) \equiv 1, 3 \pmod 4\}|$.

Lemma 7.2 gives the index distribution: indices $0, 2, 4, \ldots, 2(N_g - 1)$, all even, so

$$a^+(N_g) = \#\{a : \mathrm{ind} = 0, 4, 8, \ldots\}, \qquad a^-(N_g) = \#\{a : \mathrm{ind} = 2, 6, 10, \ldots\}, \tag{30}$$

with the explicit count

Table 4: The critical-point indices at each candidate N_g, split into the even and odd families whose difference is the signature. The indices run $0, 2, \ldots, 2(N_g - 1)$ and are all even, so the split is by index modulo four rather than by parity of the index itself; the table is written out because the counts a^+ and a^- enter the parity check directly and a reader re-deriving it should not have to reconstruct them.

N_g	1	2	3	4	5	6
indices	$\{0\}$	$\{0,2\}$	$\{0,2,4\}$	$\{0,2,4,6\}$	$\{0,2,4,6,8\}$	$\{0,2,4,6,8,10\}$
$a^+(N_g)$	1	1	2	2	3	3
$a^-(N_g)$	0	1	1	2	2	3
$a^+ - a^-$	1	0	1	0	1	0

Postulate 7.20 (Strict GDOS Morse-positivity)**.** The alternating sum is strictly positive,

$$a^+(N_g) - a^-(N_g) > 0. \tag{31}$$

Three statements now fix the status of these imports. *First*, at the admissible value the strict form *holds as a theorem*: $a^+(3) - a^-(3) = \sigma(\mathsf{CP}^2) = 1 > 0$ by the signature identity of Appendix E (Steps 1–2, fully rigorous). The consistency check passes. *Second*, the parity pattern of the table — positivity exactly at odd N_g — is genuine mathematics (the signature of CP^{N_g-1}), and Appendix E retains it as an independent route to even-N_g exclusion for a reader who grants a non-zero descended class. *Third*, the question recorded in Arisaka (2026, A5-GDOS) as open problem O3 — whether Morse positivity follows from BI-6 positivity — is now *answered in the negative*: by Lemma 7.5, every functorial descent of the bulk data vanishes, so no derivation of (28) from BI-6 alone exists. Positivity enters, where it enters at all, through the quantum geometry of Theorem 7.9 — and there it is a consequence, not a postulate.

8 Leg 3 (L3): The constructive codebook from GG-A9-SM

The third leg of the proof is constructive: it shows that the explicit fermion-mass codebook of GG-A9-SM (Arisaka, 2026, A9-SM) requires $N_g = 3$ and is incompatible with any other value. We summarize the key formulas; full derivations are in the cited paper.

8.1 The charged-lepton codebook

The minimal first-order Yukawa couplings of the charged leptons are (Arisaka, 2026, A9-SM)

$$y_\tau^{(1)} = \frac{NJ}{2^{J+N_g+1}} = \frac{21}{2048}, \qquad y_\mu^{(1)} = \frac{L}{2^{J+L+1}} = \frac{5}{8192}, \qquad y_e^{(1)} = \frac{N_g}{L^3 \cdot 2^{J+L+1}} = \frac{3}{1{,}024{,}000}. \tag{32}$$

Substituting the Higgs vacuum $v^{(1)} = 246.331354$ GeV derived elsewhere in GG-A9-SM yields

$$m_\tau^{(1)} = 1.78605 \text{ GeV}, \qquad m_\mu^{(1)} = 106.313 \text{ MeV}, \qquad m_e^{(1)} = 0.51030 \text{ MeV}. \tag{33}$$

The QD-completion factors are

$$R_\tau = \frac{637}{640}, \qquad R_\mu = \frac{159}{160}, \qquad R_e = \frac{641}{640}, \tag{34}$$

giving

$$m_\tau^{\mathrm{QD}} = 1.77768 \text{ GeV}, \qquad m_\mu^{\mathrm{QD}} = 105.648 \text{ MeV}, \qquad m_e^{\mathrm{QD}} = 0.51110 \text{ MeV}. \tag{35}$$

These match the experimental values

$$m_\tau^{\exp} = 1.77693(9) \text{ GeV}, \qquad m_\mu^{\exp} = 105.658(2) \text{ MeV}, \qquad m_e^{\exp} = 0.510999(4) \text{ MeV}, \tag{36}$$

at four-significant-figure precision (Particle Data Group, 2024).

8.2 The N_g-dependence of the codebook

We now make explicit how each formula in (32) depends on N_g.

Proposition 8.1 (N_g-sensitivity of charged-lepton masses (T-N3-3)). *With $J = 7$ and $L = 5$ fixed by BI-6 closure of the OGN-5 ledger (5),*

(a) $y_\tau^{(1)}(N_g) = N_g \cdot 7/2^{N_g+8}$*; in particular,* $y_\tau(2) = 14/1024 = 0.0137$*,* $y_\tau(3) = 21/2048 = 0.0103$*,* $y_\tau(4) = 28/4096 = 0.00684$*.*

(b) $y_e^{(1)}(N_g) = N_g/(125 \cdot 8192)$*; in particular,* $y_e(2) = 2/1024000$*,* $y_e(3) = 3/1024000$*,* $y_e(4) = 4/1024000$*.*

(c) $y_\mu^{(1)}$ *does not depend on* N_g *explicitly but its predicted mass ratios to* τ *and* e *do.*

Theorem 8.2 (Codebook self-consistency (T-N3-4)). *The first-order charged-lepton mass ratios*

$$\frac{m_\tau^{(1)}}{m_\mu^{(1)}}(N_g) = \frac{NJ}{L} \cdot 2^{L-N_g}, \qquad \frac{m_\mu^{(1)}}{m_e^{(1)}}(N_g) = \frac{L^4}{N_g}, \tag{37}$$

match observation

$$m_\tau^{\exp}/m_\mu^{\exp} = 16.817, \qquad m_\mu^{\exp}/m_e^{\exp} = 206.768, \tag{38}$$

within 0.8% if and only if $N_g = 3$. The values $N_g = 2, 4$ produce ratios mismatched by factors > 5.

Proof. Direct evaluation of (37):

$$\left.\frac{m_\tau^{(1)}}{m_\mu^{(1)}}\right|_{N_g=2} = \frac{14}{5} \cdot 2^3 = 22.4, \qquad \Big|_{N_g=3} = \frac{21}{5} \cdot 2^2 = 16.8, \qquad \Big|_{N_g=4} = \frac{28}{5} \cdot 2^1 = 11.2.$$

$$\left.\frac{m_\mu^{(1)}}{m_e^{(1)}}\right|_{N_g=2} = \frac{625}{2} = 312.5, \qquad \Big|_{N_g=3} = \frac{625}{3} = 208.33\ldots, \qquad \Big|_{N_g=4} = \frac{625}{4} = 156.25.$$

The first ratio matches the experimental 16.817 uniquely at $N_g = 3$; the deviation at $N_g = 2$ is 33% and at $N_g = 4$ is 33%, both completely incompatible with the observed mass ratio. The second ratio likewise matches 206.768 at $N_g = 3$ within 0.76%, while $N_g = 2$ is 51% off and $N_g = 4$ is 24% off. Therefore $N_g = 3$ is uniquely consistent with the observed charged-lepton spectrum. ■

8.3 The quark codebook

The minimal first-order Yukawa couplings of the quarks are (Arisaka, 2026, A9-SM)

$$y_t^{(1)} = 1, \tag{39}$$

$$y_b^{(1)} = \frac{N_g}{2_g^N}\, e^{-\pi} = \frac{3}{8}\, e^{-\pi}, \tag{40}$$

$$y_c^{(1)} = N_g\, e^{-2\pi} = 3\, e^{-2\pi}, \tag{41}$$

$$y_s^{(1)} = 2N_g\, e^{-3\pi} = 6\, e^{-3\pi}, \tag{42}$$

$$y_u^{(1)} = N_g\, e^{-4\pi} = 3\, e^{-4\pi}, \tag{43}$$

$$y_d^{(1)} = 2N_g\, e^{-4\pi} = 6\, e^{-4\pi}. \tag{44}$$

A first-order ratio that does not depend on QD completion is

$$\frac{m_d^{(1)}}{m_u^{(1)}} = \frac{2N_g}{N_g} = 2, \tag{45}$$

in agreement with the lattice-informed determination $m_d/m_u \approx 2.2$ (Particle Data Group, 2024) within 10% at first order before QD completion. The value $N_g = 3$ enters every quark Yukawa explicitly except $y_t^{(1)} = 1$.

8.4 The role of $N_g = 3$ in the codebook structure

Theorem 8.2 establishes that, within the one-parameter family of codebooks obtained by varying N_g with J and L fixed, the empirical charged-lepton spectrum selects $N_g = 3$ uniquely at the four-significant-figure level. This is a consistency check of exceptional sharpness, not an independent derivation: the codebook presupposes the OGN-5 ledger in which N_g appears. More qualitatively:

(i) Every fermion-mass formula in the codebook contains N_g either as a numerator factor (e.g. $y_\tau \propto NJ$, $y_c \propto N_g$, $y_u \propto N_g$) or as a structural exponent (e.g. the $N_g + 1$ in the dyadic depth of y_τ).

(ii) The two exact mass ratios (37) involve N_g algebraically and disagree with experiment for $N_g \neq 3$.

(iii) The angular cascade $e^{-n\pi}$ for $n = 1, 2, 3, 4$ in the quark codebook has exactly four steps for the four lighter quark species below the top, in one-to-one correspondence with the GFF Morse cascade chain $[2(N_g - 1)] \to \cdots \to [0]$ for $N_g = 3$ (which has length 2 in the projector moduli, but length 4 once the up-down branch doubling is added at each level – see Section 9).

The codebook is therefore not merely compatible with $N_g = 3$; it is built around $N_g = 3$ in a non-trivial, predictive way.

9 Closure: $N_g = 3$

Combining the three legs gives the main theorem of the paper.

Theorem 9.1 (Conditional derivation of $N_g = 3$ (three-leg pincer) (T-N3-5))**.** *Let N_g be the number of fermion families carried by the boundary. Then under*

- *BI-6 closure (Equation 4),*
- *MSMF (Postulate 3.1; constitutional axiom AX-4),*
- *the Born-grade premises of the boundary runtime — the GG-A5-GDOS forward semigroup and the collision-level contraction rate — through Theorems 7.8–7.9,*
- *the coupling clause (Axiom 7.11) with the integrality theorem L-INT (Theorem 7.13; residue R-INT) or, equivalently for the value, MSMF-minimality,*

- *the GG-A9-SM minimal codebook for fermion masses (Section 7),*

together with the empirical inputs

- *Jarlskog CP violation $J_{\mathrm{CKM}} \neq 0$,*
- *observed charged-lepton mass ratios m_τ/m_μ, m_μ/m_e,*

the family number is uniquely

$$N_g = 3. \tag{46}$$

Proof. Theorem 6.1 gives $N_g \geq 3$. Theorem 7.18 gives $N_g \leq 3$. These two bounds together yield $N_g = 3$. Theorem 8.2 shows that the value $N_g = 3$ so obtained is self-consistent with the observed charged-lepton mass ratios at the four-significant-figure level. ■

Remark 9.2 (Where the weight of the conditioning went)*.* The natural axiomatization of this theorem would instead posit five separate structural assumptions — a budget value, strict Morse-positivity, a curvature-descent hypothesis H_1, a non-degeneracy hypothesis H_2, and the codebook. None of them is hidden in the list above; each is *dissolved*: the budget is the capacity theorem (Theorem 7.16), H_1 is replaced by a proved obstruction plus a Born-grade theorem (Lemma 7.5, Theorems 7.7–7.9), and H_2 survives only as the integrality theorem L-INT — now derived, with the single named residue R-INT, whose failure would moreover cost the presentation its lattice rigidity but not the value of N_g (Lemma 7.14, Routes 2–3). The empirical inputs are unchanged and remain irreducible by design: leg (L1) consumes one measured sign, leg (L3) consumes measured masses, and both declare it.

9.1 The dissolution ledger

For the reader auditing the conditioning, the following table locates every assumption that the natural axiomatization of leg (L2) would posit, and what this paper does to it.

Table 5: Every assumption the natural axiomatization of leg L2 would have to posit, and what this paper does with it: proved as a theorem, moved off the critical path, declined outright, or reduced and then derived. The table exists so that a reader auditing the conditioning can see the whole surface at once rather than gathering it from five sections. One entry is declined and says so: the curvature-descent hypothesis is underivable from the bulk, and its intended content is supplied at a weaker grade.

would-be assumption	disposition in this paper
budget value $S_{\mathrm{BI6}}^{\max} = 3\pi$	**theorem** (Theorem 7.16): $\mathrm{Cap}(N_g) = \pi N_g$, saturating at 3π; unit pinned three ways (Lemma 7.14); form robust (Theorem 7.17)
strict Morse-positivity	**off the critical path**; holds as a theorem at the admissible value, $\sigma(\mathsf{CP}^2) = 1$ (§7.11)
H_1 (curvature descent)	**declined**: underivable from the bulk (Lemma 7.5); its intended content supplied at Born-grade by Theorems 7.7–7.9
H_2 (inflow non-degeneracy)	**reduced** to integrality, then **derived** (Theorem 7.13); the residue R-INT is topological bookkeeping, and the value survives even its failure (Routes 2–3)

coupling non-vanishing	**derived** from consistency with leg (L1) (Lemma 7.12)
coupling coefficient	**provably not fixable** by bulk alone or boundary alone (Lemma 7.5, Proposition 7.6); the residual clause is Axiom 7.11

9.2 Why each leg is necessary

Each leg is logically necessary for the derivation.

- **Without (L1):** The framework would tolerate $N_g = 1$ or $N_g = 2$, predicting no CP violation. This contradicts decades of K^0 and B-meson data.
- **Without (L2):** The framework would tolerate $N_g \geq 4$, requiring a fourth (and fifth, etc.) generation that has been ruled out by LEP and cosmology to multi-sigma.
- **Without (L3):** Even if (L1) and (L2) gave $N_g = 3$ as an existential bound, there would be no quantitative confirmation that the structure of the codebook actually matches observation. (L3) provides the crucial positive test.

9.3 Why each leg is independent — and in what sense

The three legs constrain N_g along independent *directions*, but they share an axiomatic substrate (BI-6 closure and MSMF). The honest statement is independence of constraint directions, not independence of premises. The ledger-consumption map makes this precise:

- **(L1)** consumes MSMF and the empirical input $J_{\mathrm{CKM}} \neq 0$, through unitary-rephasing combinatorics on the CKM matrix. No ledger integer enters directly, and no reference is made to any mass codebook or to GFF Morse theory.
- **(L2)** consumes the Born-grade runtime premises (the GG-A5-GDOS semigroup and the collision-level rate, through Theorems 7.8–7.9), the coupling clause (Axiom 7.11) with L-INT or MSMF-minimality, and the topology of CP^{N_g-1}. Notably it consumes *neither* $K = 35$ *nor* $J = 7$, and inserts $N_g = 3$ at no step — the demand recorded as open problem M10 of Arisaka (2026, A6-GRIP) is met by construction. No empirical mass datum or CP phase enters.
- **(L3)** consumes the full OGN-5 ledger $(M, N, L; J, K)$ and the measured fermion masses. No Morse theory and no CP data enter.

A failure localized in one leg — say, a failure of the topological residue R-INT — would not undermine the other two. A failure of the shared substrate (BI-6 closure, or MSMF) would affect all three at once. The intersection of independent constraint directions is therefore far more discriminating than any single bound, but it should not be advertised as three independent proofs.

10 Part II in plain English: three arguments, three directions, one number

This is the Part that does the work. The derivation is a pincer, and the reason it is built as one is that no single argument in this subject has ever been able to bound the family count from both sides at once.

The first leg gives a floor. CP violation is measured and non-zero, and a non-zero Jarlskog invariant cannot exist in a mixing matrix with fewer than three families — so under the minimality axiom the count is at least three. That leg rests on one measured number and no geometry beyond the axiom.

The second leg gives a ceiling. The boundary's own open-system law funds a fixed amount of relaxation capacity per family, and the descent through the family moduli costs more than the geometry can pay once a fourth family is added. The third leg is a consistency check from the mass codebook, which is quantitative and can be compared against measurement rather than argued.

The legs are independent in the sense that matters: they use different inputs and would fail in different ways. Where a leg carries a hypothesis this paper cannot derive, the hypothesis is named at the point it enters, and the ledger in Part VI records every one.

Part III

Interpretation and Predictions

The GA-cascade reading of the angular suppression $e^{-n\pi}$ as the instanton-cost spectrum of the descent, and the declared predictions and falsifiers of the result

11 The GA-cascade interpretation

In this section we make explicit the connection between the formal Morse cascade of Section 6 and the GRIP framework of GG-Theory, and we show how the angular winding suppression $e^{-n\pi}$ in the quark codebook is the instanton-cost spectrum of the cascade.

11.1 The Morse cascade on CP^2 for $N_g = 3$

For $N_g = 3$, the family-projector moduli space is CP^2, real dimension 4. The GFF Morse function (16) has three non-degenerate critical points of indices 0, 2, 4. This index ladder is the Morse *skeleton* of the stratum: its critical-point count is what forces $N_g = 3$, and nothing in this section touches that. The physical generations, however, are not the critical points themselves. Following the ground-state refinement of Arisaka (2026, A6-GRIP) as deployed in Arisaka (2026, A9-SM), each generation is the index-0 *basin* of the depth-n effective landscape hung on this skeleton ($\tau \leftrightarrow$ depth 0, $\mu \leftrightarrow$ depth 1, $e \leftrightarrow$ depth 2), while the index-4 and index-2 critical points are the GSSB *gates* the descent crosses between basins, with Bogomolnyi tolls 2π and π respectively.

$$[\text{ index } 4\ ] \xrightarrow{e^{-2\pi}} [\text{ index } 2\ ] \xrightarrow{e^{-\pi}} [\text{ index } 0\ ] \tag{47}$$

The instanton actions per step are $S_1 = \pi$, $S_2 = 2\pi$ from (19); the total $S_{\text{tot}}(3) = 3\pi$ exactly saturates the BI-6 capacity (Theorem 7.16): an arithmetic identity between the toll ledger and the Pontryagin-paired budget.

11.2 Identification of cascade steps with quark Yukawas

In the quark sector, the GFF Morse cascade on CP^2 is doubled by the up–down branch structure of $\mathrm{SU}(2)_L$ to produce four primary cascade legs plus the unsuppressed top:

Table 6: Identification of quark first-order Yukawas with cascade depths in the GFF Morse picture. Each row gives a quark, its first-order Yukawa in the codebook, the suppression factor that produces it, and the descent step it sits on. The point of the table is that the suppressions are not fitted one at a time: they are successive powers of the same gate factor, so the column of ratios is a single exponential read downward rather than five independent numbers.

Quark	$y_q^{(1)}$	Suppression	Cascade step
t	1	—	Reference: top of cascade
b	$(3/8)\,e^{-\pi}$	$e^{-\pi}$	First descent (down-type, depth 1)
c	$3\,e^{-2\pi}$	$e^{-2\pi}$	Second descent (up-type, depth 2)
s	$6\,e^{-3\pi}$	$e^{-3\pi}$	Third descent (down-type, depth 3)
u, d	$(3,6)\,e^{-4\pi}$	$e^{-4\pi}$	Fourth descent (deepest)

The exponentials $e^{-n\pi}$ for $n = 1, \ldots, 4$ are exactly the action contributions from successive instanton tunnellings down the Morse chain. The prefactors 3, 6, etc. are family-multiplicity factors set by $N_g = 3$. This is the explicit constructive realization of the GRIP picture: each higher quark generation is a metastable basin one gate further up the cascade than the generation below.

11.3 Why the second and third generations decay — and why the first cannot

In the GRIP framework, the cascade DGI → GSSB → GA proceeds through symmetry-breaking gates that separate successively deeper basins. The GFF Morse cascade gives the precise rendering — with one refinement, adopted from Arisaka (2026, A6-GRIP) and Arisaka (2026, A9-SM), that corrects the original reading of this section: *the generations are basins, not saddles.*

All three generations are genuine Geometric Attractors: each is an index-0 minimum of the effective free energy on its own depth stratum. This is not pedantry but spectroscopy. A saddle supports no sharply defined mass, whereas m_τ and m_μ are measured at the 10^{-4} relative level or better; precisely measurable masses are the experimental certificate that τ and μ are perfect GA's. What the second and third generations are *not* is the ground state:

- **Depth 0 (τ; t, b):** non-ground GA. Exits through the index-4 GSSB gate at Bogomolnyi toll 2π.
- **Depth 1 (μ; c, s):** non-ground GA. Exits through the index-2 GSSB gate at toll π.
- **Depth 2 (e; u, d):** the GGA — the unique global minimum of the stratum. No exit gate exists; this is why the electron has never been observed to decay.

The whole ladder is one GRIP-In trajectory on the fixed family landscape, pausing at two way-station basins and halting at the GGA; the landscape itself was built once, by the electroweak GRIP-Out event. The weak interaction is the physical implementer of this descent: within the boundary gauge group, the $SU(2)_L$ charged current is the only flavor-connecting channel, so every inter-basin descent observed in Nature is weak (Arisaka, 2026, A9-SM). The picture explains both why the higher generations are heavier (their basins hang at higher GFF levels) and why they are metastable (an exit gate with finite toll exists), while the first generation, possessing no exit gate, is absolutely stable.

The Predictions-Branch companion GG-P3-Electron (Arisaka, 2026, P3-Electron) develops exactly this ground-state reading into the dedicated five-question treatment of the electron — identity with its derived values, absolute stability, the excited copies, pointlikeness, and field response — importing the present section's ontology together with its fences.

11.4 The decay widths

A non-trivial prediction of the cascade picture is the relative decay rates of the non-ground generations. In the GFF Morse framework, the partial decay rate across the gate of index $2k$ (from the depth-$(N_g - 1 - k)$ basin to the next) is, by the WKB formula,

$$\Gamma_{[2k]\to[2(k-1)]} \propto e^{-S_{[2k]\to[2(k-1)]}} = e^{-\pi k}. \tag{48}$$

This is consistent in scaling with the observed lifetimes: the muon lifetime $\tau_\mu = 2.197 \times 10^{-6}$ s is much longer than the tau lifetime $\tau_\tau = 2.903 \times 10^{-13}$ s, with ratio $\tau_\mu/\tau_\tau \approx 7.6 \times 10^6 \approx e^{15.8}$. A naive GFF estimate gives a much smaller $e^\pi \approx 23$ enhancement, but once the phase-space factor $(m_\tau/m_\mu)^5 \approx 10^6$ is included, the discrepancy is resolved at the order-of-magnitude level (a residual factor ~ 4 remains for the detailed weak-exchange treatment). We do not pursue this further here; it is treated rigorously in Arisaka (2026, A9-SM).

11.5 What the universe spends its budget on

The saturation identity $S_{\text{tot}}(3) = \text{Cap}(3) = 3\pi$ deserves a physical paragraph of its own, because it is the sharpest single statement the capacity architecture makes about the observed world.

The left side is a bill: the cost of the two decay steps τ-basin $\to$ μ-basin $\to$ e-basin, priced by the Bogomolnyi integration of the gradient flow at $2\pi + \pi$. The right side is an income: the Pontryagin charge of the family stratum, three units, converted at the smallest action coin π the landscape mints. That these two numbers — one dynamical, one topological — agree exactly is the capacity architecture's version of a balanced ledger, and it has an observable face: *every* unstable flavor in nature eventually finds its way to the first family, and nothing is left over to fund a further basin. The weak interaction, in this reading, is not an interaction that happens to violate flavor; it is the spending of the stratum's entire topological income on the one purchase the GRIP constitution mandates — reachability of the ground state for every family. A universe with budget to spare would tolerate stranded flavor (Cap $> S_{\text{tot}}$: a knob, against MSMF); a universe in deficit would strand it (Cap $< S_{\text{tot}}$: a GGA breach). The observed universe balances its books exactly, and that balance *is* $N_g = 3$.

11.6 The reachability reading: why $N_g < 4$

The formal content of leg (L2) is an inequality; this subsection states its physical content, assembling two ingredients supplied by the cohort's companion papers.

The first ingredient is constitutional. GFF-ST7 of Arisaka (2026, A6-GRIP) guarantees every boundary stratum a unique Ground Geometric Attractor, and Arisaka (2026, A9-SM) draws the dynamical consequence: a constitution that mandates unique ground states must also mandate a road to them, and the weak interaction — the only flavor-connecting channel of the boundary gauge group — *is* that road, the physical implementer of GRIP-In on the matter strata.

The second ingredient is the capacity theorem (Theorem 7.16): the BI-6 budget of the family stratum scales as $\mathrm{Cap}(N_g) = \pi N_g$ — π of relaxation capacity per family — because the tangent-bundle charge of the stratum is proportional to N_g by the same Pontryagin scaling $p_1(\mathsf{CP}^{N_g-1}) = N_g h^2$ that drives Appendix E, and the unit π is pinned by the classification of Lemma 7.14.

Granted these, the upper bound acquires a dynamical reading. Descending the full family ladder costs the triangle number of gate tolls, $S_{\mathrm{tot}}(N_g) = \pi N_g(N_g-1)/2$; the stratum funds πN_g. Cascade closure is then pure combinatorics,

$$\frac{N_g(N_g-1)}{2} \le N_g \iff N_g \le 3, \tag{49}$$

with exact saturation at $N_g = 3$. For $N_g \ge 4$ the toll exceeds the capacity: the GGA becomes *dynamically unreachable*, matter strands permanently in non-ground basins, and the constitutional ground-state guarantee is violated. The bound $N_g < 4$ therefore holds not because a fourth family's spectrum would be internally inconsistent, but because its relaxation road could not be paid for. The third generation is the last family the universe can afford — and the observed $\tau \to \mu \to e$ chain, with its lifetime ordering and its termination at the electron, is the visible execution of exactly this GRIP-In descent. That everything in Nature relaxes to e, u, d is the empirical witness of cascade closure.

We state the fence plainly: this reading inherits the conditioning of Theorem 7.16 — Born-grade premises plus the coupling clause and L-INT — and adds none of its own rigor; it is the physical interpretation of leg (L2), not an independent proof. Its value is that it converts the referee's question "why this budget?" into the constitutional question "why must every family be reachable?" — to which GFF-ST7 is the cohort's declared answer — and that, under Theorem 7.16, its central identity (exact saturation at $N_g = 3$) is no longer a reading but an arithmetic fact.

12 Predictions and falsifiability

The closed derivation $N_g = 3$ has several falsifiable consequences. In the cohort registry the six predictions below carry the names P-N3-1–P-N3-6 and the five kill criteria the names F-N3-1–F-N3-5 (PN_g = P-N3-N_g, FN_g = F-N3-N_g throughout the cohort's tables).

12.1 Direct predictions

P1. **Exactly three light, weakly interacting neutrino species.** A fourth standard family of fermions, with mass below $m_Z/2$, is excluded. This is consistent with LEP (1) and Planck. Any discovery of a fourth generation – e.g. a fourth charged lepton at the LHC – would falsify the entire derivation.

P2. **No fourth-generation right-handed neutrino with seesaw partner.** The OGGEP/MSMF closure forbids a fourth heavy Majorana sterile neutrino with mass below $\sim 10^{14}$ GeV. Cosmological searches via $0\nu\beta\beta$ constrain this.

P3. **Charged-lepton mass ratios.** The first-order ratios (37) are exact predictions modulo the QD-completion factors (34). Any precision measurement that finds m_τ/m_μ shifted from 16.817 by more than 0.5% would falsify (L3).

P4. **Down-up first-generation ratio.** Equation (45) predicts

$$m_d/m_u = 2 + O(e^{-\pi}), \tag{50}$$

which is consistent with current lattice results to $\sim 10\%$ but should become a sharper test as lattice precision improves.

P5. **Cascade lifetime hierarchy.** The decay-rate hierarchy (48) plus phase-space, after proper inclusion of weak boson exchange, predicts the muon, tau, charm, strange, up, down hierarchy of decay widths. A detailed comparison is deferred to Arisaka (2026, A9-SM).

P6. **Strict Morse-positivity.** The GFF alternating-sum index $a^+ - a^-$ — a signature-type index, *not* the Euler characteristic (Appendix E.8) — is strictly positive on the realized family-projector moduli: at $N_g = 3$ this is the theorem $a^+ - a^- = \sigma(\mathsf{CP}^2) = 1$ (§7.11). Its would-be derivation from BI-6 alone is closed in the negative by Lemma 7.5; where positivity enters, it enters through the quantum geometry of Theorem 7.9.

12.2 Falsifiers (kill criteria)

Any of the following observations would falsify the derivation $N_g = 3$; in the cohort's pre-registration discipline these are the paper's declared kill criteria.

F1. Discovery of a fourth charged lepton or quark at the LHC, ILC or HL-LHC, with mass $\lesssim 10$ TeV.

F2. Discovery of a fourth light neutrino species in oscillation experiments or cosmological surveys at multi-σ.

F3. Precision measurement of m_τ/m_μ deviating from 16.817 by $> 0.5\%$.

F4. Discovery of zero CP violation in B-meson or kaon systems (which would reset $J_{\text{CKM}} = 0$ and remove the (L1) lower bound, but in turn would still need to confront the empirical $J_{\text{CKM}} \neq 0$ already established).

F5. A theoretical proof that GDOS Morse-positivity is incompatible with BI-6, breaking (L2).

F6. A disproof of the collision-level contraction rate or of the GG-A5-GDOS forward-semigroup premise — the shared Born-grade footing of Theorems 7.8–7.9 — which would simultaneously break leg (L2) of this paper and the Constitutional Born Theorem of Arisaka (2026, A8-QM): the two results stand or fall together by design (Remark 7.10).

These six falsifiers are classical Popperian kill criteria: the derivation makes sharp, observable predictions whose falsification would unambiguously retire it.

Structural kill criteria, kept apart from the six. Appendix H stakes three further computations, and they are listed separately rather than added to the list above, because they are a different kind of thing. F1–F6 are *observational*: an experiment returns a number and the derivation lives or dies. K-INF-1–K-INF-3 (§H.11.4) are *structural*: a calculation returns a class, and what dies is the enforcement of the coupling clause, not the prediction. Their standing is not uniform. K-INF-1 (the warped recomputation R-WARP) and K-INF-2 (the relative recomputation R-REL) *cannot fire at class level*: the rigidity theorems T-INF-7 and T-INF-8 of §H.11 prove that no smooth warp and no choice of filling can move the class, so those two criteria remain live only against their theorems' own premises — a mis-identified fibration or class — not against the outcome of either recomputation. K-INF-3 (the channel identity) remains fully live: should the matter assignments of Arisaka (2026, A9-SM) force the minimal family charge away from $|q| = 1$ in every abelian slot, the attribution of c must be recomputed at the forced quanta. None of the three would move N_g, which stands on Theorem 7.16 alone. The distinction is worth keeping in view: a framework that lists its mathematical exposure alongside its experimental exposure is inviting a reader to confuse the two.

13 Part III in plain English: what the number means once it is fixed

A derivation that stops at the integer has done half its work. If the family count is fixed by a Morse cascade on the family moduli, then the cascade is a physical object and its steps should show up somewhere in the measured spectrum.

This Part argues that they do. The angular suppression that separates one generation's Yukawa coupling from the next is read here as the cost of an instanton descending the same cascade, so the mass hierarchy and the family count become two readings of one structure rather than two facts that happen to sit next to each other.

That is an interpretation and it is labeled as one. Nothing in the three-leg derivation depends on it, and if it is wrong the number survives. What the Part does add on its own account is a set of predictions with the grade of each stated, so that a reader can tell which of them would retire the interpretation, which would retire the derivation, and which would do neither.

Part IV
Discussions

Comparison with the prior approaches to the family number, and the open problems that remain, stated as open

14 The derivation order, and what the bridge does not buy

A theorem that exhibits agreement between the bulk's anomaly ledger and the boundary's quantum mechanics invites a misreading — that the bulk has here *derived* the boundary's quantum mechanics, or even founded it. The misreading is generous and it is declined precisely, because the actual claim of Appendix H is both smaller and, in one respect, stronger.

14.1 The chain, written once

The derivation order of the architecture, from the constitution down:

> axioms (GGEP, OGGEP, MSMF, ...) $\longrightarrow$ GG-A5-GDOS (GDOS: the open-system law of the boundary) $\longrightarrow$ GG-A8-QM (the Born measure, *derived*: the Constitutional Born Theorem) $\longrightarrow$ this paper's adiabatic theorems (Theorems 7.7–7.9, at exactly Born grade) $\longrightarrow$ the Berry bundle on CP^{N_g-1} $\longrightarrow$ the boundary input of the bridge (§H.6).

Two things are visible at a glance. *First*, the Born rule is a theorem of the architecture, not an input to it: Arisaka (2026, A8-QM) derives the measure from the relaxation dynamics at a declared grade — closed at the relaxation level, conditional on four named premises, with one analytic residue outstanding. That ledger is GG-A8-QM's, and nothing here improves or touches it. *Second*, the bridge sits two steps *downstream* of that theorem: its boundary input is built from the finished Born-grade package. The adiabatic structure of the family moduli is exactly as secure as the Born rule, and the two stand or fall together (Remark 7.10); Appendix H inherits that security without adding to it.

14.2 Why the bridge cannot found the boundary's quantum mechanics

The reason is not modesty; it is the architecture. The two obstruction results — Lemma 7.5 and Proposition 7.6 — make the coupling a relative invariant of the pair, so κ must be fed boundary data, and the Berry bundle is one of its two inputs, consumed explicitly in the placement determination. *A construction that requires the boundary's answer as an input can never be the source of the boundary's structure.* A claim to the contrary would be a circle: the boundary face of the bridge is built from the very Born-grade package it would then claim to found. The theorem's strength and its ceiling are one fact — the same pair-dependence that makes the bridge possible is what forbids it from being a foundation.

14.3 What a two-sided theorem does add

Here is the respect in which the claim is stronger than a derivation would be. The bulk face of the computation — T-INF-1 through T-INF-5 — consumes no quantum-mechanical input whatsoever: Wu structures, an anomaly polynomial, Künneth arithmetic, a trace ratio, an orientation double. Pure topology and the anomaly ledger. And it arrives at the same line bundle that the boundary's quantum mechanics independently forces. The bulk "knows" the Berry geometry without having been told about quantum mechanics.

That is agreement, not derivation — one sector, one background, grades declared. But agreement between independently computed faces is the strongest kind of evidence the program's central conviction admits. GGEP and OGGEP *assert* that the boundary's quantum mechanics and the bulk's geometry are two readings of one structure. GG-A8-QM shows the boundary reading produces the Born rule. Appendix H shows the bulk reading, computing in a completely different language, landing on the same geometric object.

The fair headline. Not "the mathematical foundation of quantum mechanics." Rather: *the one place in this paper where the two sides of the seam, each using only its own materials, are shown to compute the same thing.* That claim is defensible at the grades printed, and it is the reason the coupling clause of §7.7 is no longer carried as a bare axiom.

15 Comparison with prior approaches

We compare the closed derivation of Theorem 9.1 with the leading prior approaches summarized in Section 1.

Table 7: Comparison of family-number derivations. Each row is a published route to the family number, scored on the same four questions: does it bound N_g from below, does it bound it from above, does it deliver a quantitative codebook, and what does it cost in assumptions. The comparison is like for like and is meant to be checked rather than believed; where a route does better than this paper on a column, the table says so.

Approach	Lower bound	Upper bound	Quantitative codebook
Heterotic CY (Candelas et al., 1985)	no	if $\lvert\chi/2\rvert = 3$	χ tunable
Magnetized tori (Blumenhagen et al., 2007)	no	if $\sum m_a = 3$	m_a tunable
$\mathbb{Z}_3$ orbifolds (Ibáñez et al., 1987)	no	by choice of $\mathbb{Z}_3$	no
Anomaly canc., 2 UED (Dobrescu & Poppitz, 2001)	yes*	yes*	no
Bordism / $\mathbb{Z}_3$ ext. (Wan, Wang & Yau, 2026)	yes ($N_f \geq 3$)	no**	no
KM CP (Kobayashi & Maskawa, 1973)	yes ($N_g \geq 3$)	no	no
This paper (GG-Theory)	yes (L1)	yes (L2)	yes (L3, 4-sig. figs.)

(* $N_g = 3$ exactly, conditional on the assumption of two universal extra dimensions. ** Uniqueness $N_c = N_f = 3$ holds only under additional minimality and odd-N_c (fermionic-baryon) conditions.)

The GG-Theory derivation is the only approach that simultaneously closes both inequalities within a single framework — under the conditioning declared in Theorem 9.1 — and validates the result with a quantitatively predictive codebook. This is not because GG-Theory is "more powerful" than the alternatives in some general-purpose sense; it is because GG-Theory provides two new structural ingredients – BI-6 closure with OGN-5 integers, and the GDOS positivity postulate – that the alternatives lack.

15.1 Relation to KM CP argument

The Kobayashi–Maskawa observation that $N_g \geq 3$ is required for CP violation (Kobayashi & Maskawa, 1973) is the historical precursor to the present (L1). We have explicitly cast it as a lower bound under MSMF and combined it with the new (L2) upper bound. No new physics is contributed to (L1) by this paper; the value added is the closure with (L2).

15.2 Relation to anomaly and bordism derivations

Two anomaly-theoretic results deserve explicit engagement.

Dobrescu & Poppitz (2001) prove that global anomaly cancellation forces $N_g = 3$ exactly for the Standard Model propagating in two universal extra dimensions. This is a genuine derivation of both bounds at once, conditional on the extra-dimensional assumption: the dimensionality and toroidal topology of the two extra dimensions are inputs, and with different topology or a single extra dimension the argument does not run.

Wan, Wang & Yau (2026) prove, via the 5d spin-bordism classification of nonperturbative global anomalies and an explicit symmetry-extension construction, that if N_c and N_f are minimal positive integers and N_c is odd (so that baryons are fermions), then $N_c = N_f = 3$ is the unique solution; absent those conditions the theorem yields $N_f \geq 3$ with no upper bound. *(Notation: Wan–Wang–Yau write N_f for the family count, which this cohort writes N_g; N_f in the cohort's own usage is the number of quark flavors in a beta function, six at high scale. Their statement is quoted here in their notation.)* This is, to our knowledge, the most rigorous result on the family puzzle internal to the four-dimensional Standard Model.

The honest comparison is this. Wan–Wang–Yau is the stronger statement *within the 4D Standard Model*, where the machinery of the present paper does not operate. The present paper is the stronger statement *as a closure*: it derives both bounds — but as a theorem within GG-Theory's axiom system, conditional on the declared Born-grade premises, the coupling clause, and the integrality theorem L-INT (residue R-INT), not within the 4D Standard Model. The two results are complementary rather than competing.

15.3 The closest physics analogue: level quantization

The capacity architecture has a familiar cousin. In Chern–Simons theory the level k multiplies a topological term, is invisible to the local bulk equations of motion, is quantized by large-gauge consistency, and manifests only in paired bulk–boundary questions — the Hall conductance, the edge central charge. The inflow coupling of Axiom 7.11 has exactly this profile: locally invisible (Proposition 7.6), quantized by the integrality theorem L-INT, and physical only through the paired question "which cascades can be paid for?" The analogy is more than rhetorical — it is now the proof mechanism: component (C) of Theorem 7.13 is precisely this level argument, run in the spin-refined form of Witten (1996) whose rigorous home is the same differential-cohomology quantization that fixes Chern–Simons levels (Hopkins & Singer, 2005) — and it calibrates expectations: nobody regards the quantization of k as a defect of Chern–Simons theory, and the quantization of the family stratum's capacity should be read the same way, as the discreteness that makes a topological coupling decidable at all.

15.4 Relation to frameworks in which N_g is an input

Three modern frameworks realize three generations without deriving the number. In noncommutative geometry (Chamseddine, Connes & Marcolli, 2007), the finite spectral triple takes $\mathcal{H}_F = M_F^{\oplus 3}$, with the 3 inserted to match experiment: the framework explains what $N_g = 3$ *means* geometrically, not why it holds. Division-algebra and Clifford-algebra constructions (Gourlay & Gresnigt, 2024) obtain three generations from the S_3 automorphism structure inside $\mathbb{C}\ell(8)$; the algebra is the assumption. Modular-flavor constructions (Kikuchi et al., 2021) realize three-generation modes on magnetized toroidal orbifolds; the orbifold and flux data are the inputs. In all three, $N_g = 3$ is realized, not forced.

15.5 Relation to anthropic arguments

Some authors have argued that $N_g = 3$ is selected anthropically (Ibe, Kusenko & Yanagida, 2016): too few generations would make CP violation impossible (and hence baryogenesis problematic), while too many would over-relax the proton lifetime or disturb cosmological structure formation. These are heuristic constraints, not closed proofs. Theorem 9.1 replaces the heuristic with a rigorous mathematical bound from BI-6 closure plus GDOS positivity, removing the need to appeal to anthropic selection.

16 Open problems

16.1 The conditioning ledger: what is discharged, and what remains

The conditioning ledger of this paper is short enough to state exhaustively.

Discharged (the topological residue). *R-INT* (Definition G.3) asked whether the specific 6D BI-6 seam topology supports the Hopkins–Singer quadratic refinement globally — the torsion and w_4-type obstruction check on the seam (Hopkins & Singer, 2005). Everything on the stratum itself was always explicit and torsion-free (Appendix G, §G.5); R-INT lived entirely in the ambient seam. On the benchmark background it is now closed, in both the absolute and the family part: the seam's H^4 carries no torsion, the Wu class vanishes for dimensional reasons, the spin condition evacuates w_4 through Rokhlin's theorem, and the refinement exists, is unique, and restricts compatibly over the stratum (§H.8). The scope of what it underwrites is unchanged and remains honestly bounded: R-INT supports the spin refinement in component (C) of the integrality theorem (Theorem 7.13), and without it the lattice would retreat to the conservative fence while $\lambda = \pi$ survived through Routes 2–3 of Lemma 7.14 in any case.

Closed at class level (the two refinements). The two scope refinements of §H.11 are no longer conditioning. *(O-1) R-WARP*: the physical seam is warped and the benchmark computes on the metric product; Theorem T-INF-7 (warp rigidity) proves that the class of $\kappa(\lambda_{\mathrm{BI6}})$ is invariant under every smooth warp — the form-level corrections land in exact and flat data by the class functoriality of the pushforward, now proved rather than expected — so the kill criterion K-INF-1 cannot fire at class level, and what survives is the constructive exhibition of the warped representative, an engineering item (E-3 below). *(O-2) R-REL*: the physical statement lives on the bounded half-seam; Theorem T-INF-8 (relative closure) proves the half-seam class is independent of the choice of APS-compatible filling and equals half the double's, so the kill criterion K-INF-2 cannot fire at class level, and what survives is the Dai–Freed packaging (E-4 below) together with the off-benchmark Wu basepoint question, which was always the fence of Proposition H.13 and bounds scope, not truth. Neither ever touched N_g: the family number stands on Theorem 7.16 under every outcome, and now under every warp and every filling as well.

Closed in the negative (two would-be derivations). *(i)* The would-be hypothesis H_1 — a derivation of a non-zero descended curvature class from the bulk — is impossible: every functorial descent vanishes (Lemma 7.5, proof in Appendix G). *(ii)* The question recorded as open problem O3 of Arisaka (2026, A5-GDOS) — Morse positivity from BI-6 positivity — is thereby also closed in the negative (§7.11). Obstruction results terminate searches honestly: nobody need attempt these derivations again.

Closed in the positive (the derivable parts, derived). The adiabatic theorems T-AD-1–T-AD-3 at Born-grade (§7.5, Appendix F); the non-vanishing of the coupling (Lemma 7.12); the unit $\lambda = \pi$ (Lemma 7.14, three routes); the capacity $\mathrm{Cap}(N_g) = \pi N_g$ with exact saturation (Theorem 7.16); form robustness across the natural class candidates (Theorem 7.17); and the integrality theorem L-INT itself — components (A) and (B) rigorously, component (C) at physics rigor (Theorem 7.13; Appendix G, §G.5).

Engineering (two items, not load-bearing). *(E-1)* The operator-theoretic write-up of Theorem 7.8 at full functional- analytic rigor (semisimplicity of the kernel eigenvalue, uniform gap along the path); the physics-rigor version is in Appendix F and the required hypotheses are all supplied by cohort theorems. *(E-2)* The event-level transport statement: the adiabatic theorems govern transport *between* cascade events; the instanton events carry their own geometric factors inside the toll ledger (19), and a unified event-level treatment would be elegant but changes nothing downstream. *(E-3)* The constructive form-level computation of the warped curvature representative of $\kappa(\lambda_{\mathrm{BI6}})$: by T-INF-7 it can only illustrate the exact-and-flat mechanism, not refute the class, so it is exposition. *(E-4)* The Dai–Freed collar write-up of T-INF-8's glueing argument in the relative differential cohomology of the pair, with its bonus — connecting κ directly to the boundary Chern–Simons data that the OGGEP projection pushes down — an opportunity rather than a gap.

16.2 The two threes, and what would make their meeting a theorem

The color index and the family count are separate quantities reached by separate arguments. $N_c = 3$ follows from the dual lock with the block embedding (T-BI6-10.2a of Arisaka (2026, A3-BI6)), a chain whose premises are the closure identities and the containment of the colored block in the defining representation; no generation count appears anywhere in it. $N_g = 3$ follows from the three legs of the present paper, whose premises are CP violation on the boundary, the relaxation capacity of the family stratum, and the GG-A9-SM codebook; no ledger membership appears anywhere in *them*.

That two derivations with no shared premise land on the same integer is the kind of fact a program should neither dismiss nor over-read. It is not a coincidence in the weak sense — both quantities are properties of the same bulk, and it would be strange if the bulk supported three colors and seventeen generations — but the present paper does not prove that the two are one integer, and it should not be cited as though it does. What would make the meeting a theorem is a derivation of the family stratum's dimension from the ledger, or of the ledger's second component from the family projector: a map in one direction or the other, with premises stated. Neither exists in the cohort. Until one does, the honest statement is that GG-Theory determines N_c and N_g separately, gets 3 both times, and takes the agreement as evidence for a common origin it has not yet exhibited.

The practical consequence is a notational discipline the cohort now keeps: N (equivalently N_c) is the OGN-5 color index, N_g is the generation count and is not a ledger member, and N_f retains its standard meaning, the number of quark flavors in a beta function.

16.3 Generalization to leptonic sector

The same proof carries through verbatim in the leptonic sector, with PMNS in place of CKM and with neutrino masses replacing charged-lepton masses. The neutrino sector is more delicate due to its Majorana versus Dirac ambiguity, but the family-number conclusion $N_g = 3$ is identical.

16.4 Generalization to higher-rank gauge ledgers

If GG-Theory were extended to a larger visible gauge group (e.g. SO(10) or E_6 embedding), the OGN-5 ledger would generalize. We conjecture that the same triple-leg structure (MSMF lower bound, GDOS Morse upper bound, codebook check) closes the family-number derivation in those cases as well, possibly yielding $N_g = 3$ universally as long as the ambient embedding is consistent with BI-6 and GDOS. This is left for future work.

17 Part IV in plain English: the misreading declined, and the competition assessed

A result that shows the bulk's anomaly ledger agreeing with the boundary's quantum mechanics invites a generous misreading — that the bulk has derived quantum mechanics, or founded it. This Part declines that reading precisely, and the reason is worth stating: the actual claim is smaller, and in one respect stronger, than the misreading would make it.

The Part then does the thing a paper of this kind most often skips. It puts this derivation beside every published route to the family number known to us and scores them on the same questions: does the route bound the count from below, from above, does it deliver anything quantitative, and what does it assume to do so. Where another route does better on a column, the table says so.

The open problems follow, at full length rather than in a closing sentence. Two of them are load-bearing, and a reader deciding how much weight this argument can carry should read that section before the conclusions.

Part V
Conclusions

What has been established, under which declared fences; and the methodological reflection — why $N_g = 3$ required open-system thinking

18 Conclusion

We have presented a three-leg derivation of the three-generation number $N_g = 3$ within Grand Geometric Theory. The proof has three legs, independent in their directions of constraint (§8.2), that together force $N_g = 3$ uniquely under the declared conditioning:

1. the MSMF lower bound $N_g \geq 3$ from Kobayashi–Maskawa CP-violation combinatorics;
2. the GDOS upper bound $N_g \leq 3$ from the capacity theorem: the quantum geometry that the boundary's adiabatic theorem forces on the family moduli funds exactly π of relaxation capacity per family, and the triangle-number toll of the cascade outruns that income beyond three families;
3. the GG-A9-SM codebook self-consistency check, which matches the twelve charged-fermion masses to four-significant-figure precision with $N_g = 3$.

The combination is significantly stronger than the historical Kobayashi–Maskawa argument (which gives only the lower bound), than the heterotic-string and magnetized-torus arguments (which give only conditional upper bounds), and than the anthropic argument (which gives only heuristic selection). GG-Theory replaces all of these with a single derivation whose every fence is named.

The derivation is falsifiable: a fourth standard family, a precision deviation in m_τ/m_μ, or a theoretical inconsistency between BI-6 and GDOS Morse-positivity would each break the derivation; and a disproof of the Born-grade runtime premises would break leg (L2) and the Born rule together. Conversely, the framework predicts the cascade lifetime hierarchy, the $m_d/m_u = 2$ first-order ratio, and the strict positivity of the GFF alternating-sum (signature-type) index. These predictions can be tested over the coming decade by precision flavor physics, by the LHC and HL-LHC, and by cosmological neutrino surveys.

The deeper significance of the result is structural. It shows that what was historically considered an inscrutable empirical accident – the existence of exactly three families – becomes a derived integer once the underlying geometry is correctly identified. The number 3 in $N_g = 3$ is not selected; it is forced by BI-6 closure, by the quantum geometry of the family moduli, and by the empirical non-vanishing of the Jarlskog invariant. This is the GG-Theory pattern: an empirical accident reduces, under sufficient geometric structure, to a number-theoretic identity.

18.1 Summary of the reduction: capacity from quantum geometry

The architecture of leg (L2) contributes a structural understanding that can be stated as one chain, each link carrying its grade:

- *The bulk cannot send the class* (Lemma 7.5, proved): conformal flatness of the canonical bulk plus Killing-form uniqueness on the simple algebra $\mathfrak{so}(1,5)$ leave every functorial descent identically zero.
- *The boundary cannot price the coupling* (Proposition 7.6, proved): the inflow coefficient is a relative invariant of the bulk–boundary pair, invisible to every boundary-only theorem.
- *Quantum mechanics supplies the geometry* (Theorems 7.7–7.9, proved at Born-grade): adiabatic transport of the family projector is Berry transport, and the quantum geometric tensor equips the moduli with the Fubini–Study package — premises shared with the Constitutional Born Theorem, nothing new staked.
- *One clause remains a dictionary entry* (Axiom 7.11): the capacity is a natural characteristic-class pairing — with non-vanishing derived from consistency (Lemma 7.12), the unit pinned three ways (Lemma 7.14), the form robust across all viable candidates (Theorem 7.17), and integrality derived (Theorem 7.13), its single residue named (R-INT), discharged on the benchmark (§H.8), and the clause's two scope refinements closed at class level (T-INF-7, T-INF-8: §H.11).
- *The arithmetic then decides* (Theorem 7.16): $\mathrm{Cap}(N_g) = \pi N_g$ against $S_{\mathrm{tot}}(N_g) = \pi N_g(N_g-1)/2$ admits $N_g \in \{1, 3\}$, saturating exactly at $N_g = 3$; with the CP lower bound, $N_g = 3$ uniquely.

The parity route stands alongside as an independent consistency check (Appendix E, Steps 1–2): $a^+ - a^- = \sigma(\mathsf{CP}^{N_g-1})$, non-zero exactly at odd N_g, with the admissible value passing strictly. What the new architecture removes is any reliance on it — and any reliance on $K = 35$ or $J = 7$, whose absence from the chain is what answers the circularity challenge (§8.2).

18.2 Pointer to the methodological reflection

The technical proof of $N_g = 3$ raises a methodological question that we take seriously. Why has this number been un-closed for nearly a century since the foundation of quantum mechanics? What did we have to assume that prior closed-system frameworks did not? Section 14 addresses these questions directly. The short answer is that closed-system QFT and GR can produce only necessary conditions on integer multiplicities, not upper bounds; upper bounds require open-system dynamics – specifically, the bulk-to-boundary projection, finite instanton-action budgets, and the cascade closure of attractors – which we encode as GDOS. We argue that this is the missing link, and that the present derivation is in equal parts a technical proof of three generations and a vindication of the open-system foundational shift.

18.3 The four integers, and who chose them

The quantitative content of Appendix H is the identity $2 \times 2/4 = 1$, and its scientific content is that nobody chose the numbers. The 4 is the price of a quadratic refinement, fixed in the mathematics literature two decades ago and not negotiable (Hopkins & Singer, 2005). The rule that the 4 must be paid as 2×2 — never 1×4 — is the Wu condition, a parity law of six-dimensional topology. The first 2 is $L - N$: the same subtraction of ledger integers that sets the BI-6 trace ratio, now found doing a third job as the family's toll charge into the one channel the topology leaves open. The second 2 is a mirror: the closed model on which the global theorems are proved is the reflection-double of the physical seam, and a mirror shows every unit of flux twice. Four integers, four different owners — a functor, a parity law, a ledger, a reflection — and one identity. The reader may decide how comfortable they are calling that an accident.

18.4 Why agreement is the right prize

It is worth being precise about why agreement is the result one should want here, rather than a derivation of one side from the other.

A derivation flows downhill: it makes one side primary and the other its consequence, and it is only as strong as its premises. An *agreement* between two computations that share no working parts is a different kind of evidence. The bulk face of the bridge uses no quantum mechanics — no wavefunctions, no operators, no Born rule; it is topology and an anomaly ledger. The boundary face uses no bulk data — it is the adiabatic theorem of §7.5 acting inside quantum mechanics. Neither could see the other's answer. That they meet, and meet exactly, after a factor audit with no adjustable entries, is either a deep property of the architecture or an extraordinary accident; the harnesses of §H.12 were built to close the accidental routes, and the wrong factorizations of the 4 are not merely absent but *forbidden*, by a parity condition this program did not invent for the occasion.

This is also why the grander reading is declined explicitly. The bridge does not found quantum mechanics; its boundary input *is* quantum mechanics, finished upstream by the cohort's own Born theorem (§14). What it founds is smaller and, for this paper, more valuable: the coupling clause — the one sentence of the family-number derivation that had remained at the axiom level — is exhibited, on the benchmark, as the output of a computation. An axiom became an answer.

19 Methodological reflection: why $N_g = 3$ required open-system thinking

This section is a methodological discussion of the proof strategy. It is not a technical addendum to the proof itself: the proof of $N_g = 3$ is complete in Sections 5–8 (granted the named technical hypotheses of Appendix E). The purpose of this section is to articulate the methodological lesson that emerged from our extensive trial of closed-system routes to $N_g = 3$ within GG-6/BI-6 alone. The lesson, in one sentence: closed-system QFT and GR can produce necessary conditions on the family number but cannot produce upper bounds, and this is precisely the structural limitation that has held $N_g = 3$ open for the century since the foundation of quantum mechanics.

The reader interested only in the technical derivation can skip this section without loss of mathematical content.

19.1 The historical pattern of failure

The three-generation puzzle has resisted closed-derivation proof for half a century. Section 1.1 of the present paper surveyed the leading prior approaches – heterotic Calabi–Yau compactifications (Candelas et al., 1985), magnetized toroidal compactifications (Blumenhagen et al., 2007), $\mathbb{Z}_3$ orbifolds (Ibáñez et al., 1987), anomaly cancellation (Dobrescu & Poppitz, 2001), and the Kobayashi–Maskawa CP argument (Kobayashi & Maskawa, 1973) – and noted that each gives only a partial constraint. But the failure pattern is more structural than that survey suggested.

Each of these approaches stays rigorously within the framework of closed-system quantum field theory and general relativity: a manifold, a connection, a partition function, a path integral. Each can produce necessary conditions on the family number N_g – Diophantine constraints from integrality, anomaly-cancellation constraints from gauge-gravitational mixed terms, index-theorem constraints from chiral asymmetry. None can produce a sufficient condition – an upper bound that excludes $N_g \geq 4$.

This is not a technical accident. It is a structural fact about closed-system frameworks, and it explains why the three-generation puzzle has remained unsolved despite half a century of sustained effort by some of the best minds in theoretical physics.

19.2 What closed-system methods can and cannot do

In a closed-system framework, time evolution is unitary. The Hilbert space carries a linear, norm-preserving, reversible representation of the time-translation group. All physical states live with equal weight in a single irreducible structure. There is no canonical notion of which states are "unstable", "decaying", "saddle-like", or "boundary-projected" – those notions require an external reference (a measurement, a coupling to an environment, a boundary condition).

Three structural consequences follow.

Consequence 1: closed-system methods produce only necessary admissibility conditions. Anomaly cancellation says "this combination of fields cannot consistently exist". Integrality says "this number must be an integer". Index theorems say "this chiral asymmetry equals this geometric invariant". All three are bilateral: they can rule out values that fail the constraint, but they cannot rule out values that satisfy the constraint with larger numbers. If $N_g = 3$ satisfies anomaly cancellation, $N_g = 6$ probably does too (with appropriate adjustment of charges). If $N_g = 3$ satisfies integrality of some Chern character, $N_g = 30$ certainly does. There is no closed-system invariant whose value is "N_g but not $N_g + 1$".

Consequence 2: closed-system methods cannot produce upper bounds on integer multiplicities. To get $N_g \leq 3$, one needs some statement of the form "for $N_g \geq 4$, something runs away, decays, fails to stabilize, exceeds a budget". Each of these statements is intrinsically open-system:

- “runs away” requires a sense of stability flow;
- “decays” requires irreversibility;
- “fails to stabilize” requires a fixed-point structure with distinguished basin of attraction;
- “exceeds a budget” requires a finite resource bound.

None of these have closed-system formalizations. In a closed system, all states evolve unitarily forever; nothing decays, nothing runs away, no budget is finite.

Consequence 3: closed-system methods cannot derive integer multiplicities from continuous data. The closed-system path integral on a fixed background is a continuous functional whose stationary points are continuous solutions to PDEs. Integer outputs (chiral asymmetries, instanton numbers, family numbers) come from the topology of the background – but the topology has to be put in by hand. There is no closed-system mechanism that takes BI-6 closure and derives the value $N_g = 3$ rather than assuming it via a chosen background topology.

The historical practice in string compactifications is to scan candidate Calabi–Yau threefolds for those with Euler characteristic $\chi = \pm 6$ (giving three families), or to choose a $\mathbb{Z}_3$ orbifold, or to fix magnetic flux $m = 3$ on each torus factor. In each case, the choice is input, not output.

19.3 The four closed-system attempts within GG-6/BI-6

We attempted four routes to derive $N_g = 3$ within closed-system GG-6/BI-6 alone, before turning to GDOS. Each failed for the structural reasons just articulated. We document them here as a record, both because they were extensive and because their failure pattern is the most direct empirical evidence for the methodological point.

Proposal 1: Chern–Simons multiplicity axiom. We attempted to show that BI-6 closure forces a unique 7-dimensional Chern–Simons level k that, on dimensional reduction to 4D, produces three families. The plan was: (i) identify the Chern–Simons action whose differential is X_4; (ii) impose integrality of the level; (iii) use the level to count boundary chiral zero modes via an index theorem. Steps (i) and (iii) are rigorous and standard. Step (ii) requires an axiom that the bulk Chern–Simons level is exactly the minimum saturating BI-6. This axiom is plausible but not derivable from BI-6 alone in a closed-system formulation; one can write down infinitely many integer levels consistent with BI-6, and the choice of the minimum is an additional principle (it is, indeed, MSMF, but MSMF is part of GG-Theory’s constitutional framework, not a closed-system QFT principle).

Proposal 2: Magnetized torus. We attempted to derive $N_g = 3$ from the index theorem on a 2-torus with magnetic flux $m = 3$ (Blumenhagen et al., 2007). The Atiyah–Singer index theorem indeed gives net chirality m on T^2 with abelian flux m. But m itself is a free integer not forced by BI-6 closure. To force $m = 3$ rather than $m = 4$ or $m = 7$, one needs an extra principle outside BI-6 – typically anomaly cancellation between the magnetic flux and a Wilson-line phase. This anomaly-cancellation argument constrains m modulo something but does not fix m uniquely.

Proposal 3: $\mathbb{Z}_3$ orbifold. We attempted to derive $N_g = 3$ via a $\mathbb{Z}_3$ orbifold of the compact direction (Ibáñez et al., 1987). The three twisted sectors correctly produce three families, but the $\mathbb{Z}_3$ choice is again input rather than output. Why $\mathbb{Z}_3$ rather than $\mathbb{Z}_2$ or $\mathbb{Z}_4$? The answer in our trial was: "because the OGN-5 ledger has $N_g = 3$". But this is circular – we are trying to derive $N_g = 3$, not to assume it.

Proposal 4: Partial GG-6 anomaly saturation. We attempted to show that BI-6 closure, taken together with a specific saturation of the SU(5)-U(1)$_B$-U(1)$_{Qi}$ mixed anomaly polynomial, forces a unique chiral spectrum with $N_g = 3$. This came closer than the others: the saturation conditions on X_4 produce a tight set of Diophantine constraints, and the OGN-5 ledger $(1, 3, 5; 7, 35)$ is the minimum solution. But "minimum solution" is again MSMF, not pure BI-6. Without MSMF, BI-6 admits infinitely many integer ledgers $(M, kN, kL; J, k^2K)$ for any integer $k \geq 1$.

In all four attempts, the final step required an additional principle beyond BI-6 itself. In Proposals 1, 2, and 4, the additional principle was MSMF-like (selecting the minimum integer solution). In Proposal 3, the additional principle was a chosen orbifold group. Each of these additional principles is itself an external input – precisely the "something more than closed-system QFT" that we eventually identified as the open-system framework.

19.4 Why Proposal 4 (within GDOS) succeeded where Proposals 1–3 failed

After concluding that the four closed-system proposals could not close the upper bound on N_g, we turned to open-system thinking. The decisive move was to ask: what is N_g as a structural quantity in the boundary theory, not in the bulk?

The answer is illuminating. N_g is the cascade depth of the boundary attractor structure – specifically, the number of unstable critical points of the GFF Morse functional f_{GFF} on the family-projector moduli space CP^{N_g-1}. This identification has four immediate consequences that make the upper bound natural:

- The family-projector moduli space CP^{N_g-1} exists only because we are projecting from the 6D bulk to the 4D boundary. In a pure closed-system bulk QFT, there is no preferred "family-projector moduli"; families are just labels on the bulk Lagrangian.
- The GFF Morse functional is a free-energy functional that includes both energetic (GEP) and entropic (CEP) contributions. The entropic part requires an open-system formulation: the entropy of a configurational class of projectors, which has no closed-system meaning.
- The cascade of critical points is a sequence of unstable saddles and one stable attractor, with the cascade closure dictated by irreversible gradient flow. The instanton budget is a finite resource available for boundary descent – a concept that has no closed-system analogue.
- The capacity $\mathrm{Cap}(N_g) = \pi N_g$ comes from the quantum geometry of the boundary moduli — Berry curvature and the Pontryagin pairing, activated by the runtime's own adiabatic theorem — with its integral quantization a Hopkins–Singer differential-cohomology statement (Hopkins

& Singer, 2005). Every ingredient requires an open-system / boundary-attractor framework to even be formulated.

In other words, Proposal 4 succeeds because it identifies N_g not as the value of a closed-system invariant but as the cascade depth of a boundary-attractor structure. The latter has an upper bound forced by the BI-6 anomaly inflow budget; the former does not.

This is the structural reason for the success. It is not a clever choice of formalism; it is a recognition that family number is, in its nature, a quantity of the boundary projection rather than of the bulk geometry.

19.5 The biggest missing link since quantum mechanics

The success of Proposal 4 invites a deeper observation. The open-system character of physical observation – the fact that all measurements are irreversible projections of bulk dynamics onto a finite-dimensional boundary – has been understood at a phenomenological level since the founding of quantum mechanics. Yet it has never been elevated to the foundational level. This, we argue, is the missing link that has held the three-generation puzzle (and many related structural questions in foundational physics) open for the century since 1925.

Quantum mechanics was established in 1925–1930 (Heisenberg, Schrödinger, Dirac, von Neumann) as a closed-system framework: a Hilbert space, a unitary evolution operator, a state vector (von Neumann, 1932). The Born rule (Born, 1926) introduced probabilities, but only at the measurement level – the unitary evolution itself is closed and reversible. Quantum field theory (Heisenberg–Pauli, Tomonaga, Schwinger, Feynman, Dyson) extended the closed-system framework to relativistic field operators on a Minkowski background. General relativity (Einstein 1915) is a closed-system theory of the gravitational field on a smooth Lorentzian manifold.

The open-system aspects of physics – measurement collapse, decoherence, irreversibility, dissipation, branching, observer-relative state reduction – were noted from the start (von Neumann's 1932 measurement chain, Heisenberg's "cuts") but were never elevated to the foundational level needed to derive structural integers like $N_g = 3$.

A long sequence of partial advances added pieces of what is needed.

Lindblad / GKLS dynamics (1976). Lindblad (1976) and Gorini, Kossakowski & Sudarshan (1976) independently derived the most general completely positive trace-preserving (CPTP) master equation for open-system quantum dynamics. This gave the framework for irreversible quantum evolution at the level of density matrices, and is the technical backbone of all modern open-system quantum theory.

Decoherence theory (Zeh 1970, Zurek 1981). Zeh (1970) and Zurek (1981) explained how open-system dynamics produces classicality from quantum entanglement: when a system couples to an environment with many degrees of freedom, off-diagonal elements of the reduced density matrix decay rapidly, leaving an effective classical mixture. This is the mechanism by which closed-system quantum mechanics produces apparently classical outcomes – but the mechanism itself is intrinsically open-system.

Holographic principle ('t Hooft 1993, Susskind 1995). 't Hooft (1993) and Susskind (1995) suggested that the fundamental degrees of freedom of a $(d{+}1)$-dimensional bulk theory project onto a d-dimensional boundary, with the bulk effectively encoded by the boundary. This was a profound suggestion that the physical content of a higher-dimensional theory is, in some operational sense, available only on a lower-dimensional boundary – precisely the open-system / projection structure.

AdS/CFT correspondence (Maldacena 1997). Maldacena (1997) gave an explicit example of holographic bulk-to-boundary descent: type-IIB string theory on $\mathrm{AdS}_5 \times S^5$ is equivalent to $\mathcal{N} = 4$ supersymmetric $\mathrm{SU}(N)$ Yang–Mills theory on the four-dimensional boundary. This was the first concrete realization of the holographic principle and provided a precise dictionary between bulk dynamics and boundary observables.

Differential cohomology (Hopkins–Singer 2005). Hopkins & Singer (2005) gave the rigorous formalism for anomaly inflow from a higher-dimensional bulk to a lower-dimensional boundary. Differential cohomology refines ordinary cohomology by incorporating both topological data (cohomology classes) and differential-form data (curvature representatives), giving the precise tool for tracking how anomaly polynomials descend holographically. This is exactly the mathematical machinery used in our Appendix E.

Each of these is a partial advance toward an open-system foundation. None of them has been elevated to the foundational role of an axiom system. They have been treated as tools – powerful tools, indispensable tools – but not as principles. Quantum mechanics in 1925 made the state vector and the unitary evolution operator into foundational principles; nobody has yet made the boundary projector and the cascade of attractors into foundational principles in an analogous way.

GDOS, as developed in Arisaka (2026, A5-GDOS) and the GG-Theory program more broadly, is the explicit assembly of these partial advances into a foundational open-system framework that elevates the bulk-to-boundary projection, the CPTP measurement structure, and the bulk-to-boundary anomaly inflow to axiomatic status. The four GG-Theory axioms (LC, GGEP, OGGEP, MSMF) include OGGEP – the Open-system Geometry-Gauge Equivalence Principle – as one of the four foundational axioms, alongside local causality, gauge-geometry equivalence, and maximum-symmetry-minimum-freedom (Arisaka, 2026, A1-Constitution).

This is the methodological move that distinguishes GG-Theory from the prior generation of closed-system attempts. And this is why GG-Theory can derive $N_g = 3$ where the closed-system frameworks of the past century could only produce partial constraints.

19.6 The four moves of the closed → open transition

Concretely, the closed → open transition involves four moves. Articulating these makes precise what has changed and what has not.

Move 1: Replace unitary $U(t)$ with the boundary projector Π. Instead of treating dynamics as unitary evolution on a Hilbert space, treat it as a CPTP map from a 6D bulk to a 4D boundary. The boundary projector $\Pi : \mathcal{M}_6 \to \mathcal{M}_4$ is the fundamental object; the bulk Hilbert space is a resource from which boundary observables are extracted, not a self-contained universe. Unitary evolution is a special case (when the bulk and boundary coincide and the projector is the identity), but it is not the general case.

Move 2: Replace the state vector with the boundary attractor (GA). Boundary observables are not arbitrary states in some Hilbert space; they are attractors of the GFF Morse functional on the boundary moduli. The discrete catalog of GAs is what we mean by "physical observables" – fermion masses, gauge couplings, vacuum expectation values are all GAs. The continuous Hilbert space is a bulk concept; the discrete attractor catalog is a boundary concept.

Move 3: Replace the closed-system Hamiltonian with the GFF. The closed-system Hamiltonian H that drives unitary evolution is replaced by the Geometric Free Functional $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$, an energetic-plus-entropic free energy on the boundary moduli. Critical points of F_{GFF} are GAs; the cascade structure of saddles and minima is the dynamics. The free-energy structure is intrinsically irreversible: gradient flow descends in F_{GFF} toward attractors and never reverses.

Move 4: Replace closed-system anomaly cancellation with bulk-to-boundary anomaly inflow. The closed-system Bianchi identity $dH_3 = 0$ is replaced by BI-6, $dH_3 = X_4$, where X_4 is the anomaly polynomial that descends to the boundary moduli. The Hopkins–Singer differential cohomology framework (Hopkins & Singer, 2005) makes this descent rigorous. BI-6 closure on $\mathcal{M}_6$ does not say "no anomaly"; it says "the bulk anomaly X_4 is exactly the bulk-to-boundary inflow into the boundary observables".

Each move converts a closed-system structure into an open-system one. Together they constitute the GDOS framework. Note that none of these moves replaces the closed-system structures on the bulk side – the bulk is still a closed-system manifold with unitary $\text{Spin}(1,5)$ frame bundle and full BI-6 closure. What changes is the relationship between bulk and boundary: that relationship is now explicitly open-system, irreversible, and CPTP.

19.7 Why GG-6/BI-6 alone wasn't enough: the structural diagnosis

The implication for the present paper is direct. Sections 5, 6, and 7 establish three legs of the proof of $N_g = 3$. Of these:

- **Leg (L1) (MSMF lower bound)** is closed-system in flavor: it uses unitary-rephasing combinatorics on the CKM matrix. But it is empirically conditioned on $J_{\text{CKM}} \neq 0$, a measurement-theoretic statement. Without the open-system measurement perspective, even (L1) is incomplete. The Born-rule projection that makes J_{CKM} measurable is itself open-system.
- **Leg (L2) (GDOS Morse upper bound)** is intrinsically open-system: it uses the GFF Morse cascade on the family-projector moduli. This is the only leg that can produce an upper bound on N_g, and it can do so only because it uses open-system tools.

- **Leg (L3) (codebook self-consistency)** is constructive – it uses the explicit GG-A9-SM codebook of fermion Yukawas, which is itself derived from open-system / boundary-projector data. The codebook formulas $y_\tau = NJ/2^{J+N_g+1}$, etc. are first-order Yukawa couplings of boundary fermions, with structural integers from the OGN-5 ledger appearing as numerators and exponents.

Each of the three legs uses the open-system / boundary-projector structure. Without GDOS, none of the three legs would close.

This is the methodological diagnosis. The "open-system perspective" is not a stylistic choice or a clever computational shortcut; it is a structural necessity for closing the derivation of $N_g = 3$. Any closed-system framework, no matter how elaborate, will produce only necessary conditions on N_g, never the upper bound.

19.8 Why this is the same diagnosis for many other unsolved structural problems

The diagnosis is not specific to the three-generation puzzle. It applies to several other long-standing structural questions in foundational physics where closed-system methods have produced only partial constraints:

- **The hierarchy problem** (why $E_{\rm EW} \ll E_{\rm Pl}$). Closed-system methods give consistency requirements (technical naturalness, anomaly cancellation) but cannot derive the hierarchy itself. GG-Theory derives the hierarchy from the bulk and the OGN-5 ledger (Arisaka, 2026, A2-GG6), with the key step being the bulk-to-boundary projection $E_{\rm Pl}/2e^{35} \simeq M_{\rm lock} \simeq 3.85$ TeV that locates the IR scale. This is open-system reasoning.

- **The Higgs mass** (m_h^2 as a hierarchy with respect to $E_{\rm GUT}$). Closed-system methods predict quadratic divergences and require fine-tuning. GG-Theory predicts m_h^2 as a boundary-projector exit from the Morse cascade with a definite OGN-5 fraction (Arisaka, 2026, A9-SM), again an open-system computation.

- **Proton stability** (why $\tau_p \gg \tau_{\rm universe}$). Closed-system methods constrain the dimension-6 baryon-violating operators and require the GUT scale to be high. GG-Theory derives the proton-instability suppression from the bulk-to-boundary projection of gauge-baryon $\mathrm{U}(1)_B$ inflow to the boundary (Arisaka, 2026, A3-BI6; Arisaka, 2026, A9-SM), an open-system computation.

- **Dark matter** (the absence of a closed-system mechanism for dark matter that does not require ad hoc particle additions). GG-Theory predicts $\sim \mu$eV axion-like dark matter as a boundary attractor of the bulk $\mathrm{U}(1)_{Qi}$ sector (Arisaka, 2026, C1-Kaxion), again open-system.

In each case, the structural integer or scale that closed-system methods cannot derive is derivable in GG-Theory because of the open-system boundary-projection structure. This is consistent with the methodological diagnosis: closed-system methods cannot derive sufficient conditions on integer multiplicities or hierarchies; open-system methods can.

19.9 The final move: from prove to decide

One methodological move behind this paper deserves its own record, because it is what unlocked the architecture of leg (L2). For most of this program the budget question was posed as *prove*: derive the cap 3π from the axioms. Two adversarial campaigns returned nulls, and the question was then re-posed as *decide*: for each needed statement, either prove it from the axioms, or prove an obstruction showing no such proof exists, and in the latter case name the minimal addition. The decision method produced, in short order, the two obstruction theorems (Appendix G), the recognition that the missing structure was the quantum geometry of the moduli — already forced at Born-grade by the adiabatic theorem — and the collapse of five hypotheses into one integrality statement, itself then derived with a single topological residue (Theorem 7.13). Obstruction results are not failures; they are the honest form of progress on questions that resist proof, because they terminate searches and localize axioms. A program that advertises falsifiability in physics should practice it in mathematics: state what would refute a derivation route, and when it is refuted, say so and move the frontier.

19.10 The bottom line

To our knowledge, $N_g = 3$ is the first integer in physics that has been closed-derivation forced – i.e., proved equal to a specific value rather than constrained to a range – with explicit lower and upper bounds inside a single declared axiom system, with the conditioning stated at every level. The Kobayashi–Maskawa CP argument has been known for over fifty years and gives only $N_g \geq 3$. The heterotic-CY and magnetized-torus arguments have been known for forty years and give only conditional upper bounds (conditional on extra geometric inputs that are themselves not derived). The anthropic argument is heuristic.

GG-Theory closes the derivation. The reason it can close where the closed-system frameworks of the past century could not is that it elevates the open-system / boundary-projector / bulk-to-boundary-anomaly-inflow perspective to foundational status.

The strict GDOS Morse-positivity reduction of Appendix E makes this concrete: the family number is forced to be odd by the mod-4 grading of differential cohomology in 6D – a parity statement that cannot even be formulated without the bulk-to-boundary descent of differential cohomology. Pure closed-system QFT does not have a notion of "differential cohomology of the bulk descended to the boundary". It has only "cohomology of the bulk". The descent is intrinsically open-system.

The methodological achievement is therefore prior to the technical result. The technical result is "$N_g = 3$"; the methodological achievement is "the open-system / boundary-projector / bulk-to-boundary-anomaly-inflow framework can do what closed-system QFT/GR cannot". In the sense that this is what was missing for a century, the present paper is in equal parts a derivation of three generations and a vindication of the open-system foundational shift.

19.11 The two fibers behind the method

The methodological lesson of this section has a geometric root, made exact by the cohort's recent campaign (GG-W1-Symmetry; Arisaka, 2026, P5-CPT; Arisaka, 2026, A5-GDOS): *holonomy requires closed loops, and an interval has none.* The two non-spatial directions of the GG-6 bulk are fibers of opposite topological kind — the compact phase circle S^1_θ carries all of the bulk's holonomy, quantized and protected, its transport a group; the scale interval I_r carries none, its transport composing only forward, a semigroup. Sorted on this anatomy, the two halves of the family-number problem live on different fibers. The *count and its protection* are circle-side: N_g is holonomy-class data — an integer written on the fiber that remembers, exact at every scale, which is why the number neither runs nor drifts. The *decision* — the upper bound — is interval-side: the capacity theorem prices the family ladder against the relaxation budget of a one-way runtime, and a budget is a semigroup concept — it exists only on the fiber that cannot revisit.

This is why the historical record has the shape documented above. Closed-system methods are, in this anatomy, *circle-only* methods: they see the quantized, the protected, and the necessary — which is why the lower bound $N_g \geq 3$ has been available since 1973 — and they are structurally blind to the interval, where the budget lives — which is why no upper bound could even be formulated for fifty years. The open-system shift of this section is, geometrically, the admission of the second fiber into the foundations; and the mod-4 parity statement of Appendix E, formulable only in bulk-to-boundary descent, is the two fibers cooperating in a single theorem — differential-cohomology data of the circle side, descended along the interval. The one-sentence summary of the method: *the family number is written on the fiber that remembers and priced by the fiber that forgets — and a method that reads only one fiber can bound it from only one side.*

19.12 Closing remark: a century-old missing link

The history of theoretical physics is, in part, a history of recognizing that what looked like an arbitrary phenomenological choice was actually forced by a deeper principle that had not yet been articulated. Newton showed that planetary motion was not arbitrary but forced by inverse-square gravity. Maxwell showed that electromagnetic phenomena were not separate but forced by his four equations. Einstein showed that gravity was not a force but a manifestation of curved spacetime. Heisenberg, Schrödinger, and Dirac showed that atomic spectra were not arbitrary but forced by quantum mechanics.

The present paper, in its small way, makes the analogous claim for the family number $N_g = 3$. The fact of three generations was a phenomenological input to the Standard Model for fifty years. The present derivation shows that it was not arbitrary: it was forced by the geometry of the 6D bulk, the integrality of BI-6, and the cascade structure of the bulk-to-boundary projection.

The principle that had not been articulated – the missing link – is the elevation of the bulk-to-boundary projection, the irreversible CPTP structure of measurement, and the cascade of attractors to foundational status. This is the open-system shift. It has been adumbrated since the early days of quantum mechanics, but never made foundational. The present paper is one piece of evidence that, when it is made foundational, the structural integers of the Standard Model can be derived rather than postulated.

We close this section with a summary table.

Table 8: Closed-system vs. open-system foundational stances on the three-generation puzzle. Row by row, the same question answered twice — once under a closed unitary description and once under the open-system projection this cohort uses. The table is here because most of the disagreement about the three-generation problem is a disagreement about which column a reader is standing in, and that is easier to see side by side than to argue in prose.

Aspect	Closed-system framework	Open-system / GDOS framework
Time evolution	Unitary $U(t)$	CPTP Π (boundary projection)
State concept	State vector $\|\psi\rangle$	Attractor catalog (GAs)
Dynamical functional	Hamiltonian H	GFF $F_{\mathrm{GFF}} = F_{\mathrm{GEP}} + F_{\mathrm{CEP}}$
Anomaly principle	Closed Bianchi $dH_3 = 0$	BI-6 inflow $dH_3 = X_4$
Family number N_g	Background topology input	Cascade depth on CP^{N_g-1}
Lower bound on N_g	KM CP $\Rightarrow N_g \geq 3$	Same
Upper bound on N_g	Not available	GFF Morse cascade $\Rightarrow N_g \leq 3$
Quantitative codebook	Tunable	GG-A9-SM codebook, 4 sig. figs.
Status of $N_g = 3$	Phenomenological input	Derived integer

The right-hand column is the GG-Theory perspective. Its operational content is exactly the open-system shift articulated in this section. Its derivational power is the technical proof of the present paper.

20 Part V in plain English: what has been established, and under which fences

The claim is that $N_g = 3$ is forced — not permitted, not preferred, but forced — by three arguments that constrain from independent directions, under conditions this paper has declared rather than assumed quietly.

That last clause is the whole of the honesty in the result. Two of the three legs carry hypotheses that are not derived here: one is imported from a companion at theorem grade, and one is declined outright, with its intended content supplied at a weaker grade and the weakness recorded. A reader who wants to know how much the derivation is worth should read the dissolution ledger before this conclusion, not after it.

What would refute the claim is stated too, and it is not hypothetical: a fourth light family would end the argument, and so would a measurement pushing the CP-violating invariant to zero. Both are the kind of result an experiment could deliver. The Part closes on why this particular integer resisted for fifty years, which is a question about method rather than about physics, and the answer offered is that a closed-system framework could not have reached it.

Part VI
Technical Appendices

The detailed calculations — the CKM phase count, Morse theory on CP^{N_g-1}, the two viable capacity formulations, the numerical charged-lepton comparison, the parity route with the record of the declined hypotheses, the adiabatic theorems of the boundary runtime, and the obstruction theorems with the capacity classification

A Detailed counting of CKM phases

Here we record the standard counting of physical CP phases in V_{CKM} for completeness.

A general $N_g \times N_g$ unitary matrix has N_g^2 real parameters. Of these, $N_g(N_g-1)/2$ are mixing angles, and $N_g(N_g+1)/2$ are phases. The N_g up-type and N_g down-type quarks can be rephased independently except for an overall phase, contributing $2N_g - 1$ unphysical phases. The number of physical CP phases is therefore

$$n_{\mathrm{phase}}(N_g) = \frac{N_g(N_g+1)}{2} - (2N_g - 1) = \frac{(N_g-1)(N_g-2)}{2}. \tag{51}$$

For $N_g = 1$: $n_{\mathrm{phase}} = 0$ (no CP phase, no mixing). For $N_g = 2$: $n_{\mathrm{phase}} = 0$ (one mixing angle, no CP phase). For $N_g = 3$: $n_{\mathrm{phase}} = 1$ (three mixing angles, one CP phase = KM phase). For $N_g = 4$: $n_{\mathrm{phase}} = 3$ (six mixing angles, three CP phases).

The Jarlskog invariant exists if and only if $n_{\mathrm{phase}} \geq 1$, i.e. $N_g \geq 3$.

B Morse theory on CP^{N_g-1} – background

We recall the basic facts about Morse theory on CP^{N_g-1} for completeness; references are Milnor (1963), Helgason (1979), and Schwarz (1993).

B.1 The height function as a perfect Morse function

Let $[z_1 : \ldots : z_N]$ be homogeneous coordinates on CP^{N_g-1}. The function

$$f([z]) = \frac{\sum_a \lambda_a |z_a|^2}{\sum_a |z_a|^2},$$

with $\lambda_1 < \cdots < \lambda_N$, is a smooth real-valued function on CP^{N_g-1}. Its critical points are the N_g homogeneous lines $[0 : \ldots : 0 : 1 : 0 : \ldots : 0]$, with the 1 in position a. The Morse index of the a-th critical point is $2(a-1)$. All critical points are non-degenerate; f is a perfect Morse function, and the Morse polynomial (17) agrees with the Poincaré polynomial of CP^{N_g-1}.

B.2 Gradient flow and the Bogomolnyi bound

The gradient flow of f with respect to the Fubini–Study metric solves

$$\dot{z}_a = -(\lambda_a - \langle\lambda\rangle)\, z_a,$$

which is a one-parameter linear ODE. The flow descends from the index-$2(N_g - 1)$ critical point through index-$2(N_g - 2)$, . . . , down to index 0. Each step satisfies a Bogomolnyi-type bound which is saturated by the gradient trajectory itself, and the saturating action is πk for the descent from index $2k$ to index $2(k-1)$ (Witten, 1982).

B.3 Transversality

Generic gradient flows are transverse (Schwarz, 1993): the unstable manifold $W^u(c_+)$ of any critical point c_+ intersects the stable manifold $W^s(c_-)$ of any other c_- generically, with intersection $W^u(c_+) \cap W^s(c_-)$ of dimension $\mathrm{ind}(c_+) - \mathrm{ind}(c_-)$. This is what makes the cascade chain (20) the dominant gradient flow on CP^{N_g-1}.

C The capacity of the family stratum: the two viable formulations and their completion

This appendix records the two formulations in which the capacity value 3π is quoted across the cohort's cross-reference ledgers, shows that they are exactly the two survivors of the natural classification (Theorem 7.17), and states what completes them. Nothing here is load-bearing — the derivation lives in §7.8 — but the agreement of the two independent formulations is itself evidence of the architecture's robustness, and the arithmetic warning below is worth any reader's attention.

Formulation A (constant form): 3π, read off the seam

The cohort's ledgers (Arisaka, 2026, A3-BI6; Arisaka, 2026, A6-GRIP) quote the budget as a constant, $S^{\max}_{\mathrm{BI6}} = \lfloor J/2 \rfloor\, \pi = 3\pi$ with $J = 7$ the OGN-5 seam integer. Two heuristic routes motivate it. *Route 1 (cell integral):* the integral of X_4 over a fundamental cell, in the canonical OGN-5 normalization, is $2\pi N_g/L = 2\pi \cdot 3/5$; multiplying by the orbifold half-length $L/2 = 5/2$ gives 3π. (The multiplier is $L/2$, *not* the seam factor $J/2 = 7/2$, which would give $21\pi/5$; the two are easily transposed and the distinction is worth restating.) *Route 2 (seam count):* reading $\lfloor J/2 \rfloor$ directly off the seam matches Route 1 numerically. The obvious referee challenge — why $\lfloor J/2 \rfloor$ rather than J, $J - 1$, or $\lfloor J/3 \rfloor$? — is unanswerable within Formulation A alone, and is answered by the classification below.

In the classification of Theorem 7.17, Formulation A is the h^2 (constant-class) capacity: viable, decisive at its MSMF-minimal lattice unit $\lambda = 3\pi$, and agreeing with the extensive form on every physical statement. Its residual blemish is a range of decisive units (3π up to $11\pi/2$) — a discrete knob that MSMF-minimality must close — which is why the extensive form is preferred.

Formulation B (extensive form): πN_g, completed as the capacity theorem

The cell-integral route, read structurally rather than numerically, says the budget scales with N_g: the cell flux of the descended class is proportional to N_g by the Pontryagin scaling $p_1(\mathsf{CP}^{N_g-1}) = N_g h^2$ (Appendix E, Step 2, rigorous), and the conversion is N_g-free — note that L cancels entirely in $(2\pi N_g/L)\cdot(L/2) = \pi N_g$, and that $L = 5$ is in any case available N_g-freely from single-tensor consistency (Arisaka, 2026, A3-BI6, T-BI6-9.0), so no circularity through the trace ratio survives. Formulation B is the p_1 (extensive-class) capacity of Theorem 7.16: its decisive unit is *unique* on the lattice ($\lambda = \pi$; Lemma 7.14), it saturates exactly at $N_g = 3$, and it carries the reachability reading of §11.6 — π of relaxation capacity per family.

What completes them

Three ingredients turn the conjectural pair into the theorem pair: *(i)* the obstruction results (Appendix G), which proved that no derivation of the cap from the prior axioms exists — terminating the search that Routes 1–2 above had been gesturing at; *(ii)* the adiabatic theorems (Appendix F), which supply the descended geometry at Born-grade; and *(iii)* the classification (Lemma 7.14 and Theorem 7.17), which shows that among all natural capacities exactly the two ledger-quoted formulations survive, and that they agree. The referee challenge dissolves: the functional form was never about J at all, and the coincidence $\lfloor 7/2 \rfloor\,\pi = 3\pi$ is a numerological accident of $J = 7$ that the architecture no longer touches.

A self-consistency remark

With the step actions $S_k = \pi k$ of (19), exact saturation at $N_g = 3$ ($S_{\text{tot}}(3) = \text{Cap}(3) = 3\pi$) is now an identity between two independently computed quantities rather than an attractive reading — the strongest self-consistency statement in the paper, and the GG-Theory analogue of criticality. The codebook of Section 7 is consistent with it.

D Numerical comparison of charged-lepton masses

We tabulate the charged-lepton masses for $N_g = 2, 3, 4$ in the GG-A9-SM codebook.

Table 9: Charged-lepton masses as a function of N_g in the GG-A9-SM first-order codebook (before QD completion). The charged-lepton spectrum recomputed at each candidate generation count, before the quark-descent completion is applied. The two ratio columns are the ones to read: they move sharply with N_g and only the $N_g = 3$ row lands near the measured values, which is a consistency check on the count rather than a derivation of it, and is stated at that strength.

N_g	$y_\tau^{(1)}$	$m_\tau^{(1)}$ (GeV)	m_τ/m_μ	m_μ/m_e
2	$14/1024 = 0.01367$	2.381	22.4	312.5
3	$21/2048 = 0.01025$	1.786	16.8	208.33
4	$28/4096 = 0.00684$	1.191	11.2	156.25
Exp.	—	1.77693(9)	16.817	206.768

The agreement of $N_g = 3$ with experiment is at the four-significant-figure level for the mass ratios; $N_g = 2$ and $N_g = 4$ are off by $\sim 33\%$, completely incompatible with observation.

E A constructive reduction of strict GDOS Morse-positivity

This appendix carries the parity route: an independent consistency check on leg (L2), together with the fate of the two technical hypotheses that the natural axiomatization of that route would posit. Steps 1–2 (the signature identity and Hirzebruch's theorem) are fully rigorous and unchanged; they are consumed by the capacity theorem (the Pontryagin scaling) and by the consistency check of §7.11. Step 3 records how the descended class is now supplied — by the quantum geometry of Theorems 7.7–7.9 rather than by the withdrawn hypothesis H_1, whose would-be derivation is closed in the negative by Lemma 7.5 (proof in Appendix G); the no-go computation of §E.5, which first exposed the vanishing, is retained in full. Step 4 states the integrality theorem L-INT, the honest successor of the withdrawn hypothesis H_2, now derived in Appendix G. Step 5 assembles the conditional parity theorem.

E.1 Statement of the result to be proved

Theorem E.1 (The parity route (T-N3-8; conditional consistency theorem))**.** *Grant a non-zero descended inflow class on the family-projector moduli $\mathcal{P}_N = \mathsf{CP}^{N_g-1}$ (supplied at Born-grade by Theorems 7.7–7.9 with the coupling clause, or granted directly by a reader who wishes to test the route in isolation). Then the strict alternating-sum inequality*

$$a^+(N_g) - a^-(N_g) > 0 \tag{52}$$

holds if and only if N_g is odd. In particular the parity route independently excludes even N_g, consistently with the capacity exclusion of Theorem 7.16; and at the admissible value the strict inequality holds unconditionally, $a^+(3) - a^-(3) = \sigma(\mathsf{CP}^2) = 1$.

The remainder of this appendix proves the theorem and records the status of its inputs.

E.2 Step 1: the alternating sum equals the topological signature

This step is fully rigorous. We claim

$$a^+(N_g) - a^-(N_g) = \sigma(\mathsf{CP}^{N_g-1}), \tag{53}$$

where $\sigma(\mathsf{CP}^{N_g-1})$ is the topological signature of CP^{N_g-1}.

Proof of (53). By Lemma 7.2, f_{GFF} on CP^{N_g-1} has N_g non-degenerate critical points with even

Morse indices $\mathrm{ind}_a = 2(a-1)$ for $a = 1, \dots, N_g$. By the definitions of $a^{\pm}(N_g)$ in Postulate 7.19,

$$a^+(N_g) - a^-(N_g) = \#\{a : \mathrm{ind}_a \equiv 0 \pmod 4\} - \#\{a : \mathrm{ind}_a \equiv 2 \pmod 4\} \tag{54}$$

$$= \sum_{a=1}^{N_g} (-1)^{(a-1)} = \sum_{k=0}^{N_g-1} (-1)^k. \tag{55}$$

This sum equals 1 if N_g is odd and 0 if N_g is even.

We compare with the known signature. By the standard topology of complex projective spaces:

- For CP^{N_g-1} with $N_g-1 = 2m$ (i.e. N_g odd, $\dim_{\mathbb{R}} = 4m$), the middle cohomology $H^{2m}(\mathsf{CP}^{N_g-1};\mathbb{R})$ is one-dimensional and the intersection form has signature $+1$. Hence $\sigma(\mathsf{CP}^{N_g-1}) = 1$.
- For CP^{N_g-1} with $N_g - 1 = 2m+1$ (i.e. N_g even, $\dim_{\mathbb{R}} = 4m+2$), the real dimension is not divisible by 4, so the signature is identically zero by definition. Hence $\sigma(\mathsf{CP}^{N_g-1}) = 0$.

The two computations agree, establishing (53). This identification can also be derived from a Morse-theoretic version of the signature for Kähler manifolds equipped with a Kähler-Hamiltonian height function; on CP^{N_g-1} with the Fubini–Study Kähler form, f_{GFF} is exactly such a function. ■

Remark E.2. The reduction (53) is the conceptual heart of this appendix. It rephrases the strict alternating-sum inequality as a topological-invariant condition on the family-projector moduli, namely the non-vanishing of the signature. The remaining steps argue that BI-6 forces this non-vanishing.

E.3 Step 2: Hirzebruch's signature theorem on CP^{N_g-1}

This step is also fully rigorous and entirely classical. We recall the relevant facts.

Theorem E.3 (Hirzebruch signature theorem, Hirzebruch, 1966)**.** *For a closed oriented $4k$-manifold M,*

$$\sigma(M) = \int_M L_k(p_1, p_2, \dots, p_k) \in \mathbb{Z}, \tag{56}$$

where L_k is the k-th Hirzebruch L-polynomial in the Pontryagin classes $p_i \in H^{4i}(M;\mathbb{R})$. The first few L-polynomials are

$$L_1 = \tfrac{1}{3}p_1, \qquad L_2 = \tfrac{1}{45}(7p_2 - p_1^2), \qquad L_3 = \tfrac{1}{945}(62p_3 - 13p_1p_2 + 2p_1^3), \ \dots \tag{57}$$

For CP^{N_g-1} with N_g odd, $N_g = 2m+1$, the manifold is $4m$-dimensional and (56) applies. The total Pontryagin class of CP^{N_g-1} is (Hirzebruch, 1995)

$$p(\mathsf{CP}^{N_g-1}) = (1+h^2)_g^N, \tag{58}$$

where $h \in H^2(\mathsf{CP}^{N_g-1};\mathbb{Z})$ is the hyperplane class normalized so that $\int_{\mathsf{CP}^{N_g-1}} h^{N_g-1} = 1$ (note: $N_g - 1$ is the complex dimension, hence the real dimension is $2(N_g-1)$, and h^{N_g-1} is the volume form). A direct computation (Hirzebruch, 1995) yields

$$\int_{\mathsf{CP}^{N_g-1}} L_m(p_1,\ldots,p_m) = 1, \qquad N_g \text{ odd}, \; N_g = 2m+1. \tag{59}$$

For N_g even, CP^{N_g-1} has real dimension $4m+2$, which is not a multiple of 4; the L-polynomial of total degree $4m+2$ does not exist (Pontryagin classes are concentrated in degrees that are multiples of 4), so the integral on the right-hand side of (56) is identically zero, and accordingly $\sigma(\mathsf{CP}^{N_g-1}) = 0$.

The dichotomy is therefore:

$$\sigma(\mathsf{CP}^{N_g-1}) = \begin{cases} 1 & N_g \text{ odd, i.e. } \dim_{\mathbb{R}} \mathsf{CP}^{N_g-1} \in 4\mathbb{Z}, \\ 0 & N_g \text{ even, i.e. } \dim_{\mathbb{R}} \mathsf{CP}^{N_g-1} \in 4\mathbb{Z}+2. \end{cases} \tag{60}$$

This is a parity constraint on the complex dimension of the family-projector moduli, dictated by the mod-4 grading of the Pontryagin classes (and hence the L-polynomial).

E.4 Step 3: the descended class — how it is supplied, and how it cannot be

The natural hypothesis to posit at this point — call it H_1: the Spin(1, 5) curvature term of X_4 pulls back to $c_R\, p_1(\mathcal{P}_N)$ with $c_R > 0$ — is exactly what this paper declines to posit. In its place stand one theorem and one obstruction, and the computation that forces the replacement.

Lemma E.4 (The descended class under the coupling clause)**.** *Under the coupling clause (Axiom 7.11) with the extensive class, the descended 4-class on the family-projector moduli is*

$$\widehat{X}_4\Big|_{\mathcal{P}_N} = c\, p_1(\mathcal{P}_N) \in H^4(\mathcal{P}_N;\mathbb{R}), \qquad c > 0, \tag{61}$$

where $p_1(\mathcal{P}_N) = N_g h^2$ is the first Pontryagin class of the stratum's own tangent bundle with its Fubini–Study structure — the geometry that Theorem 7.9 shows the boundary runtime activates at Born-grade — and $c > 0$ follows from Lemma 7.12 (non-vanishing) with the orientation fixed by the tautological sign convention of Appendix F.

Lemma E.5 (Gauge descent vanishes on family-projector moduli)**.** *The gauge-curvature contributions to X_4 pull back to zero on $\mathcal{P}_N$: the* SU(5)*,* $\mathrm{U}(1)_B$*, and* $\mathrm{U}(1)_{Qi}$ *bundles are pulled back from the gauge-orbit moduli, transverse to the family-projector moduli; their field strengths are constant along $\mathcal{P}_N$ at the level of cohomology, and their pull-backs as 4-forms vanish.*

What can *not* supply (61) is the bulk connection itself, through any construction whatsoever that uses only the axioms' data: that is Lemma 7.5 (Appendix G), of which the computation in the next subsection is the sharpest special case — the one a reader is most likely to attempt.

E.5 Interlude: why the naive spacetime pullback vanishes (c_R cannot come from the bulk curvature alone)

It is tempting to read H_1 as a routine restriction: compute $\operatorname{tr} R \wedge R$ on the canonical GG-6 bulk and pull it back to $\mathcal{P}_N$. This subsection records the result of carrying that computation out honestly: *the naive route gives exactly zero*, so H_1 cannot mean a literal pullback of the spacetime four-form.

Take the canonical scale-profile metric of Arisaka (2026, A2-GG6),

$$ds_6^2 = e^{2r}\,\eta_{\mu\nu}\,dx^\mu dx^\nu + dr^2 + d\theta^2, \tag{62}$$

with orthonormal coframe $e^a = e^r dx^a$ $(a = 0, \ldots, 3)$, $e^5 = dr$, $e^6 = d\theta$. The spin connection is $\omega^a{}_5 = e^a$ (scale profile $A(r) = r$, $A' = 1$), all other components vanishing, and the curvature two-forms are

$$R^{a5} = -\,e^a \wedge e^5, \qquad R^{ab} = -\,e^a \wedge e^b, \qquad R^{A6} = 0, \tag{63}$$

i.e. the bulk is $\mathrm{AdS}_5 \times S^1$ with all sectional curvatures of the scale profile equal to -1. (We verified (63) symbolically, component by component; note in particular that the four-dimensional sector R^{ab} does *not* vanish — a scale-profile product with non-constant scale profile curves the section.)

The first Pontryagin form is then computed directly:

$$\operatorname{tr}(R \wedge R) = \sum_{A,B} R^A{}_B \wedge R^B{}_A \equiv 0 \quad \text{on } \mathcal{M}_6, \tag{64}$$

identically and pointwise: every contributing term is a wedge of parallel two-forms ($e^a \wedge e^b$ against $e^a \wedge e^b$, or $e^a \wedge e^5$ against $e^a \wedge e^5$), and each such wedge contains a repeated one-form. The vanishing is no accident: $\mathrm{AdS}_5 \times S^1$ is conformally flat, and Pontryagin forms vanish pointwise on conformally flat backgrounds. We have confirmed (64) by independent symbolic computation of all components of the four-form.

Two consequences follow.

(i) Any construction that literally pulls the spacetime four-form $\operatorname{tr} R \wedge R$ back to the family-projector moduli yields $c_R = 0$, contradicting the requirement $c_R > 0$ of Lemma E.4. A computation that read $c_R = 1$ directly off the scale-profile metric (62) would be invalidated by (64): the naive route yields zero, exactly.

(ii) H_1, if true, must therefore be a statement about the *induced* connection on the family-projector bundle over $\mathcal{P}_N$ — the connection inherited by the rank-1 projector family from the $\mathrm{Spin}(1, 5)$ action on the quaternionic family space $\mathbb{C}_g^N$ — whose curvature lives on $\mathcal{P}_N$ itself and is not the pullback of the bulk Riemann tensor. That connection is now constructed: it is the Berry connection of the tautological family, forced at Born-grade by Theorem 7.9, with its Pontryagin form the Fubini–Study class of Lemma E.4. The generalization of the vanishing computed here — from the literal pullback to *every* functorial route — is Lemma 7.5 (Appendix G).

This no-go exposes the two-sided structure of the coupling: the failure mode of the most obvious proof attempt is not a technical inconvenience but the special case of a theorem (Lemma 7.5). The integrality statement itself is the subject of the next subsection and of Appendix G (§G.5), where it is derived; its topological residue R-INT is the declared mandate of Appendix H, met on the benchmark background by its κ-bridge theorem.

E.6 Step 4: integrality of the descended periods (L-INT)

The natural hypothesis to posit here — call it H_2, non-degeneracy of the Hopkins–Singer anomaly-inflow map — is likewise declined. As non-degeneracy it is underivable — Lemma 7.5 shows the functorial inflow is identically zero, and Proposition 7.6 shows no boundary theorem can rescue it — but its salvageable content was never non-degeneracy. It was *integrality*: the statement that the descended periods are quantized in the differential-cohomology sense of Hopkins & Singer (2005). That statement is now the integrality theorem L-INT (Theorem 7.13), derived in Appendix G (§G.5) by decomposition: pairing integrality and the toll lattice rigorously, the unit lattice at physics rigor through the level-quantization argument in its spin-refined form, with well-understood analogues in the M-theory anomaly literature (Freed & Hopkins, 2000; Witten, 1996) and a single named topological residue R-INT — the declared mandate of Appendix H, met there on the benchmark background. Its downstream role is confined to Route 1 of the unit classification (Lemma 7.14); the value $\lambda = \pi$ survives even the failure of R-INT through Routes 2–3.

E.7 Step 5: Synthesis — the parity constraint forces N_g odd

We assemble Steps 1–4 into the proof of Theorem E.1.

Proof of Theorem E.1. By Step 1, $a^+(N_g) - a^-(N_g) = \sigma(\mathsf{CP}^{N_g-1})$. By Step 2, $\sigma(\mathsf{CP}^{N_g-1}) = 1$ if N_g odd and $\sigma(\mathsf{CP}^{N_g-1}) = 0$ if N_g even. So strict positivity $a^+(N_g) - a^-(N_g) > 0$ holds if and only if N_g is odd.

It remains to connect parity to the granted descended class. By Step 3, the class is $c\, p_1(\mathcal{P}_N)$ with $c > 0$; a non-zero *integrated* inflow — a characteristic number built from a degree-4 class — exists on $\mathcal{P}_N$ only when $\dim_{\mathbb{R}} \mathcal{P}_N = 2(N_g - 1)$ is divisible by 4, i.e. when N_g is odd; for N_g even the pairing is degree-obstructed and the integrated inflow vanishes identically. Hence a non-vanishing integrated inflow forces N_g odd, and on that parity sector the strict inequality holds, $a^+(N_g) - a^-(N_g) = \sigma(\mathsf{CP}^{N_g-1}) = 1 > 0$.

At the admissible value the conditional scaffolding drops away entirely: $a^+(3) - a^-(3) = \sigma(\mathsf{CP}^2) = 1 > 0$ is an unconditional identity. ■

E.8 The spin reading of the parity route

The parity constraint of Theorem E.1 has a second face, worth recording because it ties the two exclusion legs together at the level of the stratum's geometry rather than at the level of the answer. The complex projective space CP^n is spin precisely when n is odd — $w_2(\mathsf{CP}^{N_g-1}) = N_g\, h \bmod 2$, so CP^{N_g-1} is spin precisely when N_g is *even*. The strata the parity route excludes are therefore exactly the spin strata, and the surviving odd-N_g strata are exactly the non-spin ones. The coincidence is load-bearing at the admissible value: if CP^2 were spin, Rokhlin's theorem (Rokhlin, 1952) — the signature of a closed spin 4-manifold is divisible by 16 — would force $\langle p_1, [\mathsf{CP}^2] \rangle = 3\sigma \in 48\mathbb{Z}$, and the saturation pairing $3 \notin 48\mathbb{Z}$: exact saturation $3\pi = 3\pi$ would be arithmetically impossible on a spin stratum. The parity route, in other words, was detecting the stratum's spin obstruction all along — the same non-spin-ness that permits the capacity theorem's exact saturation. (No tension with the *bulk* spin structure consumed by Theorem 7.13: bulk and stratum are different manifolds, and a spin bulk with a non-spin even-codimension stratum is precisely the setting of the shifted flux quantization of Witten, 1996.) This remark upgrades nothing — the parity route remains a consistency check — but it explains *why* the route works, and it welds it to the capacity route through the stratum's spin geometry.

E.9 What this appendix now achieves

Three things, honestly separated. *First*, the parity route is preserved as an independent consistency check: a reader who grants only a non-zero descended class — from any source — recovers even-N_g exclusion by pure topology (signature plus degree counting), in agreement with the capacity exclusion of Theorem 7.16. *Second*, the fate of the two would-be hypotheses is recorded where a reader will look for them: H_1 declined in favor of a proved obstruction plus a Born-grade theorem; H_2 reduced to integrality, which Appendix G then derives (Theorem 7.13). *Third*, the appendix's rigorous core — the signature identity of Step 1 and the Pontryagin scaling of Step 2 — is consumed upstream by the capacity theorem itself: $p_1(\mathsf{CP}^{N_g-1}) = N_g h^2$ is precisely the extensivity that gives $\mathrm{Cap}(N_g) = \pi N_g$. The cleanest reading:

> The stratum's topology pays linearly ($p_1 = N_g h^2$: one unit of charge per family) while the cascade charges quadratically (the triangle-number toll). The books balance only at $N_g \leq 3$, exactly at $N_g = 3$ — and the parity pattern of the signature, positivity exactly at odd N_g, independently ratifies the survivor.

E.10 An alternative formulation: the Witten index

It is illuminating to reformulate the Morse-positivity statement in terms of the Witten index of supersymmetric quantum mechanics on CP^{N_g-1}. Recall that for a Morse function f on a closed manifold M, Witten's supersymmetric quantum mechanics has a deformed exterior derivative $d_t = e^{-tf}\, d\, e^{tf}$, and the Witten index (the alternating sum of critical-point contributions) coincides with the Euler characteristic $\chi(M)$ in the large-t limit (Witten, 1982).

For CP^{N_g-1} with the Fubini–Study Kähler structure, f_{GFF} is Kähler-Hamiltonian, and the alternating sum $a^+(N_g) - a^-(N_g) = \sum_a (-1)^{(a-1)}$ is not the Euler characteristic (which would be $\sum_a 1 = N_g$). It is instead the signature, which arises in supersymmetric quantum mechanics with the additional grading $(-1)^F$ replaced by $(-1)^{F/2}$ (working with the holomorphic / antiholomorphic grading of the Kähler form). This is the "Hirzebruch L-genus index" rather than the "Euler index".

The strict GDOS Morse-positivity postulate is therefore equivalent to: the Witten index of the Hirzebruch L-genus on the family-projector moduli is strictly positive. And this index is non-zero exactly when the moduli supports a non-trivial L-genus – i.e., when the parity constraint of Step 2 is satisfied.

This Witten-index perspective is consistent with the integrality theorem L-INT: integral quantization of the descended periods is the differential-cohomology form of "the Witten index of the descent is a well-defined integer", and its non-vanishing on the odd-parity sector is what Step 5 uses.

E.11 A concrete check at $N_g = 3$

For $N_g = 3$, the family-projector moduli is $\mathcal{P}_3 = \mathsf{CP}^2$, real dimension 4. We verify the chain of reductions concretely.

Step 1 check. The Morse polynomial of CP^2 is $1 + t^2 + t^4$. Indices are 0, 2, 4. Hence $a^+(3) = 2$ (indices 0, 4) and $a^-(3) = 1$ (index 2). $a^+(3) - a^-(3) = 1 > 0$. ✓

Step 2 check. $\dim_{\mathbb{R}} \mathsf{CP}^2 = 4$, so the L-polynomial is $L_1 = p_1/3$. The first Pontryagin class is $p_1(\mathsf{CP}^2) = 3h^2$ where h is the hyperplane class with $\int_{\mathsf{CP}^2} h^2 = 1$. Hence

$$\sigma(\mathsf{CP}^2) = \int_{\mathsf{CP}^2} \frac{p_1}{3} = \int_{\mathsf{CP}^2} h^2 = 1. \checkmark \tag{65}$$

Step 3 check. $\widehat{X}_4|_{\mathsf{CP}^2} = c\, p_1(\mathsf{CP}^2) = 3c\, h^2$, with $c > 0$ by Lemma E.4 (via Lemma 7.12). Non-zero. ✓

Step 4 check. $\int_{\mathsf{CP}^2} \widehat{X}_4 = 3c \int_{\mathsf{CP}^2} h^2 = 3c > 0$. Non-zero. ✓

Synthesis. $\sigma(\mathsf{CP}^2) = 1 > 0$, hence $a^+(3) - a^-(3) > 0$. Strict GDOS Morse-positivity holds at $N_g = 3$. ✓

Comparison: $N_g = 2$ check. $\dim_{\mathbb{R}} \mathsf{CP}^1 = 2$. No L-polynomial of degree 2 exists; $\sigma(\mathsf{CP}^1) = 0$; $a^+(2) - a^-(2) = 0$, strict positivity fails. ✓

Comparison: $N_g = 4$ check. $\dim_{\mathbb{R}} \mathsf{CP}^3 = 6$. No L-polynomial of degree 6 exists; $\sigma(\mathsf{CP}^3) = 0$; $a^+(4) - a^-(4) = 0$, strict positivity fails. ✓

This concrete check confirms the parity-driven structure of the proof: strict positivity holds exactly at odd N_g, with $N_g = 3$ the unique admissible value once cascade closure and MSMF are imposed.

F The adiabatic theorems of the boundary runtime

This appendix proves Theorems 7.7–7.9 and fixes the conventions their use requires. Everything here is standard mathematics applied to structures the cohort already owns; the point of the appendix is to make that application airtight and to locate its premises exactly.

F.1 Inventory: what the runtime supplies

The adiabatic theorem for a driven system needs three inputs. The boundary runtime supplies all three from named cohort results:

Table 10: The three inputs the adiabatic theorem needs for a driven open system, the named cohort result that supplies each, and the grade at which it arrives. Two are theorems imported from the boundary-law paper and the third is trivial for this geometry. The table is here to make the import surface explicit: nothing in the adiabatic argument is assumed locally, and any reader who doubts a premise can follow the middle column to the place it is proved.

input	source	status
a bona fide dynamical semigroup	GG-A5-GDOS T-GDOS-3 (CPTP via Stinespring), T-GDOS-4 (GKLS)	theorem; premise (P1)
a relaxation gap g	the collision-level contraction: relative entropy contracts with ratio $(1-\mu)$ per exchange, $g \sim -\ln(1-\mu)$ (Arisaka, 2026, A8-QM)	proven at collision level
smoothness of the driven family	Fubini–Study structure of CP^{N_g-1}	trivial

Premise (P1) — the GG-A5-GDOS forward semigroup — and the rate bound $\mu > 0$ are precisely the premises of the Constitutional Born Theorem (Arisaka, 2026, A8-QM). Nothing else is consumed (Remark 7.10).

The collision picture of the gap, in one paragraph. The contraction statement of Arisaka (2026, A8-QM) is operational: each exchange with a fresh boundary port contracts the relative entropy to the basin state by the fixed ratio $(1-\mu)$, μ the guidance-mixing weight. Run k exchanges and the off-attractor content decays like $(1-\mu)^k$ — which is exactly the statement that the instantaneous generator has its non-zero spectrum uniformly damped at rate $g \sim -\ln(1-\mu)$ per collision time. The adiabatic theorem needs nothing finer: it consumes the gap as a rate, not the spectrum as a list. This is why Theorem 7.8 sits at exactly Born-grade — the same collision estimate that relaxes measurement statistics to the Born measure is the estimate that snaps the family projector onto its instantaneous attractor.

F.2 Proof of Theorem 7.7 (Kato transport)

The proof in the main text is complete; we add the two remarks that make it useful. *(i) Uniqueness:* W is the unique transport whose generator has no "in-band" component, $P\,(W'W^\dagger)\,P = P[P',P]P = 0$; any other choice differs by a gauge transformation within $\operatorname{ran} P$. *(ii) No dynamics:* the construction uses only the projector family — it is differential geometry, not physics; the physics enters in Theorem 7.8, which says the runtime's actual evolution follows this geometric transport.

F.3 Proof of Theorem 7.8 (the runtime is adiabatic)

We give the full estimate in four steps, at the level of rigor the paper promises (the operator-theoretic sharpening is engineering item E-1 of Section 12); the structure is Kato's (Kato, 1950), in the extension to gapped generators of contraction semigroups (Avron et al., 2012), which is the GDOS setting by T-GDOS-4.

Step 1 (setup and spectral splitting). In slow time $s = \varepsilon t$ the driven evolution is

$$\varepsilon\,\frac{d\rho}{ds} = \mathcal{L}(s)\,\rho, \qquad \mathcal{L}(s) := \mathcal{L}(x(s)), \tag{66}$$

with $\mathcal{P}(s)$ the spectral projector of $\mathcal{L}(s)$ at the isolated eigenvalue 0 (the instantaneous attractor manifold; for the family stratum its range is spanned by the basin state, whose spatial label is the rank-one projector $P(s)$ of Theorem 7.7). The gap hypothesis says the remainder of the spectrum lies in $\{\operatorname{Re} z \le -g\}$ with $g \sim -\ln(1-\mu)$ per collision time, $\mu > 0$ the contraction weight.

Step 2 (the ansatz and the two projected equations). Write $\rho = \mathcal{P}\rho + \mathcal{Q}\rho$ with $\mathcal{Q} = \mathbf{1} - \mathcal{P}$, and insert into (66). Using $\mathcal{L}\mathcal{P} = 0$ and differentiating $\mathcal{P}^2 = \mathcal{P}$ exactly as in Theorem 7.7:

$$\frac{d}{ds}(\mathcal{P}\rho) = [\mathcal{P}', \mathcal{P}]\,(\mathcal{P}\rho) + \mathcal{P}'\,(\mathcal{Q}\rho), \tag{67}$$

$$\varepsilon\,\frac{d}{ds}(\mathcal{Q}\rho) = \mathcal{L}(s)\,(\mathcal{Q}\rho) + \varepsilon\,[\text{drive terms}], \tag{68}$$

where the drive terms are bounded by $\sup\|\mathcal{P}'\| \cdot \|\rho\|$. Equation (67) is the exact statement that, *up to the leakage term* $\mathcal{P}'(\mathcal{Q}\rho)$, the kernel component is transported by the Kato generator.

Step 3 (the leakage is $O(\varepsilon/g)$). On the range of $\mathcal{Q}$ the generator is invertible with $\|\mathcal{L}^{-1}\mathcal{Q}\| \le 1/g$ (the gap). Solving (68) adiabatically — formally, $\mathcal{Q}\rho = -\varepsilon\,\mathcal{L}^{-1}\mathcal{Q}\big[\frac{d}{ds}(\mathcal{P}\rho)\big] + O(\varepsilon^2)$, made honest by one integration by parts in s against $e^{\mathcal{L}(s)\,\Delta s/\varepsilon}$, whose range part decays like $e^{-g\Delta s/\varepsilon}$ — gives

$$\|\mathcal{Q}\rho(s)\| \;\le\; \frac{C_1\,\varepsilon}{g} \;+\; e^{-gs/\varepsilon}\,\|\mathcal{Q}\rho(0)\|, \qquad C_1 \le \sup\|\mathcal{P}'\| + \sup\|\mathcal{P}''\|. \tag{69}$$

The memory of any initial off-attractor component dies at the relaxation rate; what survives is a standing leakage of order ε/g.

Step 4 (assembly). Substituting (69) into (67) and integrating over the traverse, the kernel component differs from pure Kato transport by at most $C\,\varepsilon/g$ with C controlled by the first two derivatives of the projector family — finite on the compact Fubini–Study geometry of CP^{N_g-1}. Hence $\|\rho_\varepsilon(s) - W(s)\rho(0)W(s)^\dagger\| \le C\varepsilon/g$, as claimed; and within the attractor bundle the transport *is* $W(s)$, whose identification with Berry transport is Theorem 7.9. □

Physics intuition. The estimate formalizes a familiar picture: a fast, gapped relaxation continually snaps the state back onto the instantaneous attractor, while the slow drive drags the attractor itself through the moduli. The state hugs the moving target to accuracy (drive rate)/(snap rate) $= \varepsilon/g$; the direction of the hug — the connection — is fixed by geometry alone, which is why the same $[P', P]$ appears whether the fast dynamics is Hamiltonian (Kato's case) or dissipative (the runtime's case).

F.4 Proof of Theorem 7.9 (the transport is Berry; the geometry is Fubini–Study)

(a) The Berry condition. Along Kato transport $|\chi(s)\rangle = W(s)|\psi_0\rangle$,

$$\langle\chi|\chi'\rangle = \langle\chi|[P', P]|\chi\rangle = \langle\chi|P'P|\chi\rangle - \langle\chi|PP'|\chi\rangle = \langle\chi|P'|\chi\rangle - \langle\chi|P'|\chi\rangle = 0, \tag{70}$$

using $P|\chi\rangle = |\chi\rangle$. The vanishing of $\langle\psi|d\psi\rangle$ along transport is precisely parallelism for the Berry connection of the tautological line bundle (Berry, 1984).

(b) Holonomy. For a closed loop, the accumulated phase is the Berry phase; on any projective line (a Bloch sphere inside CP^{N_g-1}), a circuit at colatitude θ returns the phase $-\pi(1-\cos\theta)$, i.e. $-\frac{1}{2}\times$ the enclosed solid angle. (This case is worked as a pedagogical example in Arisaka (2026,

A8-QM); here it is a theorem about the runtime's own family transport.)

(c) The quantum geometric tensor, computed explicitly. For normalized sections $|\psi(x)\rangle$ over real coordinates x^μ on CP^{N_g-1}, define

$$T_{\mu\nu} = \langle \partial_\mu \psi | (\mathbf{1} - P) | \partial_\nu \psi \rangle. \tag{71}$$

This is the Provost–Vallee construction (Provost & Vallee, 1980). We carry the computation out in full on a projective line (the general case is identical algebra index by index, and is confirmed numerically below). In the affine chart $|\psi(z)\rangle = (1, z)/\sqrt{1+|z|^2}$ of CP^1:

$$|\partial_z \psi\rangle = \frac{(0,1)}{(1+|z|^2)^{1/2}} - \frac{\bar z}{2}\,\frac{(1,z)}{(1+|z|^2)^{3/2}}, \qquad \langle \psi | \partial_z \psi \rangle = \frac{\bar z}{2(1+|z|^2)}, \tag{72}$$

so, subtracting $P|\partial_z\psi\rangle = |\psi\rangle\langle\psi|\partial_z\psi\rangle$, the projected derivative collapses to

$$(\mathbf{1} - P)|\partial_z \psi\rangle = \frac{(0,1)(1+|z|^2) - \bar z\,(1,z)}{(1+|z|^2)^{3/2}} = \frac{(-\bar z,\, 1)}{(1+|z|^2)^{3/2}}, \tag{73}$$

whence

$$\left\|(\mathbf{1} - P)\partial_z \psi\right\|^2 = \frac{1+|z|^2}{(1+|z|^2)^3} = \frac{1}{(1+|z|^2)^2}. \tag{74}$$

Hence $T_{z\bar z} = (1+|z|^2)^{-2}$: exactly the Fubini–Study metric $g_{z\bar z}$ of the area-π normalization, with constant of proportionality 1. Realifying, $\mathrm{Re}\, T$ is the FS metric and $F_{\mu\nu} = -2\,\mathrm{Im}\, T_{\mu\nu}$ is the Berry curvature two-form; on the line,

$$F_{xy} = -\frac{2}{(1+x^2+y^2)^2}, \qquad \int_{\mathbb{R}^2} F_{xy}\, dx\, dy = -2\pi = 2\pi\, c_1\big(\mathcal{O}(-1)\big), \tag{75}$$

while $\int \omega_{FS} = \pi$: the factor-2-and-sign bookkeeping of the conventions subsection, derived rather than quoted. For CP^{N_g-1} with $N_g > 2$ the same algebra gives $T_{i\bar j} = \delta_{ij}/(1+|z|^2) - \bar z_i z_j/(1+|z|^2)^2$, the FS metric in every chart; this identity was additionally confirmed numerically at randomly chosen points of CP^2 (proportionality constant 1.000000, deviations at the 10^{-11} level of the finite-difference scheme). One tensor, both geometries: the metric on state space is the fidelity metric, and its antisymmetric partner is the Berry curvature. □

Worked holonomy: the Bloch circuit, end to end. Take the family $|\psi(\theta,\phi)\rangle = (\cos\frac{\theta}{2}, \sin\frac{\theta}{2} e^{i\phi})$ on a projective line and drag ϕ once around at fixed colatitude θ. The Berry connection along the circuit is $A_\phi = \langle\psi|i\partial_\phi\psi\rangle = -\sin^2(\theta/2)$, so the holonomy phase is

$$\gamma = \oint A_\phi\, d\phi = -2\pi \sin^2\tfrac{\theta}{2} = -\pi\,(1-\cos\theta) = -\tfrac{1}{2}\,\Omega(\theta), \tag{76}$$

half the enclosed solid angle, with the tautological sign. Integrating the Kato equation (23) numerically around the same circuits reproduces (76) at $\theta = 60^\circ, 90^\circ, 120^\circ$ to 5×10^{-8}, with unitarity and projector transport maintained at the 10^{-10} level — the transport of Theorem 7.7, the connection of Berry (1984), and the geometry of Provost & Vallee (1980) agreeing to numerical precision, as they must, being three descriptions of one structure.

F.5 Conventions: the unit, the sign, and one factor of two

Three conventions must be pinned once, here, because a mismatch would propagate a wrong factor into the capacity.

- *The unit.* The step-action lemma (19) fixes the normalization in which a projective line has Fubini–Study area $\int_{\text{line}} \omega_{FS} = \pi$; the toll ledger and the capacity are both expressed in this unit, so no conversion enters the closure condition (26).
- *The sign.* The tautological bundle has $c_1 = -1$ on the line: $\int_{\text{line}} F_{\text{Berry}} = -2\pi$. Physical capacities are built from |periods|; the orientation is fixed once by declaring the descent to pair positively against the cascade's own orientation (Lemma E.4).
- *The factor of two.* $|F_{\text{Berry}}| = 2\,\omega_{FS}$ pointwise: the curvature of the line bundle is twice the Kaehler form in the area-π normalization (equivalently, the Chern integral is 2π against the area π). Every statement in this paper pairs classes in the ω_{FS} normalization of (19); the factor is stated so that no reader re-derives the capacity in the Chern normalization and reports a discrepancy of two.

F.6 Quantitative adiabaticity, and why the payload is immune

The adiabatic parameter of the cascade is $\varepsilon \sim \Gamma_{\text{inst}}/g$ with $\Gamma_{\text{inst}} \sim e^{-S}$: at the smallest toll $S = \pi$ this is $e^{-\pi} \approx 0.043$ — adiabaticity to about four percent, not exponential perfection. This costs the derivation nothing, and the reason deserves stating plainly: everything leg (L2) consumes downstream of the transport is *topological* — Chern and Pontryagin pairings, quantized periods, an admissibility comparison of two integer multiples of π — and integers do not move under four-percent deformations of the transport. The four percent is a statement about how closely a state hugs the instantaneous attractor, not an error bar on N_g.

Honestly noted, likewise: the adiabatic theorems govern transport *between* cascade events. The instanton events themselves are not adiabatic; their geometric factors are part of the toll ledger (19) that the cascade prices anyway. A unified event-level treatment is engineering item E-2 of Section 12; nothing downstream depends on it.

Physics intuition. The content of this appendix, compressed: a gapped quantum system dragged slowly through a family of states cannot help transporting them the Berry way — that is what the adiabatic theorem *is* — and a space of pure states cannot help carrying the Fubini–Study geometry — that is what the quantum geometric tensor *is*. The boundary runtime is a gapped quantum system whose slow variables are family projectors. The geometry of leg (L2) therefore rides in with quantum mechanics itself; the only creative act left to the theory is the coupling clause; the integrality statement itself is derived in Appendix G (§G.5), with one topological residue left to check.

G Rigidity, boundary relativity, and the capacity classification

This appendix proves the two obstruction results of §7.4 and the classification results of §7.8.

G.1 Proof of Lemma 7.5 (rigidity of functorial descent)

Step 1 (invariant-form uniqueness), in full. Let $\mathfrak{g}$ be a simple Lie algebra over $\mathbb{C}$ and $B(X,Y)$ any $\mathfrak{g}$-invariant symmetric bilinear form, i.e. $B([Z,X],Y)+B(X,[Z,Y])=0$ for all Z. *(i)* B defines a linear map $\beta:\mathfrak{g}\to\mathfrak{g}^*$ by $\beta(X)=B(X,\cdot)$, and invariance says precisely that β intertwines the adjoint representation with the coadjoint one. *(ii)* The Killing form κ is non-degenerate (Cartan's criterion, since $\mathfrak{g}$ is semisimple), so κ furnishes an isomorphism $\mathfrak{g}^*\cong\mathfrak{g}$ of representations; composing, $\kappa^{-1}\beta:\mathfrak{g}\to\mathfrak{g}$ intertwines the adjoint representation with itself. *(iii)* The adjoint representation of a *simple* algebra is irreducible (an invariant subspace is an ideal, and a simple algebra has none but 0 and itself), so by Schur's lemma $\kappa^{-1}\beta=c\cdot\mathrm{id}$, i.e. $B=c\,\kappa$. *(iv)* For any finite-dimensional representation ρ, the trace form $B_\rho(X,Y)=\mathrm{tr}(\rho(X)\rho(Y))$ is invariant (cyclicity of the trace), hence $B_\rho=c_\rho\,\kappa$; the constant c_ρ is by definition the Dynkin index. *(v)* The argument transfers to the real form $\mathfrak{so}(1,5)$ because invariant forms complexify and restrict: $\mathfrak{so}(1,5)\otimes\mathbb{C}=\mathfrak{so}(6,\mathbb{C})$ is simple, so the complexified form is a Killing multiple, and so therefore is its restriction. *(Numerical witness:* for the compact form $\mathfrak{so}(6)$, vector versus adjoint representations, the ratio $\mathrm{tr}_{\mathrm{adj}}(XY)/\,\mathrm{tr}_{\mathrm{vec}}(XY)$ computed over random pairs is constant to 10^{-8} and equals $n-2=4$, the textbook index ratio.)

Step 2 (Chern–Weil functoriality; Kobayashi & Nomizu, 1969). If the connection on $E=\mathcal{F}\times_\rho V$ is induced from a principal connection with curvature F, the bundle curvature is $\rho_*(F)$ and $\mathrm{tr}(\rho_*(F)\wedge\rho_*(F))=c_\rho\,\mathrm{tr}(F\wedge F)$ by Step 1 (the trace form is evaluated on the curvature's Lie-algebra values; $\mathrm{tr}\,F=0$ on the semisimple factor kills the would-be $(\mathrm{tr}\,F)^2$ term). Pullbacks along axiom-given maps commute with Chern–Weil.

Step 3 (the bulk form vanishes). On the canonical bulk the spin curvature satisfies $\mathrm{tr}(R\wedge R)\equiv 0$ *pointwise* (§E.5: every contributing term wedges a two-form against itself or its transpose partner, and $\mathrm{AdS}_5\times S^1$ is conformally flat, where Pontryagin forms vanish identically).

Combining: every functorially induced Chern–Weil 4-form on the stratum is $c_\rho\cdot 0=0$; the gauge sector vanishes separately by transversality (Lemma E.5). The scope is exactly the word *functorial*: a construction using data beyond the axioms — the tautological embedding of the moduli in state space, with its Berry structure — is not excluded, and is precisely what Theorem 7.9 supplies. □

Physics intuition. The Dynkin index is the only dial representation theory owns for a quadratic invariant, and it is multiplicative against a field that is zero. No choice of matter content, however exotic, extracts a nonzero anomaly form from a conformally flat background — just as no choice of test charge extracts a force from a vanishing field.

G.2 Proof of Proposition 7.6 (boundary relativity)

The proof is a two-model construction: we exhibit, explicitly, two consistent completions of the boundary axioms that differ in the coupling, and observe that every boundary-quantified statement takes the same truth value in both — from which no boundary theorem can separate them.

The two models. Fix the boundary data once: the Hilbert-space structure, the projection π of Arisaka (2026, A8-QM), the CPTP channel of GG-A5-GDOS T-GDOS-3, its GKLS generator (T-GDOS-4), the collision-level contraction rate, and the family moduli CP^{N_g-1} with the Berry/Fubini–Study geometry of Theorem 7.9 — all of which are computed from the state manifold and the boundary dynamics alone. Now form two joint actions

$$S_{\text{joint}}^{(\kappa)} = S_{\text{bulk}} + S_{\text{boundary}} + \kappa \int \check{X}_4 \smile \check{j}, \qquad \kappa \in \{\kappa_1, \kappa_2\},\ \kappa_1 \neq \kappa_2, \tag{77}$$

where the last term is the topological inflow pairing (differential- cohomology cup product of the anomaly class against the boundary current). A topological term has identically vanishing local variation: it contributes no term to the boundary equations of motion, no term to the GKLS generator of GG-A5-GDOS T-GDOS-4, no shift of the collision-level contraction rate, and no modification of the projector family or its Berry geometry (Theorem 7.9 is computed from the state manifold alone). Hence the two extensions induce *identical* boundary runtimes: every boundary-quantified statement takes the same truth value in both. A theorem fixing κ would distinguish them; therefore no boundary theorem fixes κ. What κ does change is the joint bookkeeping — which cascades can be paid for — and that is exactly where the capacity architecture consumes it. □

Two consequences are worth separating. *First*, this is why the would-be hypotheses H_1/H_2 resist proof: each is the one-sided shadow of a two-sided fact (the bulk cannot send the class; the boundary cannot price it; the coupling lives only in the pair). *Second*, this is not a license for arbitrariness: the coupling clause (Axiom 7.11) is pinned on every axis the pair structure leaves free — non-vanishing by consistency (Lemma 7.12), unit by the three routes of Lemma 7.14, form by Theorem 7.17 — leaving integrality to the derivation of §G.5 below.

G.3 Proof of Lemma 7.14 (the unit classification)

Admissibility reads $\pi N_g(N_g-1)/2 \leq \lambda N_g$, i.e. $N_g \leq 1 + 2\lambda/\pi$. On the derived lattice $\lambda \in \pi\mathbb{Z}$ (Theorem 7.13) — and identically on the conservative fence $(\pi/2)\mathbb{Z}$, whose extra points are all inconsistent or indecisive, as the tabulation shows: at $\lambda = \pi/2$ the admissible set (intersected with $N_g \geq 3$ from leg L1) is empty; at $\lambda = \pi$ it is $\{3\}$; at $\lambda = 3\pi/2$ it is $\{3,4\}$; and it grows monotonically thereafter. Consistency (non-emptiness) excludes $\lambda = \pi/2$; decisiveness — no residual discrete choice of N_g, which would be a hidden knob in the sense forbidden by constitutional axiom AX-4 (MSMF) — excludes $\lambda \geq 3\pi/2$. Hence $\lambda = \pi$ uniquely (Route 1). Without the lattice, consistency alone requires $3\pi \leq 3\lambda$, i.e. $\lambda \geq \pi$, and MSMF-minimality selects the infimum, attained at $\lambda = \pi$ (Route 2). Route 3 is the no-new-scales reading: the capacity quantum coincides with the smallest toll of the ledger (19), so the axiom system introduces no second action scale. □

Why this is not a scandal. Relative invariants are the normal condition of bulk–boundary physics: the level of a Chern–Simons theory is invisible to the conformal boundary algebra until one asks a paired question (a quantized Hall current, an edge central charge). The capacity is the paired question here, and the classification results below are exactly the discipline of asking it without freedom.

G.4 Proof of Theorem 7.17 (form robustness)

The natural degree-4 characteristic classes of the stratum's canonical bundles are generated by h^2 (the tautological class squared), c_1^2 and c_2 of the tangent bundle, and $p_1 = c_1^2 - 2c_2$; on CP^{N_g-1} these pair against the generator $[\mathsf{CP}^2]$ of H_4 as 1, N_g^2, $N_g(N_g-1)/2$, and N_g respectively. For each scaling $\chi(N_g)$, the minimal lattice unit consistent with $N_g = 3$ is $\lambda_{\min} = \lceil 3\pi/\chi(3)\rceil_{(\pi/2)\mathbb{Z}}$; the admissible set at $\lambda_{\min}$ is then computed directly:

- $\chi = c_2$, in full: $c_2(T\mathsf{CP}^{N_g-1})$ pairs to $\binom{N_g}{2} = N_g(N_g-1)/2$, so $\mathrm{Cap}(N_g) = \lambda N_g(N_g-1)/2$ and admissibility reads $\pi N_g(N_g-1)/2 \le \lambda N_g(N_g-1)/2$, i.e. $\lambda \ge \pi$ *independently of* N_g: cost and capacity are proportional, every $\lambda \ge \pi$ admits all $N_g \ge 3$, no λ excludes anything, and no choice is decisive. A capacity that scales exactly like the cost can never bound it. Excluded.

- $\chi = c_1^2$, in full: $c_1(T\mathsf{CP}^{N_g-1})^2$ pairs to N_g^2, so admissibility reads $N_g(N_g-1)/2 \le (\lambda/\pi)N_g^2$, i.e. $(N_g-1)/(2N_g) \le \lambda/\pi$. The left side increases from $1/3$ at $N_g = 3$ toward $1/2$; consistency (admit $N_g = 3$) requires $\lambda \ge \pi/3$, while decisiveness (exclude $N_g = 4$) requires $\lambda < 3\pi/8$. The window $[\pi/3, 3\pi/8)$ contains no point of the lattice $(\pi/2)\mathbb{Z}$ — its lower edge already exceeds no lattice point below $\pi/2$, and $\pi/2 > 3\pi/8$ — so on the lattice every choice is either inconsistent or indecisive (at $\lambda = \pi/2$, *all* N_g are admitted, since $(N_g-1)/2N_g < 1/2$ always). Excluded.

- $\chi = h^2$: $\lambda_{\min} = 3\pi$; admissible set $\{3\}$. Viable — the constant form (Formulation A of Appendix C), with a residual decisive range closed by minimality.

- $\chi = p_1$: $\lambda_{\min} = \pi$; admissible set $\{3\}$; the decisive unit is unique on the lattice. Viable — the extensive form (Formulation B), strictly the more MSMF-minimal option.

Both survivors saturate exactly at the admissible value: $\mathrm{Cap}(3) = 3\pi = S_{\mathrm{tot}}(3)$ in either form. □

The full admissibility tables. For the record, and for any reader re-deriving the classification, the complete admissible sets (with $N_g \ge 3$ from leg L1) across the lattice:

Table 11: The full classification, for the record and for any reader re-deriving it: each capacity form against each lattice unit, with the admissible set of generation counts in every cell, taken with $N_g \ge 3$ from leg L1. Read across a row and a form's behavior as the unit coarsens is visible; read down a column and the forms are compared at a fixed unit. Exactly one cell in the whole array is the singleton $\{3\}$ at a unit the geometry independently fixes.

	$\lambda = \pi/2$	π	$3\pi/2$	2π	$5\pi/2$
$\chi = h^2$ (Cap $= \lambda$)	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$
… until $\lambda = 3\pi$	$\{3\}$ at $\lambda \in [3\pi, 11\pi/2]$ on the lattice; $\{3,4\}$ from 6π				
$\chi = p_1$ (Cap $= \lambda N_g$)	$\varnothing$	$\{3\}$	$\{3,4\}$	$\{3,4,5\}$	$\{3,\dots,6\}$
$\chi = c_1^2$ (Cap $= \lambda N_g^2$)	all $N_g \ge 3$	all	all	all	all
$\chi = c_2$ (Cap $= \lambda\binom{N_g}{2}$)	$\varnothing$	all $N_g \ge 3$	all	all	all

Reading the table columnwise: the only cell anywhere that is decisive *and* whose unit is unique for its row is (p_1, π); the h^2 row becomes decisive only at 3π and stays decisive across a range (the residual knob closed by minimality); the c_1^2 row admits everything as soon as it admits anything; the c_2 row is all-or-nothing by proportionality. The arithmetic of every cell is one line of the inequality $N_g(N_g-1)/2 \le (\lambda/\pi)\chi(N_g)$.

The naturality fence, stated as such. The classification quantifies over *single* natural characteristic classes. Arbitrary rational combinations of classes could tune any admissible set and are excluded by fiat — this is a definability premise of Axiom 7.11, printed here so that no reader mistakes it for a theorem. Its justification is the usual one for naturality fences: a capacity that is not a single natural invariant of the canonical geometry would itself be a hidden structure demanding explanation, against MSMF.

G.5 The integrality theorems: proof of Theorem 7.13

The conjectural form of the integrality statement — "the descended periods are integrally quantized, so the unit lies on the half-quantum lattice" — bundles three claims with different subjects, different proof technologies, and, it turns out, different grades: a claim about the *class* (does it pair integrally?), a claim about *cycles* (does any holomorphic object carry off-lattice action?), and a claim about the *coupling* (why is the unit itself quantized?). This subsection proves them in that order. The first two are rigorous; the third is the level-quantization argument of the Chern–Simons analogy (§15.3) run quantitatively, derived at physics rigor with its single topological residue stated as Definition G.3.

Component (A): pairing integrality.

Theorem G.1 (L-INT-A)**.** *Under the descended geometry of Theorems 7.7–7.9, the Berry line bundle $L \to \mathsf{CP}^{N_g-1}$ is an honest complex line bundle, its curvature has periods in $2\pi\mathbb{Z}$, and every candidate capacity class of Axiom 7.11 pairs integrally with the generator of $H_4(\mathsf{CP}^{N_g-1};\mathbb{Z})$:*

$$\langle \chi, [\mathsf{CP}^2] \rangle \in \mathbb{Z}, \qquad \chi \in \{h^2,\ c_1^2,\ c_2,\ p_1\}. \tag{78}$$

Proof. The stratum's points are rank-one projectors realized by the runtime as operators on the fixed boundary Hilbert space (Theorem 7.8). The assignment $x \mapsto \operatorname{im} P(x)$ is therefore a rank-one *subbundle of the trivial Hilbert bundle* over CP^{N_g-1}. A subbundle of a trivial bundle is an honest vector bundle: there is no gerbe, no torsion ambiguity, no differential-cohomology subtlety to resolve, because the bundle is constructed, not postulated. Its first Chern class $c_1(L)$ is an element of $H^2(\mathsf{CP}^{N_g-1};\mathbb{Z})$ by the definition of Chern classes — integral before any curvature is mentioned. Theorem 7.9 identifies the descended connection with the Berry connection on L and its curvature with F_{Berry}; Chern–Weil then says $[F_{\text{Berry}}/2\pi]$ represents an integral class, so every period of F_{Berry} lies in $2\pi\mathbb{Z}$ — Dirac quantization, obtained rather than imposed: single-valued parallel transport around every closed loop is automatic for an honest bundle.

For the pairing: $H^*(\mathsf{CP}^{N_g-1};\mathbb{Z})$ is the truncated polynomial ring on the integral generator $h = c_1(\mathcal{O}(1))$, and L is the tautological $\mathcal{O}(-1)$, so $c_1(L) = -h$ (the sign convention of Appendix F; the numerically computed Chern number of the tautological family is exactly -1). Every candidate class is a polynomial in integral classes: $c(T\mathsf{CP}^{N_g-1}) = (1+h)^N_g$ truncated, so $c_1 = Nh$, $c_2 = \frac{1}{2}N_g(N_g-1)\,h^2$, and $p_1 = c_1^2 - 2c_2 = N_g h^2$. Pairing the degree-4 parts against $[\mathsf{CP}^2]$ gives $\langle h^2\rangle = 1$, $\langle p_1\rangle = N_g$, $\langle c_1^2\rangle = N_g^2$, and $\langle c_2\rangle = \frac{1}{2}N_g(N_g-1)$ — all integers. ∎

What made this free. Before the adiabatic theorems, the descended class was a pullback from the bulk, and its integrality was hostage to the bulk's differential-cohomology data — the hard route. The content of Theorem 7.9 is that the class is the stratum's *own*: the Berry bundle exists because the runtime's projectors exist. Component (A) is not a shortcut around Hopkins–Singer; it is the observation that, after the adiabatic theorems, this component of the problem no longer lives where Hopkins–Singer machinery is needed. *Physics intuition:* a Berry phase around a closed loop is a physical, measurable phase; if the flux through a closed two-cycle were not a multiple of 2π, transporting a state around the cycle's boundary two different ways would disagree about physics. The runtime cannot disagree with itself about a projector it constructed — that is the whole proof, stated without bundles.

Component (B): the toll lattice.

Theorem G.2 (L-INT-B)**.** *With the cohort normalization $\int_{\text{line}} \omega_{FS} = \pi$ (Appendix F), the action of every compact holomorphic curve $C \subset \mathsf{CP}^{N_g-1}$ is*

$$\int_C \omega_{FS} \;=\; \pi \cdot \deg(C) \;\in\; \pi\mathbb{Z}_{>0}. \tag{79}$$

In particular no holomorphic object with off-lattice action exists, and the step actions πk of the toll ledger (19) are points of this lattice rather than assumptions about it.

Proof. In cohomology $[\omega_{FS}] = \pi\, h$ with h the integral generator — this *is* the normalization, fixed by the quantum-geometric identification $|F_{\text{Berry}}| = 2\,\omega_{FS}$ of Theorem 7.9 together with the integral c_1 of Theorem G.1. A compact holomorphic curve defines a homology class $[C] = \deg(C)\,[\text{line}]$ with $\deg(C) \in \mathbb{Z}_{\geq 0}$ — an *integer* because $H_2(\mathsf{CP}^{N_g-1};\mathbb{Z}) = \mathbb{Z}$ is generated by the line, and non-negative because holomorphic curves pair non-negatively with the Kaehler class. Hence $\int_C \omega_{FS} = \langle \pi h, [C]\rangle = \pi \deg(C)$. By Wirtinger's theorem (holomorphic curves are calibrated by ω_{FS}) this integral is simultaneously the *area* of C: holomorphic area on the family moduli is quantized in units of π. ∎

The loophole this closes. The exclusion of $N_g \geq 4$ is an inequality between a quadratic expense and a linear income, and the expense is charged on holomorphic representatives. An adversary could ask: is there a cheaper holomorphic object — an off-lattice instanton — that saturates the budget differently? Theorem G.2 answers: there is nothing in the holomorphic category off the lattice $\pi\mathbb{Z}$, at any degree, for any N_g. The expense side of the capacity theorem is bounded below lattice-theoretically, not merely on the representatives exhibited. (Verification: the Fubini–Study areas of the degree-d rational curves $z \mapsto [1 : z : z^d]$, computed in closed form in the companion harness, are exactly $d\pi$ for $d = 1, 2, 3$.)

Component (C): the unit lattice. The remaining question is why the *coupling* λ multiplying an integral pairing should itself be quantized. This is where the Hopkins–Singer weight actually sits. Four steps.

Step (C1) — the premise. The capacity enters the runtime as an action: sector configurations are weighed by phases e^{iS} in which the family term contributes $S_{\text{fam}} = \lambda \times (\text{pairing})$. This is the reading under which the toll ledger was derived and under which the budget competes against the throughput income; it is an interpretive premise of component (C), and it is listed in the statement of Theorem 7.13.

Step (C2) — the naive lattice. Deform the configuration by the generator of the family sector's topological charge: the pairing shifts by 1 (Theorem G.1: the generator pairs to exactly 1). Single-valuedness of e^{iS} under this shift requires $\lambda \in 2\pi\mathbb{Z}$ — the standard bosonic Chern–Simons level quantization.

Step (C3) — the fork: the naive lattice is not viable. Here the classification of Lemma 7.14 turns from bookkeeping into a discovery. The decisive window for the unit is $[\pi, 3\pi/2)$: large enough that the $N_g = 3$ expense 3π fits the income 3λ, small enough that the $N_g = 4$ expense 6π does not fit 4λ. Classify the candidate lattices by their decisive sets (the lattice points whose admissible set, against $N_g \geq 3$ from leg L1, is exactly $\{3\}$):

Table 12: The three candidate lattices — bosonic, spin-refined and the conservative half-quantum fence — classified by their decisive sets, meaning the lattice points whose admissible generation count is exactly $\{3\}$ once leg L1 is imposed. The bosonic lattice has none: on it the argument cannot close at all. The spin-refined lattice has exactly one. That the decisive point is unique, and that it is the one the spin structure independently selects, is the content of the leg.

lattice	decisive set	status
$2\pi\mathbb{Z}$ (bosonic)	$\varnothing$	no decisive point exists
$\pi\mathbb{Z}$ (spin-refined)	$\{\pi\}$	decisive, uniquely
$(\pi/2)\mathbb{Z}$ (conservative fence)	$\{\pi\}$	decisive, uniquely

$2\pi\mathbb{Z} \cap [\pi, 3\pi/2) = \varnothing$: on the bosonic lattice the smallest consistent unit is 2π, which admits $\{3, 4, 5\}$ — consistent but indecisive. If step (C2) were the whole story, the lattice route (Route 1 of Lemma 7.14) would not be incomplete; it would be *wrong*, and only Routes 2–3 would survive. The refinement of the next step is therefore load-bearing for the entire lattice presentation — a fact the half-quantum phrasing of the toll ledger would disguise as a safety margin.

Step (C4) — the spin refinement, and why the runtime has already chosen it. On a spin bulk, the quantization of the relevant quadratic functional is shifted: the flux-quantization argument of Witten (1996), made rigorous by the quadratic refinement of Hopkins & Singer (2005), replaces the naive square x^2 by the refinement $2q(x) = x(x-1)$, and single-valuedness is imposed on $\lambda\, x(x-1)$, whose increment under $x \to x+1$ is $2\lambda x$ — *even* — so $\lambda \in \pi\mathbb{Z}$ suffices. The BI-6 bulk carries $\text{Spin}(1, 5)$ structure (it hosts the cohort's fermions; Arisaka, 2026, A9-SM) — an imported premise, listed in the theorem. And the cohort's own arithmetic shows the runtime uses the refined functional and not the naive square: the cascade's total toll is exactly $\pi\, q(N_g)$ with $q(x) = x(x-1)/2$ the refinement's quadratic form — the half-integer structure is not imported to rescue the argument; it is already present in the object being priced. Hence $\lambda \in \pi\mathbb{Z}$, and by step (C3) the decisive point is uniquely $\lambda = \pi$ — *invariantly*: a reader who doubts that the refinement transfers cleanly from its M-theory home to the BI-6 seam may retreat to the conservative fence $(\pi/2)\mathbb{Z}$, where the decisive point is also uniquely π. Only the non-viable bosonic lattice dissents, and it dissents by emptiness, not by preferring a different value.

The residue.

Definition G.3 (R-INT: the topological residue)**.** R-INT is the statement that the specific 6D BI-6 seam topology — the bulk M_6 with its boundary and the embedded family stratum — globally supports the quadratic refinement: the class $[X_4]$ admits the refinement q compatibly with restriction to the stratum, i.e. the relevant torsion in $H^4(M_6;\mathbb{Z})$ and the w_4-type obstruction vanish for the actual seam.

Everything on the stratum itself is explicit and torsion-free ($H^*(\mathsf{CP}^{N_g-1};\mathbb{Z})$ has no torsion; every class and pairing above is computed). R-INT lives entirely in the *ambient* seam — it is topological bookkeeping on a specific manifold, of the same kind as checking orientability, not a conceptual gap — and it is the declared mandate of Appendix H — met on the benchmark background by its κ-bridge theorem, which proves the pushforward is the Berry bundle and closes both parts of R-INT there. Its failure cost is stated once more for the record: without R-INT, step (C4) loses its geometric license, the lattice retreats to the conservative fence or falls to presentational status — and $\lambda = \pi$ survives on Routes 2–3 of Lemma 7.14 with the value of N_g unmoved.

The premises of Theorem 7.13, inventoried. Component (A): the descended geometry (Born-grade; Theorems 7.7–7.9). Component (B): the same, plus the normalization of (19). Component (C): the same, plus the action-level reading of the coupling clause (step C1) and the bulk spin structure (step C4), with R-INT the residue. No step consumes $N_g = 3$, $K = 35$, or $J = 7$: the decisive-set computations *conclude* $\{3\}$; they do not assume it. The grade vocabulary is the cohort's: components (A) and (B) are theorems given the adiabatic package; component (C) is derived-conditional at physics rigor with its residue named.

Physics intuition. The three components, compressed. (A): the stratum cannot carry fractional flux, because its states are the runtime's own projectors, and a fractional flux would make their transport multivalued. (B): holomorphic area is quantized — the landscape sells relaxation in whole tolls only, so there is no discount instanton. (C): the price of a whole toll is itself pinned by the same consistency that pins a Chern–Simons level — and the bulk's fermions, by forcing spin structure, halve the naive price lattice to exactly the one on which the family count is decided. One bundle underwrites all three: the tolls, the budget, and the unit are three readings of the Berry line bundle that the adiabatic theorem forces.

G.6 Verification record

Every numerical and structural claim of Appendices F and G, and of §7.8, is checked programmatically in the companion harnesses to this paper (Kato transport unitarity and holonomy against the half-solid-angle law; the quantum geometric tensor against the Fubini–Study metric with proportionality constant 1; the period integrals π and -2π; the signature and Morse tables for $N_g = 1, \dots, 8$; the exclusion, unit-classification, and form-robustness tables; the exact Chern number, the degree-d holomorphic areas $d\pi$, and the decisive-set classification of §G.5; and the pointwise vanishing of $\mathrm{tr}(R \wedge R)$ on the canonical bulk by full exterior-algebra recomputation). The falsifiable content of the architecture is thereby reproducible from the paper's own statements.

H The inflow bridge: how AX-B is enforced

§7.6 states the coupling clause AX-B at its true, minimal size and §7.8 spends it. This appendix supplies the pair-level mathematics that enforces it on the benchmark background of the intended GG-6 class, and discharges the topological residue R-INT (Definition G.3) there.

The instrument is the Hopkins–Singer quadratic functor: the fivebrane-pattern pushforward that turns six-dimensional Wu-structure data into a line bundle on the family moduli. Six theorems follow, T-INF-1 to T-INF-6. Their conclusion is the κ-bridge: the bulk's pushed-forward bundle *is* the Berry bundle the boundary's own adiabatic theorems construct (Theorems 7.7–7.9).

Two things should be said before the reader starts, because they bound what follows. *First*, the theorem is an agreement between the two faces of one seam, not a derivation of either face from the other; the bulk side of the computation consumes no quantum-mechanical input, and §H.9 shows the two obstruction theorems of §7.6 reappearing from inside the formalism. *Second*, the numerical scope is the benchmark background: the warped and relative refinements, R-WARP and R-REL, are stated in §H.11 and closed there at class level by the rigidity theorems T-INF-7 and T-INF-8, with what survives of each named exactly.

This appendix is written to be read on its own. It pins its conventions before it computes, and nothing in the main line depends on the machinery here — only on the fact that the clause it enforces can be enforced.

H.1 A century of quantization, in one page

The instrument of this appendix was built over ninety years, by physicists who kept discovering that an integer they had measured was topology they had not yet named, and by mathematicians who then named it. The road is worth walking before the computation starts, because every object in §H.3–§H.8 is the end of it.

H.1.1 The external road: from Dirac's monopole to the quadratic functor

1902: conservation built into geometry. The story begins, fittingly, with an identity rather than an equation of motion. Bianchi (1902) proved that the curvature of a connection automatically satisfies $dR + [A, R] = 0$: geometry polices its own bookkeeping. Everything in this appendix descends from learning, over a century, exactly when that identity can be *modified* and what the modification costs.

1931: the first topological integer of physics. Dirac (1931) showed that a single magnetic monopole anywhere in the universe forces electric charge to be an integer multiple of a unit: the wavefunction of a charged particle transported around the monopole must return to itself, and single-valuedness quantizes the flux. This is the ur-argument of the subject — *a phase that must close forces an integer* — and the reader will meet it again, essentially unchanged, at three places in this appendix: in the Dirac quantization of the Berry bundle, in the level argument behind this paper's unit lattice, and in the Wu-evenness that drives T-INF-3.

1984: the anomaly becomes geometry. Green & Schwarz (1984) discovered that a gauge theory whose quantum phase fails to close — an anomalous theory — can be rescued by *modifying the Bianchi identity*: let a tensor field H_3 satisfy $dH_3 = X_4 \neq 0$, with X_4 built from the very curvatures whose anomaly is to be canceled, and the failure of one bookkeeping exactly pays the failure of the other. The modified identity stopped being a defect and became a design principle. BI-6 is the cohort's Green–Schwarz identity, with its coefficient vector locked by OGN-5 arithmetic rather than chosen (Arisaka, 2026, A3-BI6).

1985: the third datum. A gauge field carries three kinds of information: its charge lattice (Dirac's integers), its flux density (the curvature), and its *holonomies* — the phases around closed paths that neither of the first two determines. Cheeger & Simons (1985) built the mathematical object that carries all three at once, the differential character. The third datum is the one closed-system physics can ignore and open-system physics cannot: torsion and holonomy are invisible until something wraps a cycle, and observation — in the cohort's reading — is precisely a wrapping.

1996: spin halves the lattice, and physics orders mathematics it does not yet have. Witten (1996a) found that the four-form flux of M-theory is quantized not on the integers but on integers *shifted* by half a gravitational class, $\lambda = p_1/2$, whenever the manifold is spin: the fermions' existence changes the bosonic field's lattice. In the same year Witten (1996b) showed that the fivebrane's self-dual tensor field has no Lagrangian at all — its partition function must be *defined* as a section of a line bundle over the space of boundary data, constructed from a quadratic refinement of the flux pairing. The construction sketched there needed mathematics that did not yet exist.

2005: the mathematics arrives. Hopkins & Singer (2005) supplied it: the differential cocycle (c, h, ω) — charge, holonomy, and flux as one object — and the quadratic functor κ, which takes a family of manifolds with a degree-$2k$ differential integral Wu structure and returns, in the fivebrane case, exactly the line bundle Witten's partition function needed. The modern descendants (Dai & Freed, 1994; Freed, 2000; Monnier, 2016; Sati & Schreiber, 2024) turned the construction into a doctrine: *to quantize a higher gauge field is to choose the cohomology theory its fluxes live in*, and a Green–Schwarz-modified Bianchi identity dictates a twisted differential theory. That doctrine is the frame of this appendix; the functor is its only instrument.

The thread, in one sentence. Dirac: a closing phase forces an integer. Green–Schwarz: a modified identity can trade two failures. Cheeger–Simons: keep the holonomy, not just the flux. Witten: spin shifts the lattice, and families of six-manifolds hand you line bundles. Hopkins–Singer: all of the above is one functor. This appendix: applied to the GG-6 seam, that functor's output is the Berry bundle.

H.2 The seam, the stratum, and the functor

H.2.1 The seam (from GG-A3-BI6)

The intended GG-6 background class of Arisaka (2026, A3-BI6) fixes the bulk: $M_6 = M_4 \times I_r \times S^1_\theta$ with M_4 closed, oriented, spin, $H^1(M_4;\mathbb{Z}) = 0 = H^3(M_4;\mathbb{Z})$, torsion-free H_1; APS-compatible collars at the interval ends, or the doubled closed model $M_4 \times T^2$ with the structural torus $T^2 = S^1_r \times S^1_\theta$; gauge group $\mathrm{SU}(5) \times U(1)_B \times U(1)_{Qi}$ with bundles pulled back from the base and trivial on the torus; the modified Bianchi identity $dH_3 = X_4$ with the locked coefficient vector $(1, -1, -\frac{2}{5}, -\frac{2}{5})$; and the primitive integral lift

$$\widehat{X} \;=\; \frac{t}{8\pi^2}\,X_4 \;\in\; H^4(M_6;\mathbb{Z}), \qquad \gamma = \frac{s}{t} = \frac{L-N}{L} = \frac{2}{5}, \quad (s,t) = (2,5). \tag{80}$$

The *benchmark* is $M_4 = S^4$ (the compactified physical base), on which GG-A3-BI6's reservoir existence theorem supplies the unit background flux, written here $\Omega_{\mathrm{orb}} = 1$. Cohomology of the closed benchmark (ranks, degrees 0–6): $[1,2,1,0,1,2,1]$, with $H^2(M_6) = \mathbb{Z}\,\tau$ (τ the structural-torus class, $\int_{T^2}\tau = 1$) and $H^4(M_6) = \mathbb{Z}\,\sigma$ (σ the base class, $\int_{S^4}\sigma = 1$).

H.2.2 The stratum (from §7.1)

The family-projector moduli is CP^{N_g-1} with hyperplane generator $h \in H^2(\mathsf{CP}^{N_g-1};\mathbb{Z})$. The boundary runtime's adiabatic package forces, at the Born grade of conditionality — premises shared with the Constitutional Born Theorem (Arisaka, 2026, A8-QM) — the Berry connection on the tautological bundle $\check{a}$ with

$$c_1(\check{a}) = -h, \qquad \left[\tfrac{F_a}{2\pi}\right] = -h, \qquad |F_a| = 2\,\omega_{FS}, \qquad \int_{\mathrm{line}} \omega_{FS} = \pi, \tag{81}$$

in the step-action normalization of Appendix F. These four equalities are the conventions this appendix computes against; no further sign or factor choices are made downstream.

H.2.3 The functor, and the three factors to watch

Exactly one construction in the mathematical literature has that type from six-dimensional data: the Hopkins–Singer quadratic functor (Hopkins & Singer, 2005). For an $\check{H}$-oriented family $E \to S$ of relative dimension $4k - i$ with $i \le 2$, equipped with a differential integral Wu structure λ of degree $2k$, Hopkins–Singer construct

$$\kappa_{E/S}\colon \check{H}^{2k}_{\nu}(E) \longrightarrow \check{H}^{\,i}(S), \qquad \kappa(\lambda) \equiv \tfrac{1}{8}\int_{E/S}(\lambda \cup \lambda - \check{L}_{4k}) \ \ (\text{mod torsion}), \tag{82}$$

natural under base change and transitive in composite families. In six fiber dimensions the bookkeeping $4k - i = 6$, $i \le 2$ has the unique solution $(k, i) = (2, 2)$: the refinement acts on degree-four differential cohomology, the Wu structure has degree four — and exists canonically, since the Wu class ν_4 vanishes on any six-manifold — and the value $\kappa(\lambda) \in \check{H}^2(S)$ is a *line bundle with connection on the base of the family.* This is the pattern of Witten's fivebrane partition function (Witten, 1996b), whose rigorous form was the motivation of Hopkins & Singer (2005), and of the Wu-structure field theories of Monnier (2016); its flux-quantization reading in the modern sense is Freed (2000); Sati & Schreiber (2024).

The claim of this appendix, proved as Theorem T-INF-6: on the benchmark background of the intended GG-6 class, with the seam fibered over the family-projector moduli,

$$\kappa_{E/S}(\lambda_{\mathrm{BI6}}) \;=\; \check{a} \qquad \text{— the tautological/Berry line bundle, } c_1 = -h, \tag{83}$$

with every integer in the assembly derived: the channel charge $c = L - N = 2$ from the abelian cross term of the primitive BI-6 lift at minimal quanta, and the background number $m = 2\,\Omega_{\mathrm{orb}} = 2$ from the orientation doubling of the unit reservoir. The bulk, pushed forward, is the boundary's bundle.

The functor (82) lives in the differential-character tradition of Cheeger & Simons (1985): its values carry holonomy and torsion, not curvature alone, which is exactly why it can see a coupling that every curvature-only formalism provably misses. Equation (82) carries three factor risks, each audited where it arises: the $\frac{1}{8}$ of the normalization (resolved by the cross-term 2 and the two derived 2's, §H.4–H.6); the $\check{L}$-term (vanishes identically on the product seam, Lemma H.3); and the mod-torsion caveat (empty on the benchmark, where every group in sight is torsion-free — §H.8).

H.2.4 Grade vocabulary

As throughout the cohort: *rigorous* (proof from stated premises), *physics rigor* (the standard argument of the physics literature, premises listed, no analytic gaps filled by fiat), *consistency* (agreement of independently derived structures), with every conditionality declared at the theorem statement, not in a footnote.

H.3 The seam family and the channel theorem

H.3.1 The model

Definition H.1 (The benchmark seam family)**.** The seam family is the product fibration $E = M_6 \times S \to S$, $S = \mathsf{CP}^{N_g - 1}$, $M_6 = S^4 \times T^2$, with vertical tangent bundle pulled back from M_6 and $\check{H}$-orientation fixed by the cohort conventions of §H.2. A *family Wu structure* is a differential integral Wu structure $\lambda \in \check{Z}^4(E)$ of degree four (existence canonical: $\nu_4 = 0$ in six dimensions), with integer class decomposed along Künneth as

$$[\lambda] \;=\; m\,\sigma \;+\; c\,(\tau \boxtimes h) \;+\; e\,(1 \boxtimes h^2), \qquad m, c, e \in \mathbb{Z}. \tag{84}$$

The model's premises, printed as fences: *(P-i)* the fibration is the product — the physical warping cannot move the class (Theorem T-INF-7, §H.11), and the warped representative is deferred to the engineering item E-3; *(P-ii)* family dependence enters through pullback classes, per the pullback-bundle clause of GG-A3-BI6; *(P-iii)* all signs and normalizations are those of (81), fixed before the computation.

Theorem H.2 (T-INF-1: the channel theorem)**.** *On the benchmark, the family-non-trivial part of $H^4(M_6 \times \mathsf{CP}^{N_g-1}; \mathbb{Z})$ has rank exactly two, spanned by $\tau \boxtimes h$ and $1 \boxtimes h^2$; and of these, only the mixed channel $\tau \boxtimes h$ can contribute to $\kappa_{E/S}$'s curvature. The family couples to the pushforward through the structural torus, or not at all.*

Proof. Künneth over $\mathbb{Z}$ (all groups in sight torsion-free): the $q > 0$ part of $H^4(E) = \bigoplus_q H^{4-q}(M_6) \otimes H^q(\mathsf{CP}^{N_g-1})$ receives contributions only from $q = 2$ and $q = 4$, since $H^{\text{odd}}(\mathsf{CP}^{N_g-1}) = 0$ and $H^3(M_6) = 0$: namely $H^2(M_6) \otimes H^2(\mathsf{CP}^{N_g-1}) = \mathbb{Z}\,(\tau \boxtimes h)$ and $H^0(M_6) \otimes H^4(\mathsf{CP}^{N_g-1}) = \mathbb{Z}\,(1 \boxtimes h^2)$. For the second statement: the fiber integration in (82) retains only terms of fiber degree six in the eight-form $\lambda \cup \lambda$; the cross term of $m\,\sigma$ (fiber degree 4) with $c\,(\tau \boxtimes h)$ (fiber degree 2) has fiber degree six, while every term involving $1 \boxtimes h^2$ (fiber degree 0) tops out at fiber degree four and dies under $\int_{E/S}$. (Harness §A: the rank computation; the bidegree table is Appendix A.) ■

Physics intuition. The stratum can only speak to the bulk in a language the bulk's topology understands. On the benchmark the bulk offers exactly one two-dimensional ear — the structural torus of the Kaluza–Klein sector — and the Künneth arithmetic says the family's hyperplane class must pair with it. That the family–bulk conversation runs through the compact directions is not an assumption of this appendix; it is forced by the intended class's own topology, and it lands exactly where GG-A3-BI6's structural-holonomy clause and the Kaluza–Klein architecture of Arisaka (2026, P1-KK) put the compact-direction physics.

H.4 The structure formula

Lemma H.3 (The $\check{L}$-term vanishes on the product seam)**.** *For the product fibration of Definition H.1, the vertical tangent bundle is pulled back from M_6, so every characteristic form of it — in particular the degree-eight integrand of $\check{L}_8$ — is of pure fiber degree; and an eight-form of pure fiber degree on a six-dimensional fiber vanishes identically. The correction term of (82) contributes nothing, at curvature level and exactly.*

Theorem H.4 (T-INF-2: the structure formula)**.** *For a family Wu structure (84) on the benchmark seam family, the curvature of $\kappa_{E/S}(\lambda) \in \check{H}^2(\mathsf{CP}^{N_g-1})$ is*

$$F_\kappa \;=\; \frac{c\,m}{4}\,F_a, \qquad \textit{with underlying class} \qquad c_1\big(\kappa(\lambda)\big) \;=\; -\,\frac{c\,m}{4}\,h. \tag{85}$$

Proof. By Lemma H.3, $\kappa(\lambda) = \frac{1}{8}\int_{E/S} \lambda \cup \lambda$ on the benchmark. Expand $\lambda \cup \lambda$ and keep fiber degree six (Theorem H.2): the only survivor is the cross term $2\,c\,m\,(\sigma \smile \tau) \boxtimes (h\text{-datum})$, whose fiber pairing is $\langle \sigma \smile \tau, [M_6] \rangle = (\int_{S^4} \sigma)(\int_{T^2} \tau) = 1$. Hence $F_\kappa = \frac{1}{8} \cdot 2cm \cdot F_a = (cm/4)F_a$, and the class statement follows from (81). The assembled ratio $F_\kappa/F_a = cm/4$ is verified numerically at test points of the moduli for four (c, m) pairs, with the fiber volumes computed by quadrature (harness §B; Appendix B). ■

The factor audit, first reading. Equation (85) is the honest form of the inflow: the pushforward is *not* automatically the Berry bundle — it is $(cm/4)$ of it. A derivation that proclaimed the identification here, with c and m unexamined, would be exactly the silent-factor failure the Preface promised to avoid. The next three theorems close the factor from inside.

Physics intuition. Where does the 4 in the denominator come from, and who has to pay it? The $\frac{1}{8}$ is the price of the quadratic refinement — the same $\frac{1}{8}$ that normalizes a theta characteristic, hard-wired into the functor by Hopkins & Singer (2005) and not negotiable. The cross term of the square contributes a 2, because the background and the family deformation each appear once in $\lambda \cup \lambda$'s mixed term. Net: the machine pays out one quarter of the product cm. For the output to be exactly one Berry bundle, physics must therefore supply a 4 — and the deep content of the next three theorems is that it supplies it in the only honest way possible: as 2×2, one even factor from each side of the seam, with the asymmetric splits 1×4 and 4×1 forbidden by the Wu condition. The denominator is the functor's; the numerator is the physics'; the theorem is that they meet.

H.5 The quantization theorem

Theorem H.5 (T-INF-3: κ-quantization)**.** *The Wu condition on λ — its integer class congruent to $\nu_4 = 0$ modulo two — forces m, c, e all even in (84). Consequently $cm \in 4\mathbb{Z}$:*

$$c_1(\kappa(\lambda)) \in \mathbb{Z} \cdot h \qquad \textit{always}, \tag{86}$$

and the minimal non-zero value of $|cm/4|$ is 1, attained only at $(|c|, |m|) = (2, 2)$. In particular the factorizations $(1, 4)$ and $(4, 1)$ of $cm = 4$ are excluded, and the minimal non-zero output of the pushforward is exactly the tautological bundle $\pm\check{a}$.

Proof. Evenness of the class is the definition of a differential integral Wu structure at $\nu = 0$; evenness of each coefficient follows because σ, $\tau \boxtimes h$, $1 \boxtimes h^2$ are a basis of a free group. Then $cm \equiv 0 \pmod 4$, so (86); the solution set of $|cm| = 4$ in even integers is $(\pm 2, \pm 2)$ alone (harness §D). Integrality of $c_1(\kappa)$ is guaranteed independently by the Hopkins–Singer theorem; the two routes agree, and the odd control $(c, m) = (1, 2)$ — which would give the non-integral class $-h/2$ — is displayed in the harness as the failure it is (§E). ∎

Physics intuition. This is the third appearance of one phenomenon: a spin/Wu condition halves a naive lattice, and the halving is load-bearing — the pattern whose prototype is Witten's shifted flux quantization (Witten, 1996a). This paper's unit lattice halved from $2\pi\mathbb{Z}$ to $\pi\mathbb{Z}$ (where the decisive point lives); GG-A3-BI6's tensor refinement runs on the characteristic vector $a = 1$; and here the Wu condition quantizes the pushforward on the integer multiples of the Berry class and forbids the asymmetric factorizations of 4. One refinement, three theaters — and in each, the physics sits exactly on the point the refinement creates.

H.6 The attribution theorems

H.6.1 The channel charge: $c = L - N$

Theorem H.6 (T-INF-4: attribution of c). *Let the family-locked abelian excitation of the seam carry torus flux $n_\tau \in \mathbb{Z}$ and family charge $q \in \mathbb{Z}$ under the tautological bundle. Then the abelian quadratic term of the primitive lift (80) writes into the mixed channel the coefficient*

$$c \;=\; \pm\, s\, n_\tau\, q, \qquad s = L - N = 2, \tag{87}$$

so at MSMF-minimal quanta ($n_\tau = q = 1$) the channel charge is $|c| = L - N = 2$: the first forced 2 is the numerator of the trace-ratio lock. Moreover the evenness demanded by Theorem H.5 is automatic: $s = L - N$ is even because N and L are both odd.

Proof. The abelian part of the lift is $\widehat{X} \supset -(s/8\pi^2)(F_B^2 + F_{Qi}^2)$. Writing the family-locked abelian field as $F = F_\tau + qF_a$ with $[F_\tau/2\pi] = n_\tau\, \tau$ and $[F_a/2\pi] = -h$, the cross term is $-(s/8\pi^2)\cdot 2\, F_\tau \wedge qF_a = -s\,(F_\tau/2\pi) \wedge (qF_a/2\pi)$, of class $+\,s\, n_\tau q\, (\tau \boxtimes h)$ — the mixed-channel class with coefficient $\pm s\, n_\tau q$, the sign fixed by the orientation convention. The class number is computed by explicit integration over $T^2 \times$ line in the harness (§A of the attribution harness: $-n_\tau q$ to 10^{-6} for four quanta pairs). Minimality of the quanta is the MSMF selection, stated as a premise; the parity claim is immediate. ∎

Physics intuition. The integer $s = L - N$ has been in the cohort since GG-A3-BI6: it is the numerator of the trace-ratio lock $\gamma = (L - N)/L = 2/5$, the exchange rate at which the bulk prices abelian flux against the gravitational unit — and, in the OGN-5 reading, the count of ledger slots the families leave unused. What Theorem H.6 adds is that this exchange rate is also the *family's toll charge into the mixed channel*: when the minimal family-locked excitation threads one unit of flux through the structural torus with one unit of family charge, the seam's four-form bookkeeping records exactly s units in the channel the pushforward reads. The number was never free: it is the same subtraction, $L - N$, doing its third job.

Remark H.7 (The excitation reading). The pullback-bundle clause of GG-A3-BI6 sets the torus flux to zero *in the background* — and §H.9 shows the background alone pushes forward to the trivial bundle. There is no tension: the family coupling is precisely the minimal excitation in the one sector the background clauses switch off and the topology leaves open. The clause is why the coupling is invisible to every one-sided computation; the excitation is how the pair turns it on.

Remark H.8 (The parity echo, fourth appearance). $s = L - N$ is even iff $N \equiv L \pmod 2$ — with $L = 5$, iff N is odd. For even N the minimal family channel has odd c and violates the Wu condition outright: the coupling would be forced to doubled quanta. This is consistency, not exclusion — but the odd-N condition of the signature route, the stratum-spin reading, and the unit-halving argument here surfaces a fourth time, in the channel arithmetic itself (harness §B, $N = 2$–7 table). At $N = L$ the channel coefficient vanishes identically; nothing in the cohort rides on this curiosity, and it is recorded as one.

H.6.2 The background number: $m = 2\,\Omega_{\mathrm{orb}}$

Theorem H.9 (T-INF-5: attribution of m)**.** *The closed benchmark $M_6 = S^4 \times T^2$ is the orientation double of the physical collar seam $S^4 \times I_r \times S^1_\theta$. Pairings double under doubling, so the unit reservoir flux of GG-A3-BI6's benchmark existence theorem ($\Omega_{\mathrm{orb}} = 1$ on the physical half) lands on the closed model as*

$$m \;=\; 2\,\Omega_{\mathrm{orb}} \;=\; 2. \tag{88}$$

The second forced 2 is the orientation double. Moreover $m = 2\Omega_{\mathrm{orb}}$ is even for every *reservoir value: the closed model is always Wu-compatible, so the Wu condition constrains the closed bookkeeping only, never the physical seam; and $m \neq 0$ iff the reservoir is non-trivial — the non-vanishing of the coupling has a concrete bulk carrier.*

Proof. The doubling identity $\langle \widehat{X}, [\mathrm{double}]\rangle = 2\,\langle \widehat{X}, [\mathrm{half}]\rangle$ is the reflection computation, verified to 10^{-9} on an arbitrary collar profile in the harness (§C); the identification of the closed benchmark as the double, and the unit reservoir, are the constructions of Arisaka (2026, A3-BI6) (APS doubling; benchmark reservoir existence), imported at their own grade. Evenness and the non-vanishing statement are immediate from (88). ■

Physics intuition. The second 2 is a mirror. The theorems of GG-A3-BI6 are proved on a closed model of the bulk, built by reflecting the physical collar geometry through its end — the standard Atiyah–Patodi–Singer doubling. An observer computing a flux integral on the closed model sees every physical flux twice: once directly, once in the mirror. So the unit reservoir of the physical seam appears as 2 in the closed bookkeeping — and, pleasingly, this also explains why the Wu condition never threatens the physics: *any* integer, doubled, is even. The evenness that Theorem H.5 demands of m is not a constraint on nature; it is an automatic property of reflections. Nature keeps the unit; the mirror supplies the two.

H.6.3 The placement of the physical datum, decided by the pair

Proposition H.10 (The placement fork)**.** *Two placements of the BI-6 datum in the quadratic machine are a priori available: (A) the family coupling deforms the Wu structure λ itself; (B) it rides in the refined field $\check{x}$, the physical bundle being $q^\lambda(\check{x}) = \kappa(\lambda - 2\check{x})$. Under (B), with λ gauge-fixed to the minimal even class, the pushed-forward bundle has curvature $c_x m_x F_a$ with $|c_x| = s = 2$, hence $|c_1| \geq 2$. The boundary's bundle has $c_1 = -1$ exactly (Theorem 7.9). Placement (B) is refuted by the pair; placement (A) is selected.*

This determination consumes boundary data by design. That is not a weakness to be disguised: the boundary-relativity proposition says the coupling is a relative invariant of the pair, so a pair-level matching is the only admissible fixing of the placement — and one side of the pair, the boundary, has carried its answer since the adiabatic theorems.

Physics intuition. The quadratic machine has two slots: the Wu structure λ — the *ruler*, the background datum against which fluxes are measured — and the refined field $\check{x}$ — the *measured* *length.* The placement question asks which slot the BI-6 datum occupies, and it is a genuine question because the formalism accepts either. But the two placements predict different bundles: the ruler slot delivers one Berry unit, the field slot at least two. The boundary, which has possessed its bundle since the adiabatic theorems, measures one unit. The BI-6 datum is the ruler. A formalism in which both slots gave the same answer would have made the question metaphysical; this one makes it empirical — internal to the theory, but empirical in kind.

H.7 The κ-bridge theorem

Theorem H.11 (T-INF-6: the κ-bridge)**.** *On the benchmark seam family, with the placement of Proposition H.10 and the attributions of Theorems H.6–H.9,*

$$c_1(\kappa_{E/S}(\lambda_{\mathrm{BI6}})) \;=\; -\frac{(L-N)\,\Omega_{\mathrm{orb}}}{2}\,h \;=\; -\,h \qquad \textit{at unit quanta:} \tag{89}$$

$\kappa_{E/S}(\lambda_{\mathrm{BI6}}) = \check{a}$, the tautological/Berry line bundle of the boundary, sign and normalization included. Grade: physics rigor*, conditional on three declared inputs: the channel model of Definition H.1 (itself topology-forced on the benchmark by Theorem H.2), the doubling construction and unit reservoir of GG-A3-BI6, and MSMF-minimal quanta.*

Proof. Assembly: $c = \pm(L-N) = \pm 2$ (Theorem H.6), $m = 2\Omega_{\mathrm{orb}} = 2$ (Theorem H.9), so $cm/4 = \pm 1$ in (85); the sign is fixed by the declared orientation (the pairing positive against the cascade orientation, Appendices E–F), giving $c_1 = -h$ and the curvature $F_\kappa = F_a$ on the nose. Uniqueness of the benchmark Wu structure (§H.8) ensures no residual choice hides in λ's cocycle data. ■

Physics intuition: the seam agrees with itself. Read the chain end to end. On the boundary side, quantum mechanics alone — the adiabatic theorem at the Born rule's own grade — builds a line bundle on the family moduli: the Berry bundle, curvature twice the Kähler form. On the bulk side, the BI-6 seam carries a primitive four-class whose only family-visible channel runs through the Kaluza–Klein torus; the trace-ratio lock writes the integer $L-N$ into that channel, and the orientation double writes $2\,\Omega_{\mathrm{orb}}$. The quadratic functor — the same machine that makes the fivebrane partition function well-defined — converts that bulk data into a line bundle on the same moduli. The two bundles agree exactly, factor and sign. Each side was derived with no knowledge of the other; the agreement is the pair-level fact the two obstruction theorems predicted had to exist and forbade either side to own alone. The books of the seam balance — which is, at bottom, the same sentence the capacity theorem speaks at $N=3$.

H.8 R-INT discharged on the benchmark

This paper's integrality theorem (Theorem 7.13) names as its single topological residue R-INT: that the specific six-dimensional seam supports the quadratic refinement globally — no torsion obstruction in H^4, no w_4-type obstruction, and a Wu-structure choice compatible with the stratum, uniformly in the family. On the benchmark, every part of this is now closed.

Proposition H.12 (The absolute part)**.** *On the intended class $M_6 = M_4 \times T^2$ with the clauses of §H.2:* (i) $\operatorname{Tors} H^4(M_6;\mathbb{Z}) \cong (\operatorname{Tors} H_1(M_4))^3 = 0$ *— the integral lift* (80) *is torsion-unambiguous;* (ii) $\nu_4(M_6) = 0$ *canonically ($\nu_j = 0$ for $2j > 6$; Milnor & Stasheff, 1974), so degree-four Wu structures exist, the choices a torsor over $H^3(M_6;\mathbb{Z}/2) \cong H^2(M_4;\mathbb{Z}/2)^{\oplus 2}$;* (iii) $w_4(M_6) = w_4(M_4) = \chi(M_4) \bmod 2 = \sigma(M_4) \bmod 2 = 0$*, the spin clause doing the killing through Rokhlin's divisibility (Rokhlin, 1952).*

Proof. (i) Künneth with $H^*(T^2)$ free, then universal coefficients and Poincaré duality on the closed oriented four-base: $\operatorname{Tors} H^4 = 0$, $\operatorname{Tors} H^3 \cong \operatorname{Tors} H^2 \cong \operatorname{Tors} H_1$; the clauses kill it. (ii) is the dimensional vanishing plus the standard torsor statement; the mod-2 Künneth uses $H^1 = H^3 = 0$ on the base. (iii) Whitney with the parallelizable torus, the top Stiefel–Whitney class as the mod-2 Euler characteristic, $\chi \equiv \sigma \pmod 2$, and $\sigma \equiv 0 \pmod{16}$ for spin. The non-spin control CP^2 ($\chi = 3$) shows the spin clause is load-bearing. (Harness record: Appendix B.) ■

Proposition H.13 (The relative/family part, on the benchmark)**.** *On the benchmark base $M_4 = S^4$:* (i) *the Wu torsor $H^3(S^4 \times T^2;\mathbb{Z}/2) = 0$, so the degree-four Wu structure is* unique*;* (ii) *restriction and doubling maps land in torsors over the same vanishing group, so the unique structure is automatically compatible with the APS collars;* (iii) *the product family pulls it back to a single uniform choice over CP^{N_g-1}. Off the benchmark (a $b_2 > 0$ base of K3 type, torsor of order 2^{44}) one uniform family choice always exists by pullback, and a canonical basepoint is supplied by the pullback-bundle clause; canonicity, not existence, is the residual question there.*

Corollary H.14. *R-INT holds on the benchmark background: the refinement exists, is unique, restricts compatibly, and is uniform over the stratum. This paper's conditioning ledger, restricted to the benchmark, carries no remaining topological check.*

Physics intuition. What would it have meant for R-INT to *fail*? A torsion class is a ghost flux: a field configuration invisible to every curvature measurement — it integrates to zero against all forms — yet capable of shifting quantized phases, of making two apparently identical backgrounds disagree about an integer. Had H^4 of the seam carried torsion, the BI-6 lift would have been ambiguous by exactly such ghosts, and the capacity arithmetic of §7.8 would have inherited a hidden $\pm$-ambiguity no experiment on curvature could resolve. The content of Proposition H.12 is that on the intended class there is nowhere for a ghost to live: the base's clauses evacuate the torsion, the dimension evacuates the Wu class, and the spin condition — through Rokhlin's theorem, the same theorem that polices the stratum in §7.8 — evacuates w_4. The seam is ghost-free, and the ledger's last topological worry with it.

H.9 The obstruction echoes, and the parity echo

Proposition H.15 (The rigidity echo)**.** *A family-independent Wu structure ($c = 0$ in (84)) has $\kappa_{E/S}(\lambda)$ trivial: by naturality under base change, the pushforward of an S-independent structure is pulled back from a point, and $\check{H}^2(\mathrm{pt}) = 0$. The bulk alone sends nothing to the stratum — the conclusion of the rigidity lemma (Lemma 7.5), reproduced as a two-line corollary of functoriality.*

Proposition H.16 (The relativity echo)**.** *κ consumes, by its type, the bulk Wu structure together with the fibration over the stratum; boundary-only data cannot be supplied to it. That the coupling can arise only from the pair — the boundary-relativity proposition (Proposition 7.6) — is here a type constraint, visible before any computation.*

These two echoes are the paper's quiet consistency theorem: the formalism did not have to be told the obstruction structure of the problem; it contains it. Together with the fourth appearance of the odd-N parity condition in the channel arithmetic (Remark H.8), they are the strongest internal evidence that the quadratic functor is not an imported metaphor but the native mathematics of the seam.

H.10 The bidegree bookkeeping of T-INF-2

Write forms on $E = M_6 \times S$ by bidegree (p, q) — fiber degree p, base degree q. The fiber integration $\int_{E/S}$ annihilates every component with $p \neq 6$. The family Wu structure (84) has curvature components

$$\omega_\lambda \;=\; \underbrace{m\,\omega_\sigma}_{(4,0)} \;+\; \underbrace{c\,\omega_\tau \wedge \pi_S^* F'_a}_{(2,2)} \;+\; \underbrace{e\,\pi_S^*(\text{4-form on } S)}_{(0,4)}, \tag{90}$$

with ω_σ, ω_τ the normalized representatives ($\int_{S^4} \omega_\sigma = \int_{T^2} \omega_\tau = 1$) and F'_a the curvature datum of the h-class in the mixed channel. The square:

Table 13: Every term of the square expanded by bidegree, whether it survives the integration over the fiber, and what it contributes if it does. Six of the seven vanish, and the table records why each one does — degree too high for the fiber, or a squared representative that is itself zero — rather than leaving the survivor to be taken on trust. Exactly one term survives, and it is the one the identity is built from.

term	bidegree	survives $\int_{E/S}$?	contribution
$(4,0)\cdot(4,0)$	$(8,0)$	no ($p = 8 > 6$)	—
$2\,(4,0)\cdot(2,2)$	$(6,2)$	**yes**	$2cm\,(\omega_\sigma \wedge \omega_\tau) \otimes F'_a$
$2\,(4,0)\cdot(0,4)$	$(4,4)$	no	—
$(2,2)\cdot(2,2)$	$(4,4)$	no ($\omega_\tau^2 = 0$ too)	—
$2\,(2,2)\cdot(0,4)$	$(2,6)$	no	—
$(0,4)\cdot(0,4)$	$(0,8)$	no	—
$\check{L}_8$-term	$(8,0)$	vanishes identically	Lemma H.3

Hence $\int_{E/S}[\lambda \cup \lambda] = 2cm\,F'_a$ with the fiber pairing $\int_{M_6} \omega_\sigma \wedge \omega_\tau = 1$, and κ's curvature is $\frac{1}{8} \cdot 2cm\,F'_a = \frac{cm}{4} F'_a$. Identifying F'_a with the Berry curvature F_a of (81) — the content of the placement selection, Proposition H.10 — gives (85). One subtlety deserves its own sentence: the $(3,1)$-type components absent from (90) are absent because $H^3(M_6) = 0$ and $H^{\text{odd}}(\mathsf{CP}^{N_g - 1}) = 0$ leave them no cohomology to live in; at cocycle level a $(3,1)$ exact term can be added, but it shifts κ by an exact form and moves nothing.

H.11 The two refinements, closed at class level; and the channel identity

Two refinements bound the scope of the six theorems above. The physical seam is warped, while the benchmark computes on the metric product (R-WARP); and the physical statement lives on a bounded half-seam, while the global theorems are proved on its closed reflection-double (R-REL). This subsection states both precisely and then proves that neither can move the class: two theorems, T-INF-7 and T-INF-8, close the kill criteria K-INF-1 and K-INF-2 at class level, with every step of the derivation displayed. What survives of each refinement is then stated exactly — one constructive computation and one basepoint convention — and the channel identity, which is a different kind of item and remains fully live, closes the subsection.

H.11.1 R-WARP: the warped-metric refinement, and the warp-rigidity theorem

The refinement, stated. The physical bulk metric of the intended class is warped — the scale-profile metric (62) — while every numerical statement of §H.3–§H.8 is computed on the metric product. The topological chain — T-INF-1, T-INF-3, the attributions, §H.8 — is homotopy-invariant and survives warping unchanged. The question the refinement poses is whether the *form-level* corrections from the non-flatness of the fibration can accumulate into a shift of the *class*: whether a warped recomputation could return $c_1(\kappa(\lambda_{\mathrm{BI6}})) \neq -h$.

Lemma H.17 (Class functoriality of the pushforward)**.** *The Hopkins–Singer pushforward, restricted to underlying topological classes, is metric-free: it is the topological integration (umkehr) map determined by the fibration, the orientation, and the Wu-structure data alone. A metric enters the construction only through the choice of differential representatives, and two such choices differ by exact and flat data, which the passage to classes annihilates.*

Proof. This is the defining compatibility of the differential-cohomology integration map (Hopkins & Singer, 2005): the functor fits the character diagram whose curvature leg is metric-dependent and whose class leg is the ordinary topological pushforward. Concretely, a differential class $\check{x}$ has curvature $R(\check{x})$ (a form, metric-sensitive through its representative) and underlying class $c(\check{x}) \in H^*(\cdot;\mathbb{Z})$ (metric-blind); the integration map commutes with both legs, and on the class leg it is the umkehr map of the fibration. Changing the metric changes the horizontal distribution and hence the representative by a shift lying in the image of exact forms and flat classes — precisely the data killed by c. ■

Theorem H.18 (T-INF-7: warp rigidity of the bridge class)**.** *Let g_s, $s \in [0,1]$, be any smooth family of metrics on the fixed seam manifold, interpolating from the metric product ($s = 0$) to the physical warped metric ($s = 1$), preserving the underlying smooth manifold and its fibration structure. Then the underlying class of the pushed-forward bundle is constant in s:*

$$c_1(\kappa^{(s)}_{E/S}(\lambda_{\mathrm{BI6}})) \;=\; -h \qquad \text{for every } s \in [0,1]. \tag{91}$$

In particular the kill criterion K-INF-1 cannot fire under any smooth warp: a warped recomputation that shifts the class off $-h$ would have to change the fibration topology, the class $[\lambda_{\mathrm{BI6}}]$, or pass through a singular metric, none of which a warp does.

Proof. The proof has three steps, each closing one road by which the warp could act.

Step 1 (the inputs are metric-independent). Every integer entering the bridge is topological data of the fixed manifold. The class $[\lambda] = m\,\sigma + c\,(\tau \boxtimes h) + e\,(1 \boxtimes h^2)$ and its Künneth decomposition (84) live in H^4 of the fixed manifold; the channel theorem (T-INF-1) is a rank statement in the same cohomology; the Wu-structure existence, uniqueness, and collar-compatibility clauses of §H.8 rest on $\operatorname{Tors} H^4$, ν_4, and w_4, all diffeomorphism invariants; and the attributions $c = \pm(L - N)$ (Theorem H.6, a Chern number pinned by the trace-ratio lock) and $m = 2\,\Omega_{\mathrm{orb}}$ (Theorem H.9, a count fixed by the APS doubling) are integers owned upstream of any metric. A warp is a smooth deformation of the metric on a *fixed* manifold with a *fixed* fibration: it moves none of these. (This step is the homotopy-invariance observation already recorded above, now doing load-bearing work.)

Step 2 (the pushforward cannot transmit the warp to the class). By Lemma H.17, the class leg of the pushforward is the topological umkehr map: metric-free. The warp acts on the construction only through the differential representatives — the connection on $\kappa(\lambda_{\mathrm{BI6}})$ and its curvature acquire warped corrections — and those corrections lie in exact and flat data, which the class annihilates. That warping "moves exact and flat data only, not the class" is therefore not an expectation requiring case-by-case verification: it is the class functoriality of the integration map, and it holds for every smooth metric at once.

Step 3 (continuity and integer rigidity close the gap). Suppose, for contradiction, that $c_1(\kappa^{(s)}(\lambda_{\mathrm{BI6}}))$ varied with s. Along a smooth metric family on a compact seam the construction is continuous in s: the fibration is fixed, $[\lambda_{\mathrm{BI6}}]$ is fixed (Step 1), and the warp factor $e^{\pm 2r}$ is smooth and bounded on the compact radial interval of the intended class, so no singular metric intervenes. A $\mathbb{Z}$-valued invariant depending continuously on s is constant. The only escapes from this rigidity are a discontinuity of the construction — which requires the fibration or the Wu data to jump, excluded by Step 1 — or a singular metric on the path, excluded above. Hence (91), and with it the class-level closure of K-INF-1. ■

Grade, and what survives. T-INF-7 is carried at physics rigor: its one imported ingredient is the class functoriality of the Hopkins–Singer integration map (Lemma H.17), which is that map's defining property rather than an additional hypothesis. What survives of R-WARP is exactly the constructive computation: exhibiting the warped curvature representative of $\kappa(\lambda_{\mathrm{BI6}})$ and watching its corrections land in exact and flat data, as Step 2 guarantees they must. That computation can illustrate the mechanism; by T-INF-7 it can no longer refute the class. It is accordingly reclassified as engineering (item E-3 of §12.1) and is not conditioning.

Physics intuition: the warp bends the rulers, not the ledger. A metric is a system of rulers; a warp is a smooth regrading of those rulers along the radial direction. Curvature is where the rulers live, and the warped seam's curvature representative differs from the product's at every point. But the bridge equation $c_1(\kappa(\lambda_{\mathrm{BI6}})) = -h$ is an equation between *integers* — a Chern class paired against cycles — and integers do not live in the rulers; they live in the topology the rulers are laid over. Deforming the rulers smoothly can no more change that integer than re-drawing a coastline's map grid can change the number of islands. The measured instantiation (§H.12, fourth harness): a representative deformation moving the pointwise curvature by twelve hundred percent moves the class by zero at 10^{-12}, and the model family's Chern number is -1 unchanged at warp anisotropies from e^{-7} — the cascade's own strength — to six decades, failing only when a coefficient reaches exactly zero, which is a change of topology, not of metric.

H.11.2 R-REL: the relative recomputation, and the relative-closure theorem

The refinement, stated. The doubling derivation of m (Theorem H.9) says the physical statement lives on the bounded half-seam; the fully honest object is therefore the relative differential cohomology of the pair, with the Dai–Freed collar formalism (Dai & Freed, 1994) carrying the boundary terms. The question the refinement poses is whether the relative computation could return a boundary class different from the absolute one — which would expose a factor error in the doubling bookkeeping.

Theorem H.19 (T-INF-8: relative closure on the benchmark)**.** *On the benchmark seam family, with the Wu structure of §H.8:*

- *(a)* (Filling independence.) *The half-seam pairing is well-defined modulo the period lattice, independently of the choice of APS-compatible filling of the boundary data.*
- *(b)* (Half equals half the double.) *The relative class equals one half of the double's class, and the apparent half-integer risk is void: the doubled datum is automatically even.*
- *(c)* (Class-level closure of K-INF-2.) *A correct relative computation must reproduce the absolute class: the two cannot differ modulo the lattice, because their difference is itself a closed-manifold lattice element.*

Proof. (a) Let Z denote the half-seam with boundary Y, and let Z_1, Z_2 be two extensions of the boundary data compatible with the APS collar. Glue them along Y with one orientation reversed: $X = Z_1 \cup_Y \overline{Z}_2$ is a closed manifold of the intended class. The two relative pairings differ by the absolute pairing evaluated on X. On the intended class that absolute pairing lies in the period lattice: components (A) and (B) of the integrality theorem (Theorem 7.13) are rigorous, and the topological residue R-INT is discharged on the benchmark (§H.8: the Wu structure exists, is unique, restricts compatibly to the APS collars, and is uniform over the family stratum — the collar-compatibility clause of Proposition H.13(ii) is what licenses the glueing). Hence the difference of the two relative pairings is a lattice element, and the relative class is filling-independent modulo the lattice.

(b) Take Z_2 to be the mirror image of Z_1. The glued manifold is then the reflection-double, and the doubling identity $\langle \widehat{X}, [\text{double}] \rangle = 2 \langle \widehat{X}, [\text{half}] \rangle$ (the reflection computation of Theorem H.9, harness-verified to 10^{-9}) gives the relative class as one half of the double's. The half-integer worry — that half of an odd lattice element might fall off the lattice — is void by the evenness mechanism of Theorem H.9: the mirror shows every unit of flux twice, so the doubled datum is $2m$ with m the physical integer, and half of it is m, on the lattice by construction. Any integer, doubled, is even; nature keeps the unit and the mirror supplies the two.

(c) Immediate from (a) and (b): the relative and absolute classes cannot differ modulo the lattice, because their difference is the absolute pairing on a closed glued manifold, which (a) has already placed on the lattice. A relative recomputation returning a different class would therefore not be revealing new relative physics; it would be committing an arithmetic error, and the doubling

identity that guards that arithmetic is verified to 10^{-9} in §H.12 and independently to 10^{-12} in the fourth harness. ■

Grade, and what survives. The glueing argument of (a) is the standard relative-from-absolute mechanism of the collar calculus (Dai & Freed, 1994); every load-bearing input — the lattice, the discharge of R-INT, the collar compatibility — is a benchmark result already proved above, which is why the theorem is stated on the benchmark and carries its grade. Two items survive. First, the formal Dai–Freed write-up: packaging (a)–(c) inside the relative differential cohomology of the pair, with the bonus that the relative frame connects κ directly to the boundary Chern–Simons data that the OGGEP projection pushes down — reclassified as engineering (item E-4 of §12.1), an opportunity rather than a gap. Second, and genuinely open: off the benchmark, on a $b_2 > 0$ base of K3 type, the Wu torsor has order 2^{44} and a uniform family choice exists by pullback, but the *canonicity* of the basepoint is undetermined — exactly the fence Proposition H.13 already declares. That question bounds the scope of the theorems, not their truth on the physical background, and nothing N_g-facing routes through it.

Physics intuition: two observers filling the same boundary. Two mathematicians, handed the same boundary data, will cap it off differently — and the difference between their answers is not arbitrary: it is the books of a closed world, the world obtained by glueing one cap to the other's mirror. Closed worlds keep integer books on the intended class; that is the absolute integrality theorem. So the two cappings can disagree only by an integer, which is to say, modulo the lattice, they cannot disagree at all. The half-seam's books are exactly as honest as the double's — the measured instantiation (§H.12, fourth harness): the fluxes through two different caps of the same loop differ by precisely the closed-surface integer, at four latitudes, to 10^{-12}.

H.11.3 The channel identity

Which abelian factor — $U(1)_B$, $U(1)_{Qi}$, or a combination — carries the family charge $q = 1$ is not determined here; both slots carry the same s, so the arithmetic of Theorem H.6 is identical in all cases. The identification is a physics refinement for a later edition, with the matter assignments of Arisaka (2026, A9-SM) the natural arbiter.

H.11.4 Declared exposure

A numbered member of the branch declares what would kill it. This member's exposure is *mathematical*, and it says so plainly: it carries no independent observational falsifier — its empirical exposure routes entirely through this paper's declared falsifier ledger (§12.2), since every physical claim it touches enters through the main line's conditioning. What it stakes of its own are three computations, whose standing T-INF-7 and T-INF-8 have made unequal:

K-INF-1 *(the warped computation).* By Theorem T-INF-7 the class of $\kappa(\lambda_{\mathrm{BI6}})$ cannot shift off $-h$ under any smooth warp, so this criterion can no longer fire against the outcome of the warped recomputation. It remains live against the theorem's own premises: a demonstration that the physical seam's fibration or the class $[\lambda_{\mathrm{BI6}}]$ is mis-identified would void Step 1 of the proof and reopen the computation (§H.11.1).

K-INF-2 *(the relative computation).* By Theorem T-INF-8 a correct relative computation must reproduce the absolute class, so this criterion can no longer fire against the outcome. It remains live against the arithmetic it always guarded: a factor error in the doubling bookkeeping of Theorem H.9, which the harnesses bound at 10^{-9} and 10^{-12} (§H.11.2).

K-INF-3 *(the channel identity; fully live).* If the matter assignments of Arisaka (2026, A9-SM), once brought to bear, force the minimal family charge away from $|q| = 1$ in every abelian slot, the MSMF-minimal-quanta premise of Theorem H.6 fails in the physical seam and the attribution of c must be recomputed at the forced quanta.

Any of the three would be reported as what it is. None of them touches the family number, which this appendix does not carry.

H.11.5 Honest limits

Three sentences a careful reader deserves in the main text, not a footnote. *First*, T-INF-6 does not select N: the identity $|L-N| = 2$ holds at $N = 7$ as well, and at other odd N the bridge delivers a higher multiple of the Berry class; the theorem is consistency for the coupling clause, and the family number stands on the capacity theorem of §7.8 alone. *Second*, the MSMF-minimal-quanta premise does real work; a reader who declines it retains Theorem H.5's statement that the Berry bundle is the *minimal non-zero* value of the pushforward — the bridge at its unconditional core. *Third*, the benchmark is one background; the intended class is a family of them, and Propositions H.12–H.13 say exactly which parts extend (the absolute part, everywhere on the class) and which acquire a basepoint convention ($b_2 > 0$ bases) — a division T-INF-7 and T-INF-8 sharpen but do not change: the basepoint convention is the one item of the pair-level story that remains a convention.

H.12 Verification record

Four companion harnesses, one hundred five checks, zero failures, archived in the toolchain home alongside the paper source.

`verify_hs_review.py` (35/35). The Hopkins–Singer bookkeeping this appendix stands on: the (k, i) type-check in six dimensions (unique solution $(2, 2)$; integer values would need an eight-manifold); the Wu-class vanishing table; the Künneth/torsion audit of $H^4(M_4 \times T^2)$ on S^4, K3, $S^2 \times S^2$ with controls; the parity chain $\chi \equiv \sigma \bmod 2$, Rokhlin, w_4, with the non-spin control CP^2; the ledger and coefficient arithmetic ($s = L - N = 2$, $t = 5$ clearing, $\gcd(2, 5) = 1$); and the toll ladder's exact satisfaction of the quadratic-refinement functional equation at characteristic element $a = 1$.

`verify_ckap.py` (28/28). The structure and quantization theorems: the channel ranks (T-INF-1); the assembled curvature ratio $F_\kappa/F_a = cm/4$ at moduli test points for four (c, m) pairs, with

$\mathrm{Vol}(S^4)$ by four-angle quadrature to 10^{-6} and the Berry Chern number to 10^{-8} (T-INF-2); the $\check{L}$ dimension-vanishing; the even-solution table of $cm = 4$ — the impostors $(1,4),(4,1)$ dying, $(\pm 2, \pm 2)$ surviving (T-INF-3); the integrality cross-check with the odd control; the (R-b) torsor arithmetic (benchmark 0; K3 2^{44}); and the rigidity echo.

`verify_ckap2.py` (27/27). The attributions: the cross-term class number by explicit $T^2 \times$ line integration ($-n_\tau q$ to 10^{-6}, four quanta pairs; T-INF-4); the parity table $N = 2$–7 (Wu-admissible iff N odd); the reflection-doubling identity to 10^{-9} and the evenness of $2\Omega_{\mathrm{orb}}$ for every reservoir value (T-INF-5); the placement fork ($|c_1| \geq 2$ under placement (B), refuted by the boundary's $c_1 = -1$); and the final assembly $c_1(\kappa) = -(L - N)\Omega_{\mathrm{orb}}/2\,h$ at generic odd N, with the no-insertion check ($N = 3$, $K = 35$, $J = 7$ nowhere consumed).

`open_core_check.py` (15/15). The closure theorems of §H.11, measured: class invariance under representative deformation (1200% pointwise, zero class shift at 10^{-12}); integer rigidity of the model Chern number from e^{-7} to six-decade anisotropy, the zero-coefficient control the only genuine kill; filling independence and the doubling identity, both to 10^{-12}; and the cascade-saddle link invariants.

Acknowledgments

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

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

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

Part VII
Common Objections and Responses

The strongest objections a critical reader should raise against a three-leg derivation of the family number — stated at full force and answered, each response beginning with what the objection gets right

I Common objections and responses

A paper that claims to force $N_g = 3$ from geometry will be met — rightly — with skepticism from every quarter, and a short paper invites the sharpest forms of it. This Part collects twenty-one objections — fourteen on the three-leg derivation itself, then seven on the inflow bridge of Appendix H — each stated at full force in the objector's own voice, with a response that begins by conceding whatever the objection gets right — the cohort's convention, following the closing objections section of Arisaka (2026, A2-GG6). Where the honest answer is "not yet," the response says so and points to the registered debt.

Objection 1 (O-N3-1) (the circularity skeptic)

"You prove $N_g \geq 3$ from observed CP violation. Using the world to count the world is not a derivation."

Response. Partially right, and the paper says so before the objector can: Theorem 6.1 is empirically conditioned on $J_{\text{CKM}} \neq 0$, and the conditioning is declared in §5.4, in the abstract, and at every layer it touches. Three replies. First, economy: the input is a single sign — CP violation exists — not a fitted number; every physical theory admits some empirical content, and this is among the smallest on record. Second, the conditioning touches one leg of three: the upper bound (Theorem 7.18) and the constructive realization (Theorem 8.2) owe nothing to the measured J_{CKM}. Third, the input is scheduled for retirement: the CPT companion GG-P5-CPT (Arisaka, 2026, P5-CPT) already evaluates the Jarlskog invariant at the stated layer, and the constructive closure-level derivation is the declared mandate of the CKM-extension paper (Arisaka, 2026, CKM, in preparation). If that program completes, the leg becomes over-determined and the objection dissolves; until then the paper claims exactly what it can support.

Objection 2 (O-N3-2) (the conjecture skeptic)

"The cohort's cross-reference ledgers call this cap the 'Budget Conjecture'. Renaming it a 'capacity theorem' does not make a century-old question close."

Response. The renaming is earned, and the ledger is printed. What the conjectural reading treats as one opaque assumption splits into parts with three different fates. *Proved:* the adiabatic theorems T-AD-1–T-AD-3 (at Born-grade), the non-vanishing of the coupling (Lemma 7.12), the unit $\lambda = \pi$ (three routes, Lemma 7.14), extensivity via the Pontryagin scaling, and form robustness (Theorem 7.17). *Proved impossible:* any derivation of the cap from the prior axioms alone — the two obstruction theorems of Appendix G, which are why the conjecture stood unproven for as long as it was posed that way. *Still open:* exactly one residue, R-INT (Definition G.3) — the topological check underwriting the spin refinement in Theorem 7.13 — whose failure would moreover cost the lattice presentation but not the value (Routes 2–3). A reader prepared to grant the five-assumption axiomatization is being offered strictly less to swallow; a reader who rejects the coupling clause knows precisely which single clause they are rejecting and what evidence (the classification, the saturation identity, the falsifiers) bears on it.

Objection 3 (O-N3-3) (the postulate skeptic)

"Morse positivity is assumed, not derived. Postulates dressed as machinery."

Response. The objection is answered by architecture rather than by defense: Morse positivity carries nothing in this derivation. The exclusion of $N_g = 2$ and of all $N_g \geq 4$ runs entirely through the capacity theorem (§7.8), which nowhere invokes Postulates 7.19 or 7.20; the postulates survive as an imported consistency check, and at the admissible value the strict form holds as a theorem, $a^+(3) - a^-(3) = \sigma(\mathsf{CP}^2) = 1$ (§7.11). The residue of the objection lands on the one place assumption survives — the coupling clause — and there the paper directs the skeptic to Objection 13, where that challenge is put at full force.

Objection 4 (O-N3-4) (the three-frameworks skeptic)

"Your three legs live in three different formalisms — a symmetry corollary, a Morse bound, a mass codebook. That is three sketches, not one proof."

Response. The heterogeneity is real and deliberate. The legs are not three proofs of one proposition; they are two bounds and a cross-check on one integer: leg (L1) forces $N_g \geq 3$ (Theorem 6.1), leg (L2) forces $N_g \leq 3$ (Theorem 7.18), and leg (L3) verifies that the surviving value is constructively realized, with the GG-A9-SM codebook reproducing the charged-lepton mass ratios at four significant figures (Theorem 8.2). The closure is stated once, in Theorem 9.1, with the conditionality of each leg carried along. And the heterogeneity buys falsifiability: the three legs die to different data — a CP-conserving world, a fourth basin, a failed mass ratio — so the conjunction is exposed on three independent fronts, not sheltered behind one.

Objection 5 (O-N3-5) (the modeling skeptic)

"Identifying particle generations with critical points of a projector landscape is a modeling choice, not physics."

Response. The identification is GDOS-internal, and the paper carries that fence in the statement of Theorem 7.18 itself. But within the framework the objects are not chosen ad hoc: once the boundary carries N_g families, the family-projector moduli space is forced to be CP^{N_g-1} (Section 6), and the GFF functional on it is the cohort's fixed free-energy object, inherited from (Arisaka, 2026, A6-GRIP), not tuned here. The identification also earns its keep dynamically:

the same cascade reading yields the decay-width ordering of the second and third generations and the reachability statement of §11.6, both of which match observation. A modeling choice that predicts the lifetime hierarchy it was not fitted to is the kind physics keeps.

Objection 6 (O-N3-6) (the already-excluded skeptic)

"LEP counted the neutrinos long ago. Deriving a number experiment has already settled adds nothing."

Response. The measurement (1) is settled physics, and the paper claims no discovery there. But exclusion is not explanation: the Standard Model accommodates any N_g — nothing in its structure prefers three — so the family count has stood for fifty years as a free discrete input. The derivation converts that input into a theorem with collateral obligations the measurement never imposed: the parity statement that N_g must be odd, the first-order mass ratios (predictions P3–P4), and the cascade lifetime hierarchy (P5). Those are net exposure, registered as kill criteria F1–F5 in Section 10, and they can fail tomorrow. A postdiction with new death conditions attached is a prediction about everything not yet measured.

Objection 7 (O-N3-7) (the numerology skeptic)

"$S_{\text{tot}}^{\text{max}} = 3\pi$ is numerology with a Greek letter in it."

Response. The answer is a ledger, and every line of it is checkable. What is derived: the step actions $S_k = \pi k$ follow from the cascade geometry (Section 9); the capacity's N_g-scaling follows from the Pontryagin identity $p_1 = N_g h^2$; and the unit $\lambda = \pi$ is pinned by three independent routes (Lemma 7.14) — lattice-plus-decisiveness, MSMF-minimality, and no-new-scales — none of which knows the answer. What is classified: among all natural single-class capacities, exactly two survive and both give $N_g = 3$ with exact saturation (Theorem 7.17); the old $\lfloor J/2 \rfloor$ reading is exposed as a numerical coincidence of $J = 7$ that the chain no longer touches. Numerology is a pattern selected after the fact from an unconstrained supply; here the supply is constrained to a single integer multiple of a derived unit, declared before the adjudicating mathematics is brought to bear — and that mathematics selects the declared value uniquely. That is the opposite of numerology; it is a pre-registration paying off.

Objection 8 (O-N3-8) (the axioms skeptic)

"Theorems 'within GDOS' just relocate the mystery into your axioms."

Response. We accept the criticism in the form stated: every theorem here is conditional on the declared axiom system, and a reader who rejects the constitution receives conditionals. Two replies. First, relocation with compression is what explanation is: the same axiom set that elsewhere carries the Standard-Model parameter sheet and the cosmological inventory here carries the family count — one axiom system, many former mysteries; the mystery count went down. Second, the axioms are not adjustable refuges: AX-1 through AX-4 are fixed cohort-wide, and this paper added none of its own — its two imported postulates are consistency-grade with declared provenance (Postulates 7.19 and 7.20; the import's would-be derivation from BI-6 is now closed in the negative, §7.11), its one axiom-level clause is stated at minimal size with every derivable part derived (Axiom 7.11), and its one residual check is named with its failure mode priced (R-INT, the topological residue of the integrality theorem L-INT). Nothing was smuggled in to make $N_g = 3$ come out — and the chain now consumes neither $K = 35$ nor $J = 7$, the two places a smuggler would hide.

Objection 9 (O-N3-9) (the methodology skeptic)

"Counting generations is kinematics. Why drag open-system machinery into it?"

Response. A closed-system derivation would indeed have been cleaner, and we tried to find one: four closed-system routes — a Chern–Simons multiplicity axiom, a magnetized-torus index, a $\mathbb{Z}_3$ orbifold, and partial anomaly saturation — were attempted and failed, for reasons recorded route by route in Section 14. The failure has a structural reading: N_g counts the stable readouts of the boundary runtime, and stability is an open-system notion — basins and gates, not representations. Closed-system counting delivers candidate multiplicities; it cannot say which candidates survive as registered families. That is why leg (L2) is a statement about the landscape the boundary actually runs, and why the methodological reflection is in the paper at all: the negative results are data about where the answer lives.

Objection 10 (O-N3-10) (the postdiction skeptic)

"The mass ratio 16.817 was measured decades before you 'predicted' it."

Response. Conceded without reservation: m_τ/m_μ is known physics, and the paper claims no postdictive credit. The registered exposure is forward-looking. The ratio enters as kill criterion F3: a precision value off the first-order prediction beyond the declared QD tolerance kills the codebook leg going forward, whatever the history of the measurement. The integers behind the ratios are fixed cohort-wide by the OGN-5 ledger $(1, 3, 5; 7, 35)$ — this paper could not adjust them to fit if it wanted to, which is what separates a derived ratio from a fitted one. And the paper's permanent exposure is sharper still: one fourth-generation event, at any energy, in any decade, fires F1–F2 and kills the conjunction outright. Absences are the most falsifiable claims there are.

Objection 11 (O-N3-11) (the anthropic skeptic)

"The landscape already explains $N_g = 3$: other values happen elsewhere, and observers see what observers can see."

Response. Section 11.4 gives the structural comparison without polemic. Anthropic reasoning accommodates any family count compatible with observers, at the price of a measure problem and an unobservable ensemble; it explains why we might find ourselves at $N_g = 3$ without forbidding anything nearby. The present route is a uniqueness claim with zero continuous parameters: $N_g = 3$ is forced, and a fourth generation — anywhere, ever — kills it. The two are therefore not rivals on one axis: one is a selection argument over an ensemble, the other a two-sidedly falsifiable derivation in a single universe. A reader may prefer either posture; only one of them can lose.

Objection 12 (O-N3-12) (the sociology skeptic)

"A single-author cohort, citing itself, is not independent verification."

Response. Fully conceded: self-citation is provenance bookkeeping, not evidence, and no referee should weigh it as support. What the paper offers for verification is deliberately machine-checkable without trusting the author: every inherited statement in the claim-status ledger carries a named source and a status tag that can be checked against the cited paper; the kill criteria F1–F5 are experimental statements owned by the community's instruments, not by the cohort; and the codebook cross-check of leg (L3) is arithmetic from five integers that any reader can repeat in an afternoon. The cohort structure exists to make that audit tractable — one derivation chain, consistently named across papers — not to substitute for it. Skepticism is the intended reading posture; the paper's job is to make the skepticism efficient.

Objection 13 (O-N3-13) (the axiomatized-answer skeptic)

"You proved the budget could not be derived — and then adopted it as an axiom. That is not deciding the family number; it is enshrining the failure."

Response. Stated at full force, and the paper's favorite objection, because answering it is what the architecture is for. Four replies, in descending order of weight. *First*, the necessity half is proved: the obstruction theorems show that no derivation from the prior axioms exists — so the alternatives were a named clause or a permanent conjecture, and naming is strictly more honest than conjecturing. *Second*, the clause does not mention the answer: Axiom 7.11 contains no N_g, no 3, no K, no J, and no budget value — the number three emerges from $\pi N_g(N_g - 1)/2 \leq \pi N_g$, which is arithmetic. *Third*, the clause's content is nearly all theorem: transport is Berry and the geometry is Fubini–Study at Born-grade (Appendix F); what the axiom actually adds is the pairing form, whose freedom is then closed by classification (Theorem 7.17). *Fourth*, the clause is falsifiable through its shared premises: kill criterion F6 ties it to the Born rule's own footing. What would vindicate the skeptic is a proof that the coupling clause cannot be derived from the joint bulk–boundary object either — the analogue, one level up, of what this paper proved about the prior axioms. The paper states the target rather than hiding it, and Appendix H is where that program was carried out: the residue R-INT is discharged on the benchmark, and the clause is exhibited as the output of a computation rather than assumed. What is *not* yet shown — and the objection is entitled to press it — is that no derivation from the joint object could have failed; enforcement on one background is not a necessity proof.

Objection 14 (O-N3-14) (the Born-coupling skeptic)

"Leg (L2) rests on the same premises as your Born theorem. You have made the family count hostage to your most speculative derivation."

Response. The coupling of fates is deliberate, and it runs in the skeptic's favor. The premises in question — the GG-A5-GDOS forward semigroup and the collision-level contraction — are the minimal content of "the boundary runtime is quantum mechanics with relaxation"; a reader who rejects them is not quibbling with leg (L2), they are rejecting the framework's quantum mechanics wholesale, Born rule and all, and the paper cannot and should not insulate the family count from that. What the coupling buys is economy of exposure: one premise set now carries the Born rule, the adiabatic transport, and the capacity geometry, so a single experimental or mathematical failure would break all three loudly — which is falsifiability working as designed (F6) — while a single success (the field-grade representation step, the declared residue of the Born chain) would harden all three at once. Hostage-taking would be tying the result to an *unrelated* speculation; tying it to the framework's own quantum mechanics is just honesty about what leg (L2) is: a statement about what quantum transport on the family moduli can pay for.

The inflow bridge

The seven objections that follow are addressed to Appendix H and to the enforcement claim of §7.7. They are kept together, and kept last, because a reader who takes the coupling clause as an axiom — as §7.6 states it — may stop at Objection 14 and lose nothing of the family-number argument.

Objection 15 (O-N3-15) (the model skeptic)

"Your 'channel model' is an assumption dressed as a definition. The whole theorem lives inside it."

Response. The objection is right that Definition H.1 is where the conditionality lives, and the theorem statements say so. Three replies. First, the channel itself is not modeled but computed: on the benchmark the Künneth structure admits exactly two family classes, and fiber degree kills one — there is nothing else the coupling could be, in cohomology (T-INF-1). Second, the product premise (P-i) and the pullback premise (P-ii) are the intended class's own clauses, imported from GG-A3-BI6, not invented here; the refinement they defer is closed at class level by Theorem T-INF-7, with the constructive warped representative the named survivor (E-3). Third, what escapes the model is cocycle-level (non-topological) family dependence; the paper's claim is calibrated to the class level throughout, which is where the coupling clause itself lives.

Objection 16 (O-N3-16) (the minimality skeptic)

"Strip the MSMF appeal and your bridge is a lattice statement, not an identity."

Response. Correct — and the paper prices exactly that. Without minimal quanta, T-INF-3 still holds unconditionally: the pushforward is always an integer multiple of the Berry bundle, with the Berry bundle the minimal non-zero value. The bridge identity then follows from the same two principles the unit classification of §7.8 already stakes (non-vanishing plus minimality), used in the same way. A reader who rejects MSMF here has already rejected it there; the conditional structure is shared, declared, and no worse at this link of the chain than at any other.

Objection 17 (O-N3-17) (the numerology skeptic)

"$2 \times 2/4 = 1$**. You have dressed a triviality in functors."**

Response. The arithmetic is trivial; the content is where each integer comes from and what forbids the alternatives. The 4 is the Hopkins–Singer normalization — not chosen. The evenness that excludes $(1, 4)$ and $(4, 1)$ is the Wu condition — not chosen. The first 2 is the trace-ratio numerator $L - N$, delivered by the abelian cross term of a lift fixed in GG-A3-BI6 — not chosen. The second 2 is the orientation double of a unit reservoir proved to exist there — not chosen. Four independent structures, none adjustable, meeting at 1: that is the opposite of numerology, and the harnesses verify each meeting point separately precisely so that no single "nice" coincidence carries the theorem.

Objection 18 (O-N3-18) (the circularity skeptic)

"You matched the boundary to fix the placement, then celebrate agreeing with the boundary."

Response. The objection would be fatal if the matching fixed a continuous parameter; it fixes a binary placement, and the non-trivial content survives the fixing. What the boundary decides is *which slot* of the quadratic machine carries the physical datum (Proposition H.10) — a choice the formalism offers and one side of the pair must resolve, as the boundary-relativity proposition demands. What the boundary does *not* supply is everything that then has to come out right: the channel rank, the $cm/4$ factor, the evenness, the value $L - N$, the doubling. Placement (B) was refutable and was refuted; a framework in which both placements matched would deserve the objection — this one does not.

Objection 19 (O-N3-19) (the selection skeptic)

"$|L - N| = 2$ at $N = 7$ too. Your bridge cannot tell three families from seven."

Response. Correct, stated in the paper's own honesty section (§H.11.5), and by design: this is a consistency theorem for the coupling clause, not a fourth leg of the pincer. The family number is decided by the capacity theorem of §7.8 (Theorem 7.16), which excludes $N = 7$ by the budget with room to spare. What the bridge adds at $N = 3$ is not selection but enforcement: the clause the derivation conditions on is exhibited as the output of a computation, on the one background where every ingredient is theorem-grade.

Objection 20 (O-N3-20) (the one-background skeptic)

"One background is not a theory. Everything here is $S^4 \times T^2$."

Response. Partly right, and the paper's propositions are scoped accordingly. What holds on the whole intended class: the absolute part of R-INT (Proposition H.12), the quantization theorem, the attributions, the echoes. What is benchmark-specific: the uniqueness of the Wu structure (elsewhere, a 2^{44}-torsor with a declared basepoint) and the numerical curvature check. The benchmark is not a toy: it is the background GG-A3-BI6's global anomaly theorem itself certifies, and the physical compactification of the intended geometry. Extending the uniform-choice statement across the class is the basepoint question that survives Theorem T-INF-8's closure of R-REL, stated as such in §H.11.2.

Objection 21 (O-N3-21) (the foundations skeptic)

"Your program advertises that GDOS solves the measurement problem; now your bridge leans on the Born rule. Which way does the arrow point?"

Response. Both statements are true because they live at different links of one chain, and the paper devotes §14 to writing the chain once. Downward: the axioms give GG-A5-GDOS's open-system law; GG-A8-QM derives the Born measure from it (conditionally, at its declared grade); this paper's adiabatic theorems ride on that result; and the bridge's boundary input rides on those. So yes — the Born rule is derived inside the architecture, and yes — this appendix consumes the derived result rather than re-deriving it. The arrow points one way, and this appendix sits at its downstream end. What would be incoherent is the claim the objection warns against — that the bridge *founds* quantum mechanics — and the paper declines it explicitly: its boundary input is built from the Born-grade package, and a construction fed the boundary's answer cannot be the answer's source. What it contributes instead is the two-sided agreement of §11.3, which is precisely the kind of evidence a derivation cannot supply.

Part VIII
References

Cited works in two groups — the GG-Theory cohort and the external literature — each entry carrying a relevance note and its first citation point

The reference list is organized in two sections, in line with the GG-Theory house style. **Section A** lists the Arisaka cohort entries cited in this paper — ten shipped sister papers, in the canonical order of the program, followed by the two declared in-preparation companions — and **Section B** the external literature. Every entry carries a **Relevance** paragraph stating what the work contributes to this paper and where it enters, followed by the location of its first citation; the list contains no entry that is not cited in the text. All shipped cohort preprints are available at https://preprints.arisaka-gg.org/.

A. Papers by Katsushi Arisaka (the GG-Theory cohort)

Arisaka, K. (2026). The GG-Constitution: Four Axioms and Seven Postulates of the GG-Theory Program, From Which GG-6, BI-6, OGN-5, and GDOS Descend as Theorems. Arisaka-GG-A1-Constitution.
Preprint, available at https://preprints.arisaka-gg.org/a1-constitution/.
Relevance. The constitutional charter of GG-Theory: states the four axioms (LC, GGEP, OGGEP, and MSMF — "Maximum Symmetry, Minimum Freedom", the canonical expansion used in Postulate 3.1) and the seven postulates. The ledger-minimality corollary of MSMF consumed by leg (L1) is grounded here. First cited in §3.

Arisaka, K. (2026). GG-Theory with GG-6: Geometry–Gauge Locking at the 6D Bulk (GG-6) at the Heart of Grand Geometric Theory (GG-Theory). Arisaka-GG-A2-GG6.
Preprint, available at https://preprints.arisaka-gg.org/a2-gg6/.
Relevance. The bulk paper: establishes the six-dimensional geometry, the canonical metric used in §E.5, the role of $\mathrm{Spin}(1,5)$, and — decisively for the present paper — the Diophantine uniqueness of the OGN-5 ledger $(1,3,5;7,35)$ (its Theorem 6), of which the generation count is here proved to be the physical realization. First cited in §1.

Arisaka, K. (2026). Anomaly Freedom of BI-6: The Modified Bianchi Identity (BI-6) for the Six-Force, Six-Dimensional Bulk (GG-6). Arisaka-GG-A3-BI6.
Preprint, available at https://preprints.arisaka-gg.org/a3-bi6/.
Relevance. The anomaly paper: proves BI-6 closure, the $-2/5$ trace-ratio coefficients, the integrality of $[X_4]$, and the dual-lock identity $K = LJ = M^2 + N^2 + L^2 = 35$ consumed by leg (L2). Its N_g-free block-size theorem T-BI6-9.0 and its γ-family characterization of the Abelian

coefficient are consumed by §6.2 and Appendix C: the former removes the last circularity from the extensive capacity's flux reading, the latter closes the coefficient route to N_g. First cited in §1.

Arisaka, K. (2026). GDOS: Geometric Dynamics of Open Systems — 6D GG-6 Bulk Projected to 4D Boundary for Quantum Measurement, General Relativity, and QFT. Arisaka-GG-A5-GDOS. Preprint, available at https://preprints.arisaka-gg.org/a5-gdos/.
Relevance. Develops GDOS as the open-system framework underlying boundary observation, on which leg (L2) is built. Its open problem O3 records the GFF Morse-positivity question as unsettled in any companion paper — the present paper adopts it as Postulate 7.19 and reduces its strict form in Appendix E. First cited in §3.

Arisaka, K. (2026). The GRIP on GDOS — GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind. Arisaka-GG-A6-GRIP.
Preprint, available at https://preprints.arisaka-gg.org/a6-grip/.
Relevance. Develops the GFF and GRIP pillars used in leg (L2); its cascade-closure theorem T4.1 consumes the same capacity value now delivered by Theorem 7.16, and its open problem M10 — derive the budget independent of $N_g = 3$ at every intermediate step — is met by the present architecture, which consumes neither $K = 35$ nor $J = 7$ (§8.2). The present paper supplies the dedicated three-leg treatment of $N_g = 3$ that the cohort reserves. First cited in the Executive Summary.

Arisaka, K. (2026). Quantum Mechanics Derived from GG-Theory as Projection from 6D Bulk to 4D Boundary. Arisaka-GG-A8-QM.
Preprint, available at https://preprints.arisaka-gg.org/a8-qm/.
Relevance. The boundary-quantum-mechanics companion: the Constitutional Born Theorem with its named premises — the GG-A5-GDOS forward semigroup (P1) and the collision-level contraction rate — which are exactly the premises of Theorems 7.8–7.9 (Remark 7.10), so that leg (L2)'s transport geometry is as secure as the Born rule and no more conditional. Also the pedagogical home of the Berry-phase worked example that Theorem 7.9 elevates to a statement about the runtime's own family transport. First cited in §6.

Arisaka, K. (2026). The Standard Model as Theorem of GG-Theory: Its 28 Parameters as Observations of One Six-Dimensional Geometry. Arisaka-GG-A9-SM.
Preprint, available at https://preprints.arisaka-gg.org/a9-sm/.
Relevance. Derives the 28 Standard-Model parameters from BI-6 and OGN-5; provides the explicit fermion-mass codebook (lepton and quark Yukawas, QD-completion factors, and the Higgs vacuum $v^{(1)} = 246.331354$ GeV) consumed by leg (L3). First cited in the Executive Summary.

Arisaka, K. (2026). The GG-6 Lock Cluster: A Non-Spatial Kaluza–Klein Resonance near 3.85 TeV as the Standing LHC Falsifier of GG-Theory. Arisaka-GG-P1-KK.
Preprint, available at https://preprints.arisaka-gg.org/p1-kk/.
Relevance. The physics of the structural torus through which — by T-INF-1 — the family's only coupling channel runs; the Wilson-line/holonomy sector that realizes the mixed class. First cited in Appendix H.

Arisaka, K. (2026). The Electron as the Ground State of Nature: The Six Oldest Questions about the Electron, Closed and Fenced within GG-Theory. Arisaka-GG-P3-Electron.
Preprint, available at https://preprints.arisaka-gg.org/p3-electron/.
Relevance. The sister paper of the Predictions Branch: develops the ground-state ontology of §9 into the dedicated six-question treatment of the electron (identity, stability, copies, the positron as conjugate readout, pointlikeness, sensing), importing the present paper's $N_g = 3$ theorem and its declared conditioning (the coupling clause and L-INT) verbatim, and adding the geometric pointlikeness theorem ($F \equiv 1$ as its laboratory corollary) via a readout-locality lemma derived from AX-1 of A1, with field response derived as composite geodesic motion on the AX-2 connection (per A5's Part VI) under the two-sector reading: the phase sector moves the electron, the scale sector transmutes it, and radiation is motion's dissipative channel. First cited in §4.

Arisaka, K. (2026). The Real Geometry Beneath the Complex Numbers of Physics: Positive Frequency is Holomorphy — from Cardano's Cubics to Twistor Space, via the Sixth Dimension of GG-Theory. Arisaka-GG-P4-Real.
Preprint, available at https://preprints.arisaka-gg.org/p4-real/.
Relevance. The origin-of-i member paper of the Foundations Branch: derives the boundary's imaginary unit from two real structures of the GG-A2-GG6 arena — the compact phase circle and the causal clock — the complex structure that the unitary mixing matrices of leg (L1) presuppose at boundary level. Cited at roster level in §4.1. First cited in §4.

Arisaka, K. (2026). The Mirrors of Matter: Why Every Particle Has a Twin, Why Nature's Mirrors Break, and Why the Universe Is Made of Matter. Arisaka-GG-P5-CPT.
Preprint, available at https://preprints.arisaka-gg.org/p5-cpt/.
Relevance. The CPT member paper of the Foundations Branch, and the Foundations member nearest to the present paper: C, P, and T as orientation reversals of the three real bulk factors, with a flavor-holonomy calculus that runs on the three-generation structure derived here — for $N_g = 3$ its chiral-pentagon reading ties the CP-odd quartets of the mixing matrix to the single surviving Jarlskog phase, the same empirical non-vanishing that leg (L1) consumes. Cited in §4.1. First cited in §4.

Arisaka, K. (2026). The Kaxion: Cold Dark Matter as the Non-Spatial Compact-Angle Pseudoscalar of GG-6. Arisaka-GG-C1-Kaxion.
Preprint, available at https://preprints.arisaka-gg.org/c1-kaxion/.
Relevance. Cited in Section 14 as a further example of an open-system structural prediction (a $\sim \mu$eV axion-like dark matter candidate) that closed-system frameworks cannot derive. First cited in §14.

Arisaka, K. (2026). GG-Theory: CKM, PMNS, and CP violation from BI-6. (In preparation.)
Relevance. Provides the predicted value $J_{\mathrm{CKM}}^{\mathrm{pred}} \approx 3.06 \times 10^{-5}$, removing any residual circularity in (L1). First cited in §5.

Arisaka, K. (2026). From Euclid to GG-6: The Realization of Euclid's Elements in Six Dimensions — GU-5, OGN-5, and a Parameter-Free Reconstruction of Modern Physics. Arisaka-GG-A4-Euclid.

Preprint, available at https://preprints.arisaka-gg.org/a4-euclid/.
Relevance. The parameter audit that this paper's claim has to survive. A4 records every dimensional constant of the framework as a conversion factor and every dimensionless integer as an output of a Diophantine condition, and N_g is one of the integers on its list. Cited because a derivation of the family number is only worth what the audit says it is worth: if N_g were tunable anywhere in the chain, A4 is where that would show.

Arisaka, K. (2026). The Qi-4 Pillars — Qi-QUA, Qi-EP, Qi-CON, Qi-HAL: The Quantum-Information Backbone of GDOS. Arisaka-GG-A7-Qi4.
Preprint, available at https://preprints.arisaka-gg.org/a7-qi4/.
Relevance. The quantum-information pillars, and the runtime vocabulary in which leg L2 is stated. The capacity argument here is a statement about what a boundary runtime can fund, and the terms it uses — relaxation, gap, descent — are A7's rather than this paper's. Cited for the vocabulary; the capacity theorem itself is proved here.

Arisaka, K. (2026). The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry. Arisaka-GG-A10-Cosmos.
Preprint, available at https://preprints.arisaka-gg.org/a10-cosmos/.
Relevance. The cosmological reading of the same geometry, and an independent consumer of the family number: the light-neutrino count that primordial nucleosynthesis and the microwave background constrain is the same N_g this paper derives. The agreement is a consistency check across two very different measurements and is worth naming as one. No step of the derivation here depends on A10.

Arisaka, K. (2026). The Geometric Inevitability of Cosmos, Life, and Mind: The Self-Referential Closure of the GG-Theory Program. Arisaka-GG-A11-Origin.
Preprint, available at https://preprints.arisaka-gg.org/a11-origin/. No result of the present paper depends on it.
Relevance. The capstone, and the one companion whose arrow points the other way: A11 consumes this paper's result rather than supplying anything to it. Its identity cascade runs from quarks upward and takes the family count as fixed, so a failure here propagates into it and not the reverse. Cited so that the dependency is visible from this end as well as from A11's.

Arisaka, K. (2026). The Hubble Tension in the Light of GG-Theory: The Hubble Constant as One Reading of a Derived Expansion History $H(t)$. Arisaka-GG-C2-Hubble.
Preprint, available at https://preprints.arisaka-gg.org/c2-hubble/.
Relevance. A phenomenology paper on the astrophysical side of the same geometry. It carries no result this paper uses and this paper carries none it uses. Cited at roster level in the cohort map, and the note says roster level rather than manufacturing a connection.

Arisaka, K. (2026). Black-Hole Entropy under GG-Theory: The GDOS Reading of the Horizon as an OGGEP Port — with the Exact Bekenstein–Hawking $A/4G_N$ Derived from Induced Gravity. Arisaka-GG-C3-BH.
Preprint, available at https://preprints.arisaka-gg.org/c3-bh/.

Relevance. The second phenomenology paper, and the other end of the open-system reading whose boundary law leg L2 leans on. What the two share is GDOS: C3 asks what the projection does at a horizon, this paper asks what it funds on the family moduli. Nothing here turns on the horizon result. Cited at roster level.

Arisaka, K. (2026). The Proton as the Ground Geometric Attractor of Matter: The Eleven Questions of the Proton, Closed and Fenced within GG-Theory. Arisaka-GG-P6-Proton.
Preprint, available at https://preprints.arisaka-gg.org/p6-proton/.
Relevance. The other Predictions-Branch paper whose whole content is an integer statement holding exactly. P6 argues that baryon number is gauged and its anomaly canceled, so the proton never decays; this paper argues that the same ledger admits exactly three families. Both are all-or-nothing claims about integers, and both would be retired by a single contrary measurement. Cited at roster level in the cohort map.

Arisaka, K. (2026). Symmetry, Conservation, and Geometry: Seven Steps from Noether, Bianchi, to Hopkins–Singer, and the Ledger the Open-System Turn Pays Out. Arisaka-GG-W1-Symmetry.
Preprint, available at https://preprints.arisaka-gg.org/w1-symmetry/.
Relevance. The survey of the century this paper's method belongs to. The upper bound in leg L2 is not a symmetry argument — it is a capacity argument, and it reaches a conservation-style conclusion without a conserved current to hang it on. W1 is the cohort's account of exactly that shift. It supplies no number and no theorem used here. Cited at roster level.

Arisaka, K. (2026). The Circle That Remembers, the Interval That Forgets: GG-Theory and the Sixteen Problems of Particle Physics and Cosmology. Arisaka-GG-W2-Holonomy.
Preprint, available at https://preprints.arisaka-gg.org/w2-holonomy/.
Relevance. The classification the lattice argument belongs to. The unit that makes leg L2 decisive is a spin-refined quantization condition — protection by topology rather than by symmetry — and that is W2's subject. Cited at roster level in the cohort map; the lattice derivation itself is complete in Part VI and borrows nothing from it.

B. Other references

ALEPH, DELPHI, L3, OPAL Collaborations & LEP Electroweak Working Group.
Precision electroweak measurements on the Z resonance. *Physics Reports* **427** (2006) 257–454.
Relevance. Establishes the empirical bound on the number of light neutrino species, which is the most direct experimental confirmation of $N_\nu = 3$. Underpins prediction P1 of this paper. First cited in §1.

Blumenhagen, R., Körs, B., Lüst, D., Stieberger, S. Four-dimensional string compactifications with D-branes, orientifolds and fluxes. *Physics Reports* **445** (2007) 1–193.
Relevance. Comprehensive review of magnetized toroidal compactifications; provides the index-theorem family-counting argument that the present paper compares against. First cited in §1.

Candelas, P., Horowitz, G. T., Strominger, A., & Witten, E. Vacuum configurations for superstrings. *Nuclear Physics B* **258** (1985) 46–74.

Relevance. The seminal Calabi–Yau compactification paper; the $|\chi/2| = 3$ family-counting condition compared against in Section 11. First cited in §1.

Gorini, V., Kossakowski, A., & Sudarshan, E. C. G. Completely positive dynamical semigroups of N-level systems. *Journal of Mathematical Physics* **17** (1976) 821–825.

Relevance. The GKLS theorem characterizing the most general completely positive trace-preserving (CPTP) generator of open-system quantum dynamics. Together with Lindblad (1976), this is the mathematical backbone of open-system quantum theory and the prototype for the CPTP boundary projection that GDOS elevates to foundational status. First cited in §14.

Green, M. B., & Schwarz, J. H. Anomaly cancellations in supersymmetric $D = 10$ gauge theory and superstring theory. *Physics Letters B* **149** (1984) 117–122.

Relevance. The original modified Bianchi identity construction; the six-dimensional descent of which is BI-6 of the present framework. First cited in §3.

Green, M. B., Schwarz, J. H., & Witten, E. *Superstring Theory, Vols. 1–2.* Cambridge University Press (1987).

Relevance. Standard reference for superstring compactifications and anomaly cancellation; provides the broader context for the BI-6 construction. First cited in §1.

Born, M. Zur Quantenmechanik der Stoßvorgänge. *Zeitschrift für Physik* **37** (1926) 863–867.

Relevance. The Born rule that translates closed-system unitary evolution into open-system probabilistic outcomes. The Born rule is the historical seed of the open-system perspective; cited in Section 14 as one of the foundational moments where open-system aspects entered quantum mechanics without being elevated to foundational status. First cited in §14.

Freed, D. S., & Hopkins, M. J. On Ramond–Ramond fields and K-theory. *Journal of High Energy Physics* **2000**(05) (2000) 044.

Relevance. Foundational paper on the differential-cohomology classification of higher-form gauge fields and anomaly inflow. Provides the technical machinery underlying the non-degeneracy of the Hopkins–Singer inflow map (hypothesis H_2 in Appendix E). First cited in Appendix E.

Helgason, S. *Differential Geometry, Lie Groups, and Symmetric Spaces.* Academic Press (1979).

Relevance. Standard reference for the Fubini–Study metric on CP^{N_g-1} and its geodesic structure, used in Lemma 7.4. First cited in §6.

Hirzebruch, F. *Topological Methods in Algebraic Geometry* (3rd ed.). Springer-Verlag, Berlin (1966).

Relevance. The foundational reference for the Hirzebruch signature theorem and the L-polynomial computation, used throughout Step 2 of Appendix E. First cited in Appendix E.

Hirzebruch, F. *Topological Methods in Algebraic Geometry* (Classics in Mathematics ed.). Springer-Verlag, Berlin (1995).

Relevance. Reissue of the classic monograph; provides the explicit Pontryagin and L-class computations on CP^{N_g-1} used in Steps 2–4 of Appendix E. First cited in Appendix E.

Hopkins, M. J., & Singer, I. M. Quadratic functions in geometry, topology, and M-theory. *Journal of Differential Geometry* **70** (2005) 329–452.
Relevance. The differential cohomology framework underlying BI-6 closure, and the quadratic refinement behind component (C) of the integrality theorem L-INT (Theorem 7.13) and its residue R-INT. First cited in §6.

Kato, T. On the adiabatic theorem of quantum mechanics. *Journal of the Physical Society of Japan* **5** (1950) 435–439.
Relevance. The adiabatic theorem in its projector form, and the transport construction $W' = [P', P]W$ used verbatim in Theorem 7.7. The three-line intertwining proof in §6 is Kato's. First cited in §6.

Berry, M. V. Quantal phase factors accompanying adiabatic changes. *Proceedings of the Royal Society of London A* **392** (1984) 45–57.
Relevance. The geometric phase of adiabatic transport and the connection on the tautological bundle that Theorem 7.9 identifies with the runtime's family transport. First cited in §6.

Provost, J. P., & Vallee, G. Riemannian structure on manifolds of quantum states. *Communications in Mathematical Physics* **76** (1980) 289–301.
Relevance. The quantum geometric tensor: the Fubini–Study metric as the fidelity metric on state manifolds, with the Berry curvature as its antisymmetric partner — the identification used in Theorem 7.9(c) and proved in Appendix F. First cited in §6.

Avron, J. E., Fraas, M., Graf, G. M., & Grech, P. Adiabatic theorems for generators of contracting evolutions. *Communications in Mathematical Physics* **314** (2012) 163–191.
Relevance. The adiabatic theorem for gapped generators of contraction semigroups — the exact setting of the GDOS runtime (CPTP/GKLS by GG-A5-GDOS T-GDOS-3–T-GDOS-4) — used in the proof of Theorem 7.8 in Appendix F. First cited in Appendix F.

Kobayashi, S., & Nomizu, K. *Foundations of Differential Geometry, Vol. II.* Wiley-Interscience, New York (1969).
Relevance. Standard reference for Chern–Weil theory and the identification of $\operatorname{tr} R \wedge R$ with the first Pontryagin class on $\mathrm{SO}(2k)$ frame bundles, used in Lemma E.4 of Appendix E. First cited in Appendix E.

Ibáñez, L. E., Nilles, H. P., & Quevedo, F. Orbifolds and Wilson lines. *Physics Letters B* **187** (1987) 25–32.
Relevance. The $\mathbb{Z}_3$-orbifold construction with three twisted sectors; the most natural prior model giving exactly three generations, compared against in Section 11. First cited in §1.

Ibáñez, L. E., & Uranga, A. M. *String Theory and Particle Physics: An Introduction to String Phenomenology.* Cambridge University Press (2012).

Relevance. Standard reference for D-brane and intersecting-brane family counting; provides the broader phenomenological background. First cited in §1.

Dobrescu, B. A., & Poppitz, E. Number of fermion generations derived from anomaly cancellation. *Physical Review Letters* **87** (2001) 031801.
Relevance. Proves that global anomaly cancellation forces exactly three generations for the Standard Model in two universal extra dimensions — a genuine derivation of both bounds, conditional on the extra-dimensional assumption. Engaged in Sections 1, 11, and 14. First cited in §1.

Lindblad, G. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics* **48** (1976) 119–130.
Relevance. The Lindblad master equation, derived independently of Gorini, Kossakowski & Sudarshan (1976), gives the most general CPTP generator of Markovian open-system quantum dynamics. Cited in Section 14 as the foundational technical backbone of modern open-system quantum theory. First cited in §14.

Maldacena, J. The large N_g limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics* **2** (1997) 231–252.
Relevance. The original AdS/CFT correspondence: the first explicit example of holographic bulk-to-boundary descent. Cited in Section 14 as the prototype for the bulk-to-boundary projection that GDOS elevates to foundational status. First cited in §14.

Kobayashi, M., & Maskawa, T. CP-violation in the renormalizable theory of weak interaction. *Progress of Theoretical Physics* **49** (1973) 652–657.
Relevance. The original Kobayashi–Maskawa observation that $N_g \geq 3$ is required for CP violation in the quark sector; the historical precursor to (L1) of the present paper. First cited in §1.

Kikuchi, S., Kobayashi, T., & Uchida, H. Modular flavor symmetries of three-generation modes on magnetized toroidal orbifolds. *Physical Review D* **104** (2021) 065008.
Relevance. A modern modular-flavor construction realizing three generations conditionally on chosen orbifold and flux inputs; representative of the N_g-as-input class compared against in Section 11. First cited in §1.

Milnor, J. *Morse Theory.* Annals of Mathematics Studies 51, Princeton University Press (1963).
Relevance. The foundational reference for Morse theory; the height function on CP^{N_g-1} as a perfect Morse function is the textbook example that drives Lemma 7.2. First cited in §6.

NOvA Collaboration. Improved measurement of neutrino oscillation parameters by the NOvA experiment. *Physical Review D* **106** (2022) 032004.
Relevance. Provides the empirical evidence for non-trivial PMNS mixing and a non-zero leptonic CP phase, supporting the parallel application of (L1) in the leptonic sector. First cited in §5.

Particle Data Group. Review of Particle Physics. *Progress of Theoretical and Experimental Physics* **2024** (2024) 083C01.

Relevance. The standard reference for measured Standard-Model parameters, including all charged-lepton and quark masses against which the GG-A9-SM codebook of (L3) is compared. First cited in §1.

Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. *Astronomy & Astrophysics* **641** (2020) A6.
Relevance. Provides the cosmological measurement $N_{\text{eff}} = 2.99 \pm 0.17$, an independent confirmation of $N_\nu = 3$ complementing the LEP electroweak result. First cited in §1.

Rokhlin, V. A. New results in the theory of four-dimensional manifolds. *Doklady Akademii Nauk SSSR (N.S.)* **84** (1952) 221–224.
Relevance. The divisibility theorem — the signature of a closed spin 4-manifold is divisible by 16 — behind the spin reading of the parity route: a spin stratum could not host the exact saturation $3\pi = 3\pi$ (it would force $\langle p_1 \rangle \in 48\mathbb{Z} \not\ni 3$), so the non-spin-ness of CP^2 is load-bearing (Appendix E, §E.8). First cited in Appendix E.

Susskind, L. The world as a hologram. *Journal of Mathematical Physics* **36** (1995) 6377–6396.
Relevance. The holographic principle in its string-theory formulation: the suggestion that bulk degrees of freedom project onto a lower-dimensional boundary. Cited in Section 14 as one of the foundational steps toward the open-system / boundary-projection perspective elevated to foundational status by GDOS. First cited in §14.

't Hooft, G. Dimensional reduction in quantum gravity. arXiv:gr-qc/9310026 (1993). Reprinted in *Salamfestschrift* (World Scientific, 1993).
Relevance. The original holographic-principle proposal: that the fundamental degrees of freedom of a $(d+1)$-dimensional gravitational theory project onto a d-dimensional boundary. Cited in Section 14 as the seminal articulation of the holographic boundary-projection idea. First cited in §14.

von Neumann, J. *Mathematische Grundlagen der Quantenmechanik.* Springer-Verlag, Berlin (1932). English translation (1955), *Mathematical Foundations of Quantum Mechanics*, Princeton University Press.
Relevance. The foundational mathematical formulation of quantum mechanics as unitary evolution on a Hilbert space, together with the projection postulate for measurement. Cited in Section 14 as the moment when closed-system QM was formalized; the open-system aspects of measurement were already present but were not elevated to foundational status. First cited in §14.

Zeh, H. D. On the interpretation of measurement in quantum theory. *Foundations of Physics* **1** (1970) 69–76.
Relevance. The first articulation of decoherence: the mechanism by which open-system entanglement with the environment produces effective classicality from quantum dynamics. Cited in Section 14 as a foundational step toward recognizing the open-system character of quantum measurement. First cited in §14.

Zurek, W. H. Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse? *Physical Review D* **24** (1981) 1516–1525.
Relevance. The pointer-basis selection mechanism in decoherence theory: open-system dynamics privileges a particular basis ("pointer states") for the measurement outcomes. Together with Zeh (1970), provides the technical machinery underlying the open-system measurement perspective discussed in Section 14. First cited in §14.

Schwarz, M. *Morse Homology.* Birkhäuser Progress in Mathematics 111 (1993).
Relevance. Modern reference for transversality of gradient flows between Morse critical points; underlies Lemma 7.3 on saddle decay paths in CP^{N_g-1}. First cited in §6.

T2K Collaboration. Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations. *Nature* **580** (2020) 339–344.
Relevance. Empirical evidence of leptonic CP violation; provides independent support for (L1) in the leptonic sector. First cited in §5.

Witten, E. On flux quantization in M-theory and the effective action. *Journal of Geometry and Physics* **22** (1996) 1–13.
Relevance. The flux-quantization argument whose spin-refined level lattice drives component (C) of the integrality theorem L-INT (Appendix G, §G.5, step C4); the prototype setting for the shifted quantization on a spin bulk with a non-spin stratum. First cited in §6.

Witten, E. Supersymmetry and Morse theory. *Journal of Differential Geometry* **17** (1982) 661–692.
Relevance. The foundational paper relating supersymmetric quantum mechanics to Morse theory; provides the instanton-action computation $S = \pi k$ used in Lemma 7.4. First cited in §6.

Witten, E. Global anomalies in string theory. In: *Symposium on anomalies, geometry, topology* (Argonne, IL, 1985), 61–99.
Relevance. Reviews global-anomaly constraints and their relation to family number; provides one of the precursors to the BI-6 anomaly polynomial. First cited in §1.

Wan, Z., Wang, J., & Yau, S.-T. Three Families and Pontryagin Class. Preprint, arXiv (submitted January 29, 2026).
Relevance. Bordism-classified global-anomaly theorem: with N_c, N_f minimal and N_c odd (fermionic baryons), $N_c = N_f = 3$ is unique; absent those conditions, $N_f \geq 3$ with no upper bound. The most rigorous family-puzzle result internal to the 4D Standard Model; complementary to the present upper bound (Section 11). First cited in §11.

Chamseddine, A. H., Connes, A., & Marcolli, M. Gravity and the standard model with neutrino mixing. *Advances in Theoretical and Mathematical Physics* **11** (2007) 991–1089.
Relevance. The canonical spectral-triple reconstruction of the Standard Model in noncommutative geometry, in which the generation number enters as the input multiplicity $\mathcal{H}_F = M_F^{\oplus 3}$; compared against in Section 11 as the leading N_g-as-input framework. First cited in §11.

Gourlay, L., & Gresnigt, N. Algebraic realization of three fermion generations with S_3 family and unbroken gauge symmetry from $\mathbb{C}\ell(8)$. *European Physical Journal C* **84** (2024) 1129.

Relevance. Division-algebra/Clifford-algebra route to three generations via the S_3 automorphisms of the sedenions embedded in $\mathbb{C}\ell(8)$; the algebra choice is the input. Compared against in Section 11. First cited in §11.

Ibe, M., Kusenko, A., & Yanagida, T. T. Why three generations? *Physics Letters B* **758** (2016) 365–369.
Relevance. The modern anthropic argument: baryogenesis plus dark matter motivate $N_g \geq 3$ heuristically, with no upper bound. Engaged in Section 11. First cited in §11.

Esteban, I., González-García, M. C., Maltoni, M., Schwetz, T., & Zhou, A. The fate of hints: updated global analysis of three-flavor neutrino oscillations. *Journal of High Energy Physics* **2020**(09) (2020) 178.
Relevance. Global-fit source (NuFIT) for the leptonic Jarlskog invariant best-fit value quoted in Section 5.2. First cited in §5.

Bianchi, L. Sui simboli a quattro indici e sulla curvatura di Riemann. *Rendiconti della Accademia Nazionale dei Lincei* **11** (1902) 3–7.
Relevance. The original identity whose century of modifications — Green–Schwarz's above all — is the lineage of BI-6; the point of departure of §2's external road. First cited in Appendix H.

Dirac, P. A. M. Quantised singularities in the electromagnetic field. *Proceedings of the Royal Society of London A* **133** (1931) 60–72.
Relevance. The first topological integer of physics: a phase that must close forces an integer — the ur-argument that reappears in the Berry bundle's quantization, this paper's unit lattice, and the Wu-evenness of T-INF-3. First cited in Appendix H.

Cheeger, J., & Simons, J. Differential characters and geometric invariants. In *Geometry and Topology*, Springer LNM 1167 (1985) 50–80.
Relevance. The origin of the differential-character formalism in which the seam's holonomy and torsion data live; the ancestor of the Hopkins–Singer refinement used throughout. First cited in Appendix H.

Dai, X., & Freed, D. S. η-invariants and determinant lines. *Journal of Mathematical Physics* **35** (1994) 5155–5194.
Relevance. The collar formalism of the intended class's global-anomaly ladder, and the frame in which the relative closure theorem T-INF-8 is packaged (engineering item E-4). First cited in Appendix H.

Witten, E. Five-brane effective action in M-theory. *Journal of Geometry and Physics* **22** (1997) 103–133; arXiv:hep-th/9610234.
Relevance. The fivebrane partition function — the physical origin of the $(k, i) = (2, 2)$ pattern by which a six-dimensional seam pushes forward to a line bundle on a family base, the very type of the bridge. First cited in Appendix H.

Witten, E. On flux quantization in M-theory and the effective action. *Journal of Geometry and Physics* **22** (1997) 1–13; arXiv:hep-th/9609122.

Relevance. The prototype of every spin-shifted integrality statement in the architecture — the lattice halving whose Wu-condition form drives T-INF-3. First cited in Appendix H.

Freed, D. S. Dirac charge quantization and generalized differential cohomology. arXiv:hep-th/0011220 (with an appendix by M. J. Hopkins).
Relevance. Charge quantization of abelian gauge fields as cocycles in generalized differential cohomology — the doctrinal frame in which the BI-6 data are differential cocycles at all. First cited in Appendix H.

Milnor, J., & Stasheff, J. *Characteristic Classes.* Annals of Mathematics Studies 76, Princeton University Press (1974).
Relevance. The Wu formula and the dimensional vanishing $\nu_j = 0$ for $2j > \dim$, which make degree-four Wu structures canonical in six dimensions and set up the evenness that drives T-INF-3. First cited in Appendix H.

Monnier, S. Topological field theories on manifolds with Wu structures. *Reviews in Mathematical Physics* **29** (2017) 1750015; arXiv:1607.01396.
Relevance. The invertible theories carrying degree-four Wu data whose boundary is the six-dimensional self-dual field — the closest external sibling of the seam structure quantized here. First cited in Appendix H.

Sati, H., & Schreiber, U. Flux quantization. *Encyclopedia of Mathematical Physics*, 2nd ed.; arXiv:2402.18473.
Relevance. The modern axiomatization of which cohomology theory quantizes a given higher Bianchi identity; the frame in which BI-6's Green–Schwarz structure is read as a twisted differential cocycle. First cited in Appendix H.