

A Falsifiable Topological Correspondence for the Standard Model Parameters from a G_2 -Holonomy Manifold

Brieuc de La Fournière

brieuc@bdelaf.com

Article

Keywords: G2 holonomy, exceptional Lie algebras, Standard Model parameters, falsifiability, formal verification, foundations of physics, structural realism

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11059128/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

A Falsifiable Topological Correspondence for the Standard Model Parameters from a G_2 -Holonomy Manifold

Brieuc de La Fournière
Independent researcher, Beaune, France
brieuc@bdelaf.com June 2026

Abstract

The Standard Model requires nineteen experimentally determined parameters that it does not explain. We investigate, in an exploratory and explicitly falsifiable spirit, whether these dimensionless parameters can be expressed as topological invariants of a *conjectural* compact seven-dimensional G_2 -holonomy manifold K_7 with Betti numbers $(b_2, b_3) = (21, 77)$, coupled to an $E_8 \times E_8$ structure, with no continuously adjustable physical parameter. We do not exhibit this manifold: its topological and lattice data are certified, while the full smooth analytic compact construction remains open, and every claim below is conditional on this clearly delimited hypothesis. The framework's load-bearing claim is a *parameter-free core* of exact rational relations among topological integers: for instance $\sin^2 \theta_W = b_2/(b_3 + \dim G_2) = 3/13$, the Koide relation $Q = \dim(G_2)/b_2 = 2/3$, $m_s/m_d = 20$, and $\Omega_{\text{DM}}/\Omega_b = 43/8$, which cannot be tuned and are individually correct or wrong; all are machine-checked in Lean 4 (zero **sorry**). The geometric substrate is independently anchored by an interval-rigorous, machine-checked bound on the Calabi–Yau residual of an explicit closed-form metric on the K3 fibre (a variance $\text{Var}(\log R) \leq 1.32 \times 10^{-4}$, certified box-locally over the surface, not merely at sample points). Against a deliberately generous null model that draws random algebraic formulas from the *same* constants, no set among three million random formula sets reproduces the framework's joint precision profile, indicating strong over-determination rather than curve-fitting. The decisive falsifier is the leptonic CP phase $\delta_{CP} = 7 \times 14 + 99 = 197^\circ$, compatible with the spread of recent global fits (which have ranged over 177° – 212° between releases) though no longer at the central value of the latest, and to be tested definitively by DUNE; a fourth fermion generation, or a measured δ_{CP} outside $[182^\circ, 212^\circ]$ at 3σ , would falsify the framework. We claim no correctness: only a sharp, testable correspondence, and we are explicit about where geometric derivation ends and empirical correspondence begins.

Keywords: G_2 holonomy · exceptional Lie algebras · Standard Model parameters · falsifiability · formal verification · foundations of physics · structural realism

Epistemic status at a glance

Component	Status
Parameter-free core (33 exact algebraic relations)	Machine-checked in Lean 4 (algebraic identities)
K_7 Betti/lattice data, $(b_2, b_3) = (21, 77)$	Topological/lattice route certified (Joyce–Karigiannis $\mathbb{Z}_2^3$); full smooth compact analytic construction open
Closed-form K3-fibre CY residual	Interval-rigorous, machine-checked bound (box-local)
Physical interpretation of ratios as SM parameters	Conjectural / falsifiable
Dynamical mechanisms (gauge running, mass hierarchy)	Exploratory, not part of the parameter-free core

1. Introduction

1.1 The parameter problem

The Standard Model describes the electroweak and strong interactions with extraordinary precision, yet its minimal form contains nineteen measured inputs, three gauge couplings, the Yukawa couplings spanning five orders of magnitude, the CKM mixing data, and the Higgs sector, fixed only by measurement, with no accepted theoretical origin [1]; the neutrino sector (PMNS mixing, including δ_{CP} , and the mass splittings) adds several more. We refer to the nineteen below as a convenient anchor, not a closed count. Several long-standing puzzles sharpen the question of whether a deeper structure underlies these numbers: the absence of any mechanism fixing the number of generations or the mass hierarchy [5], and the Koide relation $Q = (m_e + m_\mu + m_\tau)^2 / [(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2 \cdot \frac{3}{2}] \approx 2/3$, empirically stable for four decades without explanation [6].

This paper asks a deliberately narrow question. *Suppose* the dimensionless parameters are not free but are topological invariants of a compact internal geometry; what geometry would be required, what relations would follow, and are they borne out? We do not argue that nature realizes this geometry. We argue that a specific, rigid choice: a compact G_2 -holonomy seven-manifold K_7 with $(b_2, b_3) = (21, 77)$ and an $E_8 \times E_8$ structure, produces a set of exact, unfittable relations whose joint agreement with experiment is statistically difficult to attribute to chance, and which make sharp, near-term-testable predictions.

1.2 What this paper claims, and at which register

We call the framework set out here Geometric Information Field Theory (GIFT). It is useful to separate three registers, because the support it offers is uneven across them.

Predictive register. The framework specifies a finite list of topological inputs and derives from them a *parameter-free core* of exact rational relations among integers (§3). These are the only claims on which

the paper stakes itself: each is individually falsifiable, none can be tuned, and they are machine-checked as algebraic identities (§3.4). The decisive empirical test is the neutrino CP phase (§3.3, §7).

Architectural register. The choice of a G_2 -holonomy seven-manifold with $E_8 \times E_8$ structure is *motivated* but not *forced*: G_2 is the unique compact exceptional holonomy admitting the stabilizer chain $G_2 \supset SU(3) \supset SU(2) \supset U(1)$, and the topological constraints fix $(b_2, b_3) = (21, 77)$ and $N_{\text{gen}} = 3$ (§2). We argue this architecture is natural; we do not argue it is necessary.

Ontological register. Whether geometry, information, and energy are three views of a single configuration: a reading resonant with Wheeler’s “it from bit” programme [35], is compatible with the framework and historically motivated it, but is neither required nor demonstrated by any prediction. In the terms of Worrall [36] and Ladyman–Ross [37], the framework is best read as a *moderate structural-realist orientation*: structure carries predictive weight independently of any further ontological commitment. The framework’s fate is decided by experiment in the predictive register, not by adjudication in the ontological one. This is *orientation, not ontology*: we map predictive structure and remain agnostic about what, if anything, it is structure *of*.

1.3 Structure of the argument

Section 2 fixes the geometric architecture ($E_8 \times E_8$, K_7 , the G_2 three-form, and the topological constraints). Section 3 presents the parameter-free core: the exact algebraic relations, the leptonic-CP falsifier, and the scope of formal verification, which is where the paper makes its only positive claims. Section 4 states the rigorous mathematical anchors that make the geometric substrate more than a postulate: an interval-rigorous, machine-checked bound on the Calabi–Yau residual of an explicit closed-form K3-fibre metric, and the certified topological/lattice realization of $(21, 77)$. Section 5 quantifies over-determination against an algebraic null model. Section 6 presents, as explicitly *exploratory* and outside the parameter-free core, the dynamical mechanisms (gauge-coupling running and the lepton-mass hierarchy). Section 7 collects the falsifiable predictions, discusses limitations, and concludes.

2. Geometric architecture

2.1 From octonions to G_2

The octonions $\mathbb{O}$ are the largest normed division algebra; the division-algebra sequence $\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}$ terminates at dimension eight. Their automorphism group $G_2 = \text{Aut}(\mathbb{O})$ is the compact fourteen-dimensional exceptional Lie group and acts irreducibly on the imaginary octonions $\text{Im}(\mathbb{O}) = \mathbb{R}^7$. A seven-manifold of G_2 holonomy is, in this sense, the natural Riemannian home for octonionic structure, just as a Kähler manifold is the natural home for $\mathbb{C}$. Concretely, G_2 is the stabilizer in $\text{GL}(7, \mathbb{R})$ of the standard three-form

$$\varphi_0 = e_{012} + e_{034} + e_{056} + e_{135} - e_{146} - e_{236} - e_{245},$$

in which only seven of the $\binom{7}{3} = 35$ triples carry a nonzero (± 1) coefficient, and its Lie algebra is the kernel of

$L_{\varphi_0} : \mathfrak{gl}(7) \rightarrow \wedge^3(\mathbb{R}^7)^*$, $X \mapsto \mathcal{L}_X \varphi_0$, giving $\dim(\mathfrak{g}_2) = 49 - \text{rank}(L_{\varphi_0}) = 49 - 35 = 14$. The seven coefficients and the dimension count are machine-checked in Lean 4 (`G2ThreeForm.lean`, by `native_decide`).

2.2 $E_8 \times E_8$ and the compactification

E_8 is the largest exceptional simple Lie algebra (dimension 248, rank 8), and $E_8 \times E_8$ (dimension 496) is one of the two anomaly-free gauge structures of the heterotic string [16]. The framework adopts an $E_8 \times E_8$ -motivated exceptional gauge architecture over $\text{AdS}_4 \times K_7$, with the first E_8 containing the Standard Model gauge group through the chain

$$E_8 \rightarrow E_6 \times SU(3) \rightarrow SO(10) \times U(1) \rightarrow SU(5) \rightarrow SU(3) \times SU(2) \times U(1),$$

and the second E_8 a decoupled hidden sector. We stress that $E_8 \times E_8$ enters here as a *gauge structure*, not as a representational container into which fermions are embedded; this distinction is exactly what evades the Distler–Garibaldi obstruction [26] (§7.3).

2.3 K_7 and its topology

K_7 is a compact, simply connected seven-manifold of G_2 holonomy, hence Ricci-flat, with $b_1 = 0$ and $N = 1$ supersymmetry in four dimensions [11], with Betti numbers $(b_2, b_3) = (21, 77)$. Two topological consequences are load-bearing.

- **Generation number.** The chiral-index constraint $(\text{rank}(E_8) + N_{\text{gen}}) b_2 = N_{\text{gen}} b_3$, i.e. $(8 + N_{\text{gen}}) \cdot 21 = N_{\text{gen}} \cdot 77$, has the unique solution $N_{\text{gen}} = 3$.
- **Field-content bounds.** $b_2 = 21$ and $b_3 = 77$ bound the gauge- and matter-sector multiplicities. The non-abelian Standard Model gauge group arises from the algebraic/spectral structure $(\mathfrak{g}_2 \subset \mathfrak{so}(7)$, and the $\mathfrak{so}(8) = \mathfrak{g}_2 \oplus L \oplus R$ triality), not from ADE singularities, consistent with a smooth metric at a generic point of G_2 moduli space.

The pair $(21, 77)$ does not appear among the catalogued compact G_2 manifolds (Joyce orbifolds and the twisted-connected-sum database of Corti–Haskins–Nordström–Pacini [19]); orthogonal TCS is excluded by parity, since $b_2 + b_3 = 98$ is even. It is, however, realized at the topological/lattice level by a Joyce–Karigiannis $\mathbb{Z}_2^3$ orbifold construction [27]. The certified status of that realization, and a rigorous residual bound on the associated K3-fibre metric, are the subject of §4. Figure 1 summarizes the chain and the topological identities it fixes.




Figure 1. The G_2 stabilizer chain $G_2 \supset SU(3) \supset SU(2) \supset U(1)$ (G_2 is the unique compact exceptional holonomy admitting it) and the Standard-Model sector each step supplies (left), alongside the load-bearing topological identities that the Betti data $(b_2, b_3) = (21, 77)$ fix along the chain (right): $N_{gen} = b_2 / \dim K_7 = 3$, $\sin^2 \theta_W = b_2 / (b_3 + \dim G_2) = 3/13$, $Q_{Koide} = \dim G_2 / b_2 = 2/3$, and $\delta_{CP} = \dim K_7 \cdot \dim G_2 + H^* = 197^\circ$.

2.4 The determinant normalization

The metric determinant $\det(g) = 65/32$ plays two distinct roles, and neither is a prediction. As the *overall scale* of the numerical K_7 metric it is a free normalization: a Ricci-flat metric carries a free overall volume, and the entire dimensionless spectrum is invariant under it. Rescaling the metric to $\det(g) = 2$ leaves every spectral ratio unchanged and shifts absolute eigenvalues uniformly by 0.22%, as we verify directly. As an *algebraic value* it enters six closed-form relations of §3; there the load-bearing element is the denominator $\det_{den} = 32 = 2^5$, a structural integer appearing in four of the six, while the numerator 65 is used by only two and is the single component of the normalization not fixed by topology. We are explicit that the value $65/32$ is *imposed*, not derived: it is algebraically suggestive, $65 = \text{Weyl} \times \alpha_{sum}$ and $32 = b_2 + \dim(G_2) - N_{gen}$, so that $\det(g) = 2 + 1/32$, and is fixed as the numerical metric's normalization, but no dimensionless spectral prediction depends on it.

2.5 The topological coupling λ_T

A second structural constant, $\lambda_T := b_3 - \dim(G_2) - p_2 = 77 - 14 - 2 = 61$, enters four closed-form relations of §3. We fix its name and status here to forestall a natural objection. It is a *topological coupling*: a fixed integer combination of the third Betti number, the dimension of the holonomy group, and the second Pontryagin number, and not the intrinsic torsion of the G_2 structure: the Fernández–Gray torsion classes

$W_1 \oplus W_7 \oplus W_{14} \oplus W_{27}$ of K_7 have total dimension $1 + 7 + 14 + 27 = 49$, a distinct invariant. (Its reciprocal $\kappa_T = 1/\lambda_T = 1/61$ is the torsion *capacity* used in the supplements; earlier internal notation that called it a “torsion coefficient” is dropped accordingly.)

What protects the value 61 from the charge of being tuned to one observable is over-determination by two independent witnesses. The cosmological ratios $Y_p = (1 + \dim G_2)/\lambda_T = 15/61$ and $\Omega_m/\Omega_\Lambda = (1 + \dim J_3(\mathbb{O}))/\lambda_T = 28/61$ carry fixed topological numerators and no adjustable content beyond the integer itself, yet the measured primordial-helium fraction and matter-to-dark-energy ratio independently select $\lambda_T = 15/Y_p \approx 61.2$ and $28/(\Omega_m/\Omega_\Lambda) \approx 60.7$. They pin the value to 61 within one unit with no reference to any lepton mass; the structural form of m_τ/m_e of §3.4 is then a consequence of the same integer, not its motivation. The fine-structure relation $\alpha^{-1} = 137 + (65/32)/\lambda_T$ also contains λ_T , but only as a small additive correction to 137 (a unit change in λ_T moves α^{-1} only in its fifth decimal) so we treat it as a *consumer* of the value fixed by the cosmological ratios, not as an independent selector of it.

3. The parameter-free core

This section contains the only claims on which the paper stakes itself. Each is a relation among topological integers (and a small set of standard transcendentals); none contains a continuously adjustable physical parameter; each is individually falsifiable. We present the exact relations, the decisive empirical falsifier, and the precise scope of the formal verification.

Here “parameter-free” is meant precisely: every relation below is fixed by topological integers, with the single shared normalization $\det(g) = 65/32$ of §2.4 entering only six of them, and the calibrated dynamical mechanisms of §6 explicitly excluded. In particular, the ≈ 667 coefficients of the explicit metric approximation (§4.1) enter only the residual-bound certificate; no prediction formula in this section depends on them.

3.1 Exact relations among topological integers

The relations below are *unfittable*: they equate a measured dimensionless quantity to a ratio of fixed integers (or a fixed closed form), so each is simply correct or wrong to within experimental precision. The integers are the structural constants of K_7 and the exceptional algebras, $b_2 = 21$, $b_3 = 77$, $\dim(G_2) = 14$, $\dim(K_7) = 7$, $\dim(E_8) = 248$, $\text{rank}(E_8) = 8$, $H^* = b_2 + b_3 + 1 = 99$, $p_2 = \dim(G_2)/\dim(K_7) = 2$, Weyl = 5, $\lambda_T = 61$ (§2.5), together with π , $\sqrt{2}$, $\ln 2$, and the Riemann ζ values.

Relation	GIFT	Experiment (2024–2025)	Deviation
$\sin^2 \theta_W = b_2/(b_3 + \dim G_2) = 3/13$	0.230769	0.23122 (PDG 2024)	0.20%
$Q_{\text{Koide}} = \dim(G_2)/b_2 = 2/3$	0.666667	0.666664 (PDG 2024)	0.0005%
$N_{\text{gen}} = b_2/\dim(K_7) = 3$	3	3 (observed)	0%

Relation	GIFT	Experiment (2024–2025)	Deviation
$m_s/m_d =$ $p_2^2 \cdot \text{Weyl} = 20$	20	20.0 ± 1.0 (PDG/FLAG)	0.0%
$\Omega_{\text{DM}}/\Omega_b =$ $(1 + 2b_2)/\text{rank}(E_8) =$ $43/8$	5.375	5.364 (Planck 2018)	0.2%
$\alpha_s(M_Z) =$ $\sqrt{2}/(\dim G_2 - p_2) =$ $\sqrt{2}/12$	0.11785	0.1180 ± 0.0009 (PDG)	0.13%

The Koide relation [6], empirically stable for four decades and without an accepted derivation, here reduces to a two-integer identity $Q = \dim(G_2)/b_2 = 14/21 = 2/3$, with no fitting. The Weinberg angle is a ratio of the gauge-counting Betti number to the matter-plus-holonomy count. The dark-to-baryonic ratio $43/8$ and the strange-to-down ratio 20 are exact rationals: a future lattice determination of m_s/m_d to ± 0.5 , or a sub-percent measurement of $\Omega_{\text{DM}}/\Omega_b$, tests them directly. A handful of further ratios (e.g. $m_b/m_t = b_0/2b_2 = 1/42$, the boson-mass ratios $m_H/m_W = 81/52$ and $m_W/m_Z = 37/42$, and several PMNS/CKM entries) belong to the same class but carry scheme- or convention-dependent comparisons; they are tabulated in the supplement.

The full set of 33 core relations is listed in the supplement; six representative ones are shown above. Experimental references are PDG 2024 [1], the FLAG lattice averages, and Planck 2018 [3]; the mass-ratio and coupling rows are compared at standard scheme and scale (α_s in $\overline{\text{MS}}$ at M_Z , quark-mass ratios as scale-invariant ratios of $\overline{\text{MS}}$ running masses, m_τ/m_e from pole masses), as detailed there.

3.2 The leptonic-CP falsifier

The single sharpest prediction is the leptonic CP-violating phase,

$$\delta_{CP} = \dim(K_7) \cdot \dim(G_2) + H^* = 7 \times 14 + 99 = 197^\circ,$$

a pure topological integer with no corrections, decomposing into a local fibre-holonomy term (98) and a global cohomological term ($H^* = 99$). The experimental central value has been notably *unstable* across global oscillation fits, 197° (NuFIT 5.2, 2022), $177^\circ \pm 20^\circ$ (NuFIT 6.0, 2024), and, for normal ordering, 207° (without SK) / 212° (with SK) in NuFIT 6.1 (2025) [32], so the GIFT value 197° remains within the current 1σ interval, although the central value has moved away from it (the long-baseline inputs include the joint T2K–NOvA analysis [34]). We therefore treat δ_{CP} not as a present discriminator but as the decisive *near-term* falsifier: DUNE (2028–2040) is expected to resolve it to between a few degrees and $\sim 15^\circ$ [33]. **Falsification criterion:** the sharp prediction is the single value $\delta_{CP} = 197^\circ$; the interval $[182^\circ, 212^\circ]$ is not a theoretical uncertainty band but a near-term operational window set by the expected $\sim 15^\circ$ experimental resolution. A measured δ_{CP} outside this window at 3σ would place the framework in serious tension, and once next-generation precision reaches a few degrees, exclusion of 197° itself falsifies it. Figure 2 places the prediction against the recent global-fit history and the DUNE timeline.




Figure 2. The sharp prediction $\delta_{CP} = 7 \times 14 + 99 = 197^\circ$ against the central values (with 1σ bands) of successive NuFIT global fits, 197° (5.2, 2022), 177° (6.0, 2024), and $207^\circ/212^\circ$ (6.1, 2025, without/with Super-Kamiokande): the near-term operational window $[182^\circ, 212^\circ]$, and the expected DUNE resolution. 197° was the central value in 2022, drifted, and remains within the current spread; it is to be settled by DUNE (2028–2040), not by today’s fits.

A second clean falsifier requires no flexibility at all: $N_{\text{gen}} = 3$ (§2.3); a fourth chiral generation falsifies the framework immediately.

3.3 Scope of formal verification

The exact relations of §3.1 are machine-checked in Lean 4 [30,31] as algebraic identities: for example

```
theorem weinberg_relation : b2 * 13 = 3 * (b3 + dim_G2) := by native_decide
theorem koide_relation    : dim_G2 * 3 = b2 * 2           := by native_decide
```

The formalization rests on **four substantive axioms on the prediction chain**: one literature bundle (Crowley–Goette–Nordström 2024 [13] and Joyce 2000 [17]), the standard Cheeger inequality, and two geometric Rayleigh/Cheeger bounds on the TCS neck. A separate K3-block module contributes eleven further declarations (five opaque real constants, five externally re-derivable interval-arithmetic brackets of width $\lesssim 10^{-12}$, and one meta-level placeholder) that do **not** enter any gauge, mass, or coupling prediction. The repository declares fifteen axioms in total, with zero **sorry** across 143 files, and the full build completes with zero errors; the master certificate bundles roughly 140 leaf conjuncts across its three pillars (Foundations, Predictions, Spectral).

We are precise about what this establishes. The proofs certify the *algebraic* content: given the topological integers, the numerical relations hold exactly. They do **not** establish the existence of K_7 , the physical

interpretation of the ratios as Standard Model parameters, or the uniqueness of the formula assignments. Formal verification here removes arithmetic and bookkeeping error from a large, cross-coupled set of claims; it is a guarantee of internal consistency, not of physical truth.

3.4 A weaker prediction, flagged

Two frequently cited relations are individually accurate but each invokes an element beyond the topological-integer vocabulary, and we therefore hold them to a weaker standard than the integer identities of §3.1 and exclude them from the unfittable core. (i) $m_\mu/m_e = 27^\varphi = \dim(J_3(\mathbb{O}))^\varphi \approx 207.0$ (experiment 206.768) uses the golden ratio $\varphi = (1 + \sqrt{5})/2$, which arises through the McKay correspondence $E_8 \leftrightarrow$ binary icosahedral group but requires this additional step. (ii) $n_s = \zeta(11)/\zeta(5) \approx 0.96486$ (experiment 0.9649) invokes the Riemann ζ -function, which lies outside the integer grammar and is not machine-checked. Each agrees with experiment, but each rests on a non-structural ingredient, so neither is counted among the unfittable integer relations. (A third relation often quoted in an inelegant base-ten form, $m_\tau/m_e = 3477$, in fact admits the purely structural expression $(b_3 - b_2)(\lambda_T + 1) + \text{Weyl}$ in the topological coupling λ_T of §2.5 (itself fixed independently by the cosmological ratios) machine-checked in Lean, and remains in the core; it is simply not among the representative relations of §3.1.)

4. Rigorous mathematical anchors

The relations of §3 presuppose a geometric substrate: a compact G_2 manifold K_7 with $(b_2, b_3) = (21, 77)$. A reader is entitled to ask whether that substrate is more than a convenient postulate. It is. Two parts of the construction are anchored by machine-checked mathematics that is logically independent of any physical claim: a certified residual bound on an explicit K3-fibre metric, and a certified topological/lattice realization of $(21, 77)$. We state them, and we state plainly what remains open.

4.1 A certified bound on the Calabi–Yau residual

The Ricci-flat Kähler (Calabi–Yau) metric on the $\mathbb{Z}_2^3$ -equivariant K3 fibre admits an explicit, finitely parametrised closed-form approximation (order three, 667 parameters). Its quality is controlled by a machine-checked theorem on the variance of the log-ratio R between the approximate and exact volume forms:

$$\text{Var}(\log R) \leq \varepsilon'_3 = \frac{1321}{10^7} \approx 1.321 \times 10^{-4} < 10^{-3} \quad (\text{safety factor} \times 7.57),$$

and, in the strengthened *box-local* form, this bound holds uniformly over 4000 open boxes of radius $r = 10^{-8}$ covering a frozen, reproducible witness on the surface: not merely at sample points. The certificate is assumption-free in a precise sense: the K3 points are Krawczyk–Rump-certified, the closed-form metric is evaluated in forward interval arithmetic, and Lean itself recomputes the variance envelope from the raw rational box enclosures (via a minimiser bound $\text{Var} \leq \frac{1}{N-1} \sum_n \max((L_n^{\text{lo}} - c)^2, (L_n^{\text{hi}} - c)^2)$), rather than ingesting a pre-computed numerical value. The strict containment of the Krawczyk operator, $K(X) \subsetneq \text{int}(X)$ (28 000 strict integer inequalities) and an independent rational-interval engine that reproduces all

28 000 enclosure endpoints in exact agreement with the floating interval library together narrow the trust boundary to the soundness of the interval evaluation alone. The whole chain is formalized in Lean 4 (K3ClosedFormWitness, K3ClosedFormBoxEnclosures, K3KrawczykContainment; native_decide; zero sorry; no axiom added on this sub-chain) and distributed in the public core package [A, D]; it is the subject of a dedicated companion paper currently under review.

Honest scope. The bound holds over the box-local neighbourhood of the frozen witness. Promoting it to a bound over the entire K3 surface is an open constraint-aware global-positivity (Positivstellensatz/SOS) problem.

Figure 3 shows the box cover and the safety margin of the certified bound.




Figure 3. The certified Calabi–Yau residual bound on the K3 fibre, box-local over 4000 Krawczyk boxes of radius 10^{-8} . (a) Coverage: each box placed by its interval-midpoint $\det(g)$ and $J \wedge J \wedge H$ values, coloured by relative interval width: the certificate spans the surface, not a handful of sample points. (b) Margin: sorted per-box relative widths against the 10^{-3} admissibility threshold; the production envelope $\epsilon = 1.33 \times 10^{-4}$ clears it with a $> 7\times$ safety factor, and the Lean-recomputed variance envelope of the theorem above tightens this to $\epsilon'_3 = 1321/10^7$ (safety $\times 7.57$).

4.2 The topological realization of (21, 77)

The pair (21, 77) is realized at the topological/lattice level by a Joyce–Karigiannis $\mathbb{Z}_2^3$ orbifold construction [27]. A computer-assisted audit closes the Betti count exactly: a V_4 symplectic screen, an anti-symplectic obstruction check, an existence proof for the required K3 lattice via Mukai [28] and Garbagnati–Sarti [29], and the Betti formula $b_2 = 0 + 21$, $b_3 = 22 + 55$. This is Lean-formalized. The route *extends* the catalogue of known compact G_2 manifolds rather than fitting into it: the nearest twisted-connected-sum neighbour [19] lies at distance 7.6 in (b_2, b_3) , and orthogonal TCS is excluded by parity (§2.3).

4.3 What remains open

We claim neither more nor less than the above. An explicit polynomial coordinate model for the $\mathbb{Z}_2^3$ -equivariant K3, the Joyce–Karigiannis 2017 analytic gluing certificate, and therefore the *complete smooth*

analytic compact G_2 construction (holonomy exactly G_2 over the whole manifold) remain open. The framework’s geometric substrate is thus anchored at two levels: the topological/lattice realization of (21, 77) and a certified residual bound on the explicit K3-fibre metric, both machine-checked; the full analytic existence proof is a stated open problem, not an asserted result. This is the credibility the physical correspondence of §3 rests on, and the boundary past which it does not reach.

5. Statistical over-determination

The natural objection to any framework that matches many numbers is selection: with enough formulas, some will fit. We address it on four fronts, all computed on the exact core of §3: a deliberately generous algebraic null model with its constant set declared in full (§5.1); a search-grammar ranking that locates each relation among *all* admissible formulas (§5.2); the observation that the framework does not minimize deviation but leaves accuracy on the table (§5.3); and structural over-determination diagnostics (§5.4).

5.1 An algebraic null model

The null draws candidate formulas from the *same* twenty structural constants the framework uses, nineteen integers ($b_2 = 21$, $b_3 = 77$, $\dim K_7 = 7$, $\dim G_2 = 14$, $p_2 = 2$, $\dim E_8 = 248$, $\text{rank } E_8 = 8$, $\dim E_6 = 78$, $\dim F_4 = 52$, $\text{fund. } E_7 = 56$, $\dim J_3(\mathbb{O}) = 27$, $N_{\text{gen}} = 3$, $\text{Weyl} = 5$, $D_{\text{bulk}} = 11$, $H^* = 99$, $|\text{PSL}(2, 7)| = 168$, $\lambda_T = b_3 - \dim(G_2) - p_2 = 61$ (§2.5), and the two determinant integers 65 and 32) together with the golden ratio φ . Declaring the set in full matters: an undisclosed constant is an undisclosed degree of freedom. Over these, with five binary operations (addition, subtraction, multiplication, division, exponentiation) and five unary transforms (identity, inverse, square root, square, natural log), we generated 4 187 310 distinct formula values by exhaustive depth-1/2 enumeration ($\approx 1.8\text{M}$) plus $\approx 3\text{M}$ random expression trees of depth 2–4. This null is far stronger than a uniform-deviation null precisely because it is built from the framework’s own vocabulary.

The decisive figure is set-level and assumption-free: among 3 000 000 random formula *sets*, none reproduces the framework’s joint profile (mean deviation, exact-match count, and within-1% count together): an empirical upper bound of order 10^{-6} at this sampling, with no independence assumption. A per-observable product over the exact relations gives the far smaller $P \approx 10^{-134}$, but because the relations are cross-coupled (§5.4) this product is *not* a calibrated joint probability and we do not rely on it; we report it only as an order-of-magnitude indicator of non-genericity. (A uniform-deviation null, which treats predictions as arbitrary reals, gives a vastly smaller probability still; we regard it as far too generous and discard it.) Figure 4 shows both the set-level result and the per-observable improbabilities.




Figure 4. The algebraic null model (NuFIT 6.1 / PDG 2024 / Planck 2018, 33 core relations; 4 187 310 formula values; 3 000 000 random formula-sets). (a) Set-level and assumption-free: no random formula-set among three million reproduces the framework’s joint profile: the best random set reaches at most one exact match and three within 1%, against GIFT’s four and twenty-nine. (b) Per-observable $\log_{10} P$ that a random formula matches GIFT’s deviation; their product $\approx 10^{-134}$ is an indicator of non-genericity, not a calibrated joint probability, because the relations are cross-coupled (§5.4).

5.2 Where the relations sit in the admissible space

A sharper test fixes a declared formula grammar (an invariant set, a maximum tree depth, a maximum integer coefficient, and a per-class complexity budget), enumerates *every* admissible formula for each observable, and asks where the framework’s formula ranks. Scored by deviation alone, the GIFT formula is the error-minimizer for the majority of the core relations and lies on the error–complexity Pareto frontier for all of them. A robust subset (the Koide ratio, the generation count, and the fine-structure relation) remains not merely error-optimal but *composite*-optimal (jointly lowest in deviation and complexity) even when the grammar is widened to the full constant set and depth four, where tens of thousands of formulas compete: the Koide relation, for instance, ranks first of 2 446 admissible formulas by both criteria. We claim no more than this. Raw error-rank is *not* robust to grammar width (opening the grammar lets accidental near-misses accumulate), so the load-bearing statement is Pareto-membership for the core and composite-optimality of the robust subset, not a universal “rank one.”

5.3 The framework leaves accuracy on the table

This is the strongest single argument against post-hoc fitting, and it runs opposite to the usual suspicion. A curve-fitter would adopt, for each observable, the closest formula its grammar admits. The framework does not. Bound to a single fixed structural assignment, it forgoes available accuracy: in the declared grammar a strictly more accurate admissible formula exists for, among others, the reactor angle θ_{13} (the GIFT formula ranks tenth by deviation), the atmospheric angle θ_{23} (third), the muon-to-electron mass ratio (second), the dark-energy fraction (third), and the Higgs self-coupling (seventh). Were the aim to maximize agreement, switching to these error-minimizers would push the mean deviation well below 0.1%

and invite a “theory-of-everything” claim; the framework declines, because it is selected for structural coherence, not fitted for deviation. Accuracy *foregone* is the signature of a structure-first selection, not of curve-fitting.

5.4 Over-determination, cross-coupling, and complexity

Three further diagnostics corroborate this. First, a singular-value decomposition of the 20×33 constant-usage matrix gives an effective rank ≈ 15.5 : the 33 relations rest on roughly fifteen effective constant-combinations, about twice as many constraints as effective degrees of freedom: read qualitatively, the signature of an over-determined rather than a fitted system. Second, the relations are not independent silos: of the 528 observable pairs, 155 share a structural constant and all 33 lie in a single connected component (the Weinberg angle and the Koide ratio both rest on $(b_2, \dim G_2)$; m_b/m_t and m_μ/m_τ both on $2b_2 = 42$). Third, precision is not bought by complexity: a random-forest regression of $|\text{deviation}|$ on formula features (constants used, largest constant, depth, sector), under leave-one-out cross-validation across the 33 relations, yields $R^2 < 0$, complex formulas fit no better than simple ones. Finally, the over-determination does not lean on the determinant value: $\det(g)$ is in part an overall metric scale (§2.4), of the six relations that use it only two use its numerator, and excluding all six normalization-dependent observables leaves these diagnostics unchanged.

6. Exploratory dynamical mechanisms

Beyond the parameter-free core, the framework includes dynamical mechanisms that link the topology to running couplings and the mass hierarchy. We present them separately and label them clearly: unlike §3, these involve calibrated quantities and are therefore *not* unfittable. They are offered as exploratory evidence of geometric coherence: not as part of the paper’s load-bearing claims, and deliberately excluded from the statistical case of §5.

6.1 Gauge-coupling running

The topological boundary values $\sin^2 \theta_W = 3/13$ and $\alpha_s = \sqrt{2}/12$, evolved with the two-loop MSSM renormalization-group equations, reproduce the measured couplings to a few percent ($\sin^2 \theta_W$ to 2.8%, α_s to 3.7% with split-spectrum threshold matching). Notably, the topological values themselves are *more* accurate than their RGE-evolved counterparts, which suggests they may behave as infrared fixed points rather than ultraviolet boundary conditions. A by-product (the “B-test” identity) expresses the unification-scale α_{em}^{-1} as $91\sqrt{2} = \dim(\Lambda^2 \mathfrak{g}_2) \cdot \sqrt{2}$ and yields the holonomy sequence $\alpha_1^{-1} : \alpha_2^{-1} : \alpha_3^{-1} = \dim(G_2) : \dim(K_7) : p_2 = 14 : 7 : 2$. We record this as a structural observation, with its residual $\sim 0.5\%$ discrepancy, not as a precision claim.

6.2 The lepton mass hierarchy

Two complementary geometric mechanisms each reproduce the five-order lepton hierarchy to a few percent: Wilson-line eigenvalue splitting on the K3 fibre (which fixes the relative structure) and associative 3-cycle

volume differences (which fix the exponential scale). Combined through the geometric normalization $\alpha = e^K = \hat{V}^{-3}$ (the Kähler potential of K_7 , not a fit parameter), they reach sub-percent agreement: $m_\tau/m_e = 3485$ (0.24%), $m_\tau/m_\mu = 16.69$ (0.75%), $m_\mu/m_e = 208.8$ (0.97%). We state plainly that the Wilson-line coupling $c = 0.452$ and the cycle-to-generation assignment are *calibrated*; these results are corroborative, not unfittable.

6.3 A remark on the supersymmetric spectrum

A 4D $N = 1$ supergravity reduction yields a G_2 -MSSM-type spectrum (gravitino ~ 0.17 TeV, moduli ~ 3 TeV) via hidden-sector gaugino condensation. We record this only as an internal consistency check: much of the corresponding simplified-model parameter space is already constrained by collider searches, and the spectrum is viable only in compressed or suppressed-coupling realizations. We do not present it as a prediction.

7. Discussion, falsifiability, and conclusion

7.1 Falsifiable predictions

Test	Prediction	Window / timeline	Status
DUNE	$\delta_{CP} = 197^\circ$	outside $[182^\circ, 212^\circ]$ at 3σ falsifies; 2028–2040	critical
Any collider	$N_{\text{gen}} = 3$	a fourth generation falsifies	immediate
Lattice QCD	$m_s/m_d = 20$	target ± 0.5 by ~ 2030	strong
DUNE / Hyper-K	θ_{23} in the upper octant	octant resolution, 2030s	live bet (§7.2)
FCC-ee	$\sin^2 \theta_W = 3/13$, $Q_{\text{Koide}} = 2/3$	$\sim 10^{-5}$ precision, 2040s	definitive

7.2 The θ_{23} octant: a falsifiable bet

The framework predicts $\theta_{23} = \arcsin((b_3 - p_2)/H^*) = \arcsin(75/99) = 49.25^\circ$, i.e. the *upper* octant ($\sin^2 \theta_{23} \approx 0.574$). We present this as a falsifiable bet, not a settled success. The θ_{23} octant is one of the genuinely unresolved questions in neutrino physics, and the global-fit best fit has itself crossed maximal mixing between recent releases: the upper octant in NuFIT 6.0, and a lower-octant best fit near 43.3° in NuFIT 6.1. GIFT commits to the upper octant; DUNE and Hyper-Kamiokande are expected to resolve the octant decisively. We flag this as the framework’s most exposed near-term prediction and decline to soften it.

7.3 Why GIFT evades the E_8 no-go

Distler and Garibaldi proved that no single E_8 contains a representation decomposing into the chiral Standard Model fermion content [26]. GIFT is not subject to this no-go: $E_8 \times E_8$ enters as a *gauge*

structure, not as a container into which fermions are embedded. The Standard Model gauge group emerges from a breaking chain (§2.2), the generation number from the topological index constraint (§2.3), and chirality from the compactification. The relationship between E_8 and the particle content is cohomological (emergent from geometry), not representational (embedding into an algebra).

7.4 Limitations

We list the framework’s principal open points without minimizing them.

- The complete *smooth analytic* compact G_2 construction is open (§4.3); what is certified is the topological/lattice realization of (21, 77) and the K3-fibre residual bound.
- The formula-selection principle is not derived from first principles. The observed relations are statistically non-generic within the declared grammar (§5), but a deductive rule fixing *which* relation holds for *which* observable remains open.
- The dynamical mechanisms of §6 are calibrated, not parameter-free, and are excluded from the statistical case.
- The δ_{CP} central value is experimentally unsettled, and the θ_{23} octant is a live bet (§7.2).
- The second (hidden) E_8 is a candidate dark sector with no direct observable coupling; this is speculative and carries no weight in the predictive register.

7.5 Conclusion

We have set out a framework in which the dimensionless parameters of the Standard Model appear as topological invariants of a compact G_2 -holonomy manifold K_7 with $E_8 \times E_8$ structure, with no continuously adjustable physical parameter. Its load-bearing content is a parameter-free core of exact, machine-checked relations among topological integers (§3), non-generic under a declared algebraic null (no random set among three million reproduces its joint precision profile, §5), anchored by a certified residual bound on an explicit K3-fibre metric and a certified topological realization of (21, 77) (§4), and falsifiable through δ_{CP} , N_{gen} , and the θ_{23} octant (§7). We do not claim the framework is correct: it may be a genuine geometric insight, an effective approximation, or an elaborate coincidence, and only experiment can decide between them. What we claim is narrower, and we hope more useful: that the correspondence is sharp, the inputs explicit, the geometry anchored where we say it is anchored and open where we say it is open, and the predictions testable within a decade. The ultimate arbiter is experiment.

Declarations

Competing interests. The author declares no competing interests.

Funding. This research received no external funding.

Data availability. All Lean 4 source, Python scripts, and certificate data are openly available in the GIFT core repository (github.com/gift-framework/core). The observable dataset and the supplementary

reproducibility bundles are archived on Zenodo (DOIs [A]–[D] below). No new experimental datasets were generated.

Use of generative AI and AI-assisted technologies. During the preparation of this work the author used AI-based tools (Claude, Anthropic; GPT, OpenAI) to assist with drafting and refining text, with derivations, and with formal-verification scaffolding. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication. No generative-AI tool is an author. Figures, where present in the final version, are produced by deterministic plotting code from the cited computations; no AI-generated images are included.

References

- [1] Particle Data Group, *Phys. Rev. D* **110**, 030001 (2024).
- [2] S. Weinberg, *Phys. Rev. D* **13**, 974 (1976).
- [3] Planck Collaboration, *Astron. Astrophys.* **641**, A6 (2020), arXiv:1807.06209.
- [4] A. G. Riess et al., *ApJL* **934**, L7 (2022).
- [5] C. D. Froggatt, H. B. Nielsen, *Nucl. Phys. B* **147**, 277 (1979).
- [6] Y. Koide, *Lett. Nuovo Cim.* **34**, 201 (1982).
- [7] C. Furey, *Standard Model Physics from an Algebra?*, PhD thesis, U. Waterloo (2015), arXiv:1611.09182.
- [8] C. Furey, M. J. Hughes, *Phys. Lett. B* **831**, 137186 (2022), arXiv:2210.10126.
- [9] R. A. Wilson, “On possible embeddings of the Standard Model in E_8 ,” arXiv:2404.18938 (2024). (*real form* $E_8(-24)$)
- [10] T. P. Singh et al., arXiv:2206.06911 (2024).
- [11] B. S. Acharya, S. Gukov, *Phys. Rep.* **392**, 121 (2004).
- [12] L. Foscolo, M. Haskins, J. Nordström, *Duke Math. J.* **170**(15), 3323 (2021).
- [13] D. Crowley, S. Goette, J. Nordström, “An analytic invariant of G_2 manifolds,” *Invent. Math.* **239**, 865 (2025).
- [14] T. Dray, C. A. Manogue, *The Geometry of the Octonions*, World Scientific (2015).
- [15] J. F. Adams, *Lectures on Exceptional Lie Groups*, U. Chicago Press (1996).
- [16] D. J. Gross, J. A. Harvey, E. Martinec, R. Rohm, *Nucl. Phys. B* **256**, 253 (1985).
- [17] D. D. Joyce, *Compact Manifolds with Special Holonomy*, Oxford U. Press (2000).
- [18] A. Kovalev, *J. Reine Angew. Math.* **565**, 125 (2003).
- [19] A. Corti, M. Haskins, J. Nordström, T. Pacini, *Duke Math. J.* **164**, 1971 (2015).
- [20] R. Harvey, H. B. Lawson, *Acta Math.* **148**, 47 (1982).
- [21] B. S. Acharya, *Class. Quant. Grav.* **19**, 5619 (2002).

- [22] B. S. Acharya, E. Witten, arXiv:hep-th/0109152 (2001).
- [23] E. Witten, *Nucl. Phys. B* **471**, 135 (1996).
- [24] B. S. Acharya, K. Bobkov, G. Kane, P. Kumar, J. Shao, *Phys. Rev. D* **76**, 126010 (2007).
- [25] M. Atiyah, E. Witten, *Adv. Theor. Math. Phys.* **6**, 1 (2002).
- [26] J. Distler, S. Garibaldi, *Commun. Math. Phys.* **298**, 419 (2010).
- [27] D. Joyce, S. Karigiannis, *J. Differential Geom.* **117**, 255 (2021), arXiv:1707.09325.
- [28] S. Mukai, *Invent. Math.* **94**, 183 (1988).
- [29] A. Garbagnati, A. Sarti, *J. Algebra* **318**, 323 (2008), arXiv:math/0603742.
- [30] L. de Moura, S. Ullrich, *CADE* 28, 625 (2021).
- [31] mathlib Community, github.com/leanprover-community/mathlib4.
- [32] NuFIT 6.1 (2025), www.nu-fit.org; I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz et al., *JHEP* **12** (2024) 216, arXiv:2410.05380 (NuFIT 6.0).
- [33] DUNE Collaboration, arXiv:2103.04797 (2021).
- [34] T2K and NOvA Collaborations, “Joint neutrino oscillation analysis from the T2K and NOvA experiments,” *Nature* **646**, 818–824 (2025), arXiv:2510.19888.
- [35] J. A. Wheeler, “Information, physics, quantum: the search for links,” in *Complexity, Entropy, and the Physics of Information* (W. H. Zurek, ed.), Addison-Wesley (1990), pp. 3–28.
- [36] J. Worrall, “Structural realism: the best of both worlds?,” *Dialectica* **43**, 99–124 (1989).
- [37] J. Ladyman, D. Ross, *Every Thing Must Go: Metaphysics Naturalized*, Oxford U. Press (2007).
- [38] J. C. Baez, “Octonions and the Standard Model,” n-Category Café (2020–2025).

Datasets and code (Zenodo).

- [A] B. de La Fournière, “A Certified Torsion-Free G_2 Structure on a TCS Neck Model via Computer-Assisted Proof,” Zenodo 10.5281/zenodo.19892350 (2026).
- [B] B. de La Fournière, “Spectral Geometry of an Explicit G_2 Metric: Laplacian Spectrum and Harmonic Forms,” Zenodo 10.5281/zenodo.19893371 (2026).
- [C] B. de La Fournière, “Newton–Kantorovich diagnostics on a Donaldson K3 metric: two β estimates, machine-checked inequalities,” Zenodo 10.5281/zenodo.19708916 (2026).
- [D] B. de La Fournière, “An Explicit Closed-Form G_2 Ansatz on a K3-Coassociative Neck with Hyperkähler Rotation and Picard–Lefschetz Wirtinger Certificate,” Zenodo 10.5281/zenodo.20039066 (2026).