Is the 13D Shape the Simplest? (The Minimality Case — Selected, Not Proven) — Fable_Version rendered package. Rendered from `SHAPE_MINIMALITY_PHD.md`; frozen technical content unchanged by rendering.

Is the 13D Shape the Simplest? (The Minimality Case — Selected, Not Proven)

> **Honest ceiling, stated up front.** This document explains the *minimality case* for the Fable program's frozen 13-dimensional internal geometry. It is a **serious candidate / partial unification — NOT a validated theory, and NOT a proof of simplicity.** The shape is an **axiom/anchor**, not a derivation. What is "derived" here is always *conditional*: consequences that follow **given the shape** and **given the observed Standard-Model spectrum E** (we write this **DERIVED-GIVEN-E**). The geometry is **DECLARED-FROZEN with minimality OPEN** — *selected within a declared category, not forced absolutely*. Throughout: **given-E ≠ derivation of E; selection ≠ derivation; category-relative ≠ absolute; axiom-closed ≠ proven.**

---

1. The object whose simplicity is in question

The Fable program freezes a single internal geometry — the "active branch" — and evaluates every downstream claim against that frozen object. Only one layer of it carries metric dimension:

$$\mathfrak{B}_{\text{active}} = \underbrace{\left[\mathcal{M}_4 \times K_6 \times S^2 \times S^1_Y/\mathbb{Z}_2\right]_\times}_{\text{STAGE — metric geometry, } D=13} \oplus \underbrace{\left[F^+_{\text{finite}} \oplus C_{\text{admiss}}\right]_\oplus}_{\text{RULEBOOK — 0 dims}} \otimes \underbrace{\left[E_{\text{matter}} \oplus E_{\text{gauge}} \oplus E_{\text{Higgs}} \oplus E_{\text{pro}}\right.}_{\text{ACTORS — 0 dims}}$$

with the dimension ledger

$$D = 4 + 6 + 2 + 1 = 13,$$

and the internal color carrier the **flag manifold** $K_6 = SU(3)/T^2$. The factors route to forces as **isometries of the internal spaces** (Kaluza–Klein):

Internal factor	Routes to	Mechanism
$K_6 = SU(3)/T^2$	$SU(3)_c$ color	left-isometry algebra $\mathfrak{su}(3)$
$S^2 = SU(2)/U(1)$	$SU(2)_L$ weak	isometry algebra $\mathfrak{su}(2)$
$S^1_Y/\mathbb{Z}_2$	$U(1)_Y$ hypercharge	translation isometry + orbifold chirality filter

Three further structural facts are read off this object: three generations from a **spin-$\mathbb{C}$ index** $\chi(K_6, E) = -3$; charge quantization from a global $\mathbb{Z}_6$ identification gluing the three centers; and a Wilson-line (Hosotani) Higgs. The object is content-addressed and frozen (branch hash `dcc66f1b2685`, manifest meta-hash `a5b1e6f9d951`), so the geometry a reviewer attacks is byte-identical to the one the claims were evaluated against.

The question of this document is narrow and specific: **is this 13D shape the *simplest* object that could do the job?** The honest answer has a strong part and an open part, and the discipline of this program is to keep them strictly separated.

---

2. What "minimal" can and cannot mean here

There are two radically different claims one might attach to the word "minimal," and conflating them is the cardinal error:

1. **Absolute minimality** — *no object, anywhere, in any framework, is a shorter description of the same physics*. This is a **universal negative** over all conceivable architectures.
2. **Category-relative (selector-) minimality** — *inside a pre-declared search category (a grammar of admissible structures + a description-length metric + a finite dimension ladder), the 13D branch is the shortest survivor.*

Claim (1) is not available, and the program does not assert it. "No competitor is shorter, anywhere" is the Kolmogorov complexity $K(T)$ of the constant set being described, which is **uncomputable** — there is no algorithm that certifies a universal shortest-description. Worse, the comparison ledger that scores any geometry against rivals carries the observed spectrum **E** on *both* sides (every candidate is required to reproduce the same Standard-Model content), so even in principle the strongest reachable statement is "shortest generator of the flavor/coupling data **given E**" — never "the geometry from nothing." This is the **bottoms-on-E** wall.

Claim (2) is what the minimality case actually argues, and it is a real, technical, falsifiable claim — but it is explicitly **selected, not forced absolutely**. The geometry-specification gate is graded **DECLARED-FROZEN**: a *freeze-and-reproduce certificate, not a uniqueness theorem*. It certifies three things and only three — **specificity** (the branch is fully, unambiguously specified before any downstream gate is scored), **no-layer-smuggling** (no gate is closed by content not declared in its proper $\times/\oplus/\otimes$ layer), and **reproducibility** (every primitive is hashed; a reproducer regenerates each hash). It does **not** certify that the object is unique, forced, or derived.

So when the rest of this document says a carrier is "forced," it always means **forced within the declared grammar** — the category whose defining axiom is *gauge forces are the isometries of the internal factors*. A reviewer who works in a different category (e.g. gauge groups sourced from bundle structure groups or branes, as in string/M-theory) is outside the scope of these forcing arguments, and the program declares this openly.

With that fence in place, here is the substance — and it is genuinely strong where it is strong.

---

3. The three forcing arguments (the genuine internal content)

The Standard-Model gauge **outcome** $SU(3)_c \times SU(2)_L \times U(1)_Y$ is, by itself, **not** a discriminating result. String, M-, F-theory, noncommutative geometry, and lattice constructions all recover it by their own routes; recovering the SM gauge group is a filter *every* serious framework passes. On the outcome, the program **ties**. The genuine Fable-internal content is not the outcome but the **carrier-forcedness**: *given* that you build gauge from isometries, *which* internal spaces are forced, and which cheaper-looking competitors collapse. Three results carry this, and all three are hand-checkable structural facts, not fitted numbers.

3.1 F1 — the weak carrier S^2 is forced (no abelian space carries $SU(2)$)

Fact F1: flat/abelian geometries cannot supply a non-abelian force. The isometry group of any torus T^n is $U(1)^n$, which is **abelian** and contains no $SU(2)$; the same holds for torus orbifolds. Under coset dimensional reduction, an abelian-isometry carrier yields only abelian gauge (the centralizer of an abelian group supplies no non-abelian survivor).

Therefore the non-abelian weak force $SU(2)_L$ **cannot** be carried by any abelian factor of any dimension. It requires a genuinely non-abelian-isometry carrier, and the minimal one is the two-sphere

$$S^2 = SU(2)/U(1), \qquad \dim S^2 = 2, \qquad \mathrm{Isom}(S^2) = SU(2).$$

This is a clean classification statement that closes the *cheaper direction* over an entire shelf of candidates, not against a single named rival. Status: a stated fact, promotable to a closed no-go theorem (it is a bounded Lie-theory argument); honestly **AXIOM-CLOSED, promotable to DERIVED within the grammar**. It is category-relative: a bundle/brane framework routes around F1 by sourcing $SU(2)$ from structure groups, which is a *different category*.

3.2 F2 — the hypercharge carrier must be the orbifold $S^1_Y/\mathbb{Z}_2$ (no mirror fermions)

Fact F2: a closed odd-dimensional factor produces no net chirality. The chiral (Dirac) index on a *closed* odd-dimensional manifold vanishes identically, so a **bare** circle S^1_Y mirrors every fermion — for each left-handed mode it keeps a right-handed partner. A full mirror sector would have shown at LEP: the measured light-species count is

$$N_\nu = 2.984 \pm 0.008 \quad (\text{LEP/SLD } Z\text{-width}),$$

with no room for mirror generations. The bare circle is therefore **falsified by measurement**.

The repair is the $\mathbb{Z}_2$ **orbifold fold**: $S^1_Y/\mathbb{Z}_2$ has fixed-point boundaries, and boundaries re-open a one-sided chirality channel (Atiyah–Patodi–Singer index). Hypercharge $U(1)_Y$ is the translation generator along the parent circle, quantized by its topology. So the hyper carrier is forced to be the **folded** circle, not the bare one — again a forcing over a whole shelf, anchored to a real number (2.984 ± 0.008). (The full *uniqueness* of the $\mathbb{Z}_2$ fold versus other boundary data is a downstream chirality question; for the carrier identity, $S^1_Y/\mathbb{Z}_2$ with $U(1)_Y$ surviving is what is owed and delivered.) Status: **AXIOM-CLOSED** for the carrier identity.

3.3 The color carrier $K_6 = SU(3)/T^2$ is the unique *clean* $SU(3)$ carrier

This is the strongest and most interesting result, because it is **architecture-neutral within the coset-reduction grammar** — it follows from representation theory plus the CSDR centralizer rule, with no fitted number anywhere.

Under coset (Forgács–Manton / CSDR) reduction, a homogeneous carrier G/H delivers G as gauge **only if the isotropy subgroup H does not itself act as gauge** — the surviving 4D gauge content is governed by the centralizer of the isotropy embedding. The decisive structural fact for $G = SU(3)$:

$$\text{the maximal torus } T^2 \text{ is the unique } \textit{purely abelian} \text{ isotropy of } SU(3), \qquad C_{SU(3)}(T^2) = T^2 \text{ (Cartan}$$

Because T^2's centralizer is just T^2 itself, it injects **no spurious non-abelian gauge factor**: $K_6 = SU(3)/T^2$ is a **clean** color carrier. Its full isometry algebra is exactly $\mathfrak{su}(3)$, and the recovered gauge algebra of $K_6 \times S^2 \times S^1_Y$ is the multiset $\{\mathfrak{su}(3), \mathfrak{su}(2), \mathfrak{u}(1)\} - 8 + 3 + 1 = 12$ generators, rank 4, **equality** (no extra unbroken factor, no missing factor), not mere containment.

The cheaper-looking competitor and why it dies. The naive economy move is to use the lower-dimensional $\mathbb{CP}^2 = SU(3)/U(2)$ (which would shave dimensions). But the isotropy $U(2) = (SU(2) \times U(1))/\mathbb{Z}_2$ is **non-abelian** and sits inside color $SU(3)$. By the CSDR centralizer rule it is **gauge-active**, forcing a lose–lose fork (the manuscript states it disjunctively):

- **Keep** the S^2 and S^1 factors → the active $U(2)$ isotropy over-produces an **extra unwanted** $SU(2) \times U(1)$, so the equality clause fails and **Gate 2 (gauge recovery) breaks**; *or*
- **Drop** them → the weak/hyper structure gets **isotropy-locked inside color $SU(3)$**, violating the matter-routing binding.

Either branch fails. This is not a hand-wave: the **11-dimensional $\mathbb{CP}^2$ route was built end-to-end** (as an adversarial sandbox) and **breaks at the gauge-recovery gate**. (A separate, weaker historical reason — that $\mathbb{CP}^2$'s family count is "tunable" — has been *retired* as unsound, since its three-family count is a discrete spin-$\mathbb{C}$ index $r(r+1)/2 = 3$, not a continuous dial; the abelian-isotropy argument above is the correct, stronger, target-blind reason.)

This is the part of the case that is genuinely **not a tie**: a Fable-internal elimination of the cheaper carrier, from group theory alone. Status: **AXIOM-CLOSED on the named shelf** (it is a *theorem* for the clean class, with the one concrete competitor built and broken), reducible to **DERIVED within CSDR** by a bounded enumeration of the (small) subgroup lattice of $SU(3)$.

3.4 What the three results jointly establish — and do not

Jointly, F1, F2, and the abelian-isotropy theorem say: **given the "forces = isometries" grammar and given E, the weak, hypercharge, and color carriers each close their cheaper direction over whole shelves of alternatives, and the one concrete cheaper color competitor is explicitly killed.** That is a strong, specific, technical case. It is what the gauge-recovery gate is graded **DERIVED-GIVEN-E**: a rigid algebra recovery plus carrier-forcedness, conditional on E and on the selected geometry.

What it does **not** establish: that the carriers are forced *across all geometries* (the recovery is an existence-and-equality statement on **one** frozen branch — cross-geometry uniqueness is not claimed and cannot be obtained from this route); and that the forcing is *architecture-neutral* against a reviewer who rejects the isometry grammar (the neutrality is asserted, not proven, outside the category).

---

4. The four honest reasons "no simpler competitor anywhere" stays OPEN

Even granting §3 in full, the leap from "the carriers are forced within the grammar" to "the 13D shape is the simplest object, period" requires four further things that the program candidly lists as **open** — these are the corpus's own declared first review targets, not externally imposed objections.

4.1 Absolute irreducibility is uncomputable (the Kolmogorov wall)

As in §2, "no competitor anywhere is shorter" is a universal negative equivalent to computing $K(T)$ — provably unreachable in full. The honest move is not to attempt it but to **declare a finite grammar** of role-mechanisms and assert minimality *relative to that grammar*, which converts an uncomputable universal negative into a finite, decidable lower-bound problem. That conversion is a real epistemic gain (a *sharper-OPEN*), **not** a closure. The absolute target is retired as the wrong target; the right target (close the grammar-relative matrix) is §4.2.

4.2 The dimension ladder is a survey, not a certified classification

Under the program's description-length (MDL) metric, the 13D branch wins **every considered rung** of a dimension ladder ($D = 4 \dots 12$ plus non-dimensional alternatives): in the current matrix, 10 candidates LOSE to 13D, 1 fails structurally, 0 refute it. **But this is a first-pass survey, not a certified classification theorem.** Closing it requires a *role-mechanism normal-form / exhaustion theorem*: prove that every architecture satisfying the target reduces, *without cost increase*, to one of finitely many normal-form classes (over axes like gauge-origin, chirality-origin, family-count, flavor-origin, scale-origin), then lower-bound each class. That theorem is **OPEN** — bounded and tractable-in-principle, but not done. Until it is, "no competitor below 13D is shorter" is *surveyed, not proven*. A genuine outcome of completing it could be **REFUTED** (some class has a strictly shorter recipe) — which would be an honest, valuable discovery, not a failure.

4.3 The economy metric itself is unproven (MDL vs dimension-first — the decisive seam)

Every "13D wins" verdict holds **under the MDL / description-length metric**, where an independently-measured real anchor costs $\Theta(\log(1/\Delta_0))$ bits (large) and a discrete structural choice — a coset, a $\mathbb{Z}_6$, an integer index — costs $O(1)$ bits (small), so *anchor-burden dominates dimension-burden*. **Under a different, equally natural metric — dimension-first lexicographic order — a clean 4D effective field theory wins** ($k_{\text{dim}} = 4 < 13$) *regardless of how many reals it injects*, and the entire ladder folds for everyone.

Which metric is correct cannot be read off any current measurement. It must be decided by an architecture-neutral principle (the program's candidate: *granularity ⇒ a finite-record universe ⇒ MDL is the natural simplicity measure*), and that bridge theorem is **UNPROVEN**. Crucially, any such proof must charge the **geometry→observables generator map**: if that map is itself a large injected object, the 13D win evaporates and the 4D EFT wins. This is the single seam that decides **forced-given-E vs**

merely-selected, and a faithful **"the 4D EFT is genuinely cheaper"** outcome remains a **live, honest possibility** — explicitly not to be feared or suppressed.

4.4 The honest input cost is ~13–14 reals, not "4 in"

A public-facing program must not overstate its economy. The headline "4 inputs → 22 outputs" (the four declared anchors $\{M_{\rm Pl}, \alpha_i(M_Z), y_t, |V_{us}|\}$) is **overstated**. The honest charged cost is approximately **4 anchors + 9–10 injected reals ≈ 13–14 reals total** — the additional reals being fitted normalizations (N_d, N_e, N_ν), the threshold triple $\delta = (+4.8424, -3.1112, -1.7313)$, and the Wilson-line angle $\theta^\star_H$. (The opposite over-correction — treating frozen within-sector ratios as independently injected, giving "~14 in" — is *also* wrong and is forbidden by the program's own normalization rules. The honest figure is the middle one.) The economy is **real but modest**, not the advertised factor. This is a disclosure correction, not a physics result.

---

5. Putting it together: a scorecard

Claim	Status	Why
The branch is fully specified, frozen, reproducible	**certified** (DECLARED-FROZEN)	content-addressed hashes + reproducer; specificity + no-smuggling + reproducibility
Recovers $SU(3) \times SU(2) \times U(1)$ (equality, $\{8, 3, 1\}$, rank 4)	**DERIVED-GIVEN-E** but a **rival TIE**	every serious framework recovers it; not discriminating
S^2 forced for weak (F1)	**AXIOM-CLOSED in-grammar** → DERIVED reachable	no abelian space carries $SU(2)$; category-relative
$S^1_Y/\mathbb{Z}_2$ forced for hypercharge (F2)	**AXIOM-CLOSED in-grammar**	bare circle mirrors → LEP $N_\nu = 2.984$ excludes it
$K_6 = SU(3)/T^2$ is the **unique clean** $SU(3)$ carrier	**AXIOM-CLOSED on named shelf** (a theorem); CP² built-and-broken	$C_{SU(3)}(T^2) = T^2$; $U(2)$ isotropy over-produces gauge
K_6 is the unique $SU(3)$ carrier over **all** carriers	**OPEN / AUDIT**	shelf-completeness (incl. non-homogeneous) not certified
13D is the minimal/forced geometry, absolutely	**OPEN — not claimed**	uncomputable; survey not classification; metric unproven; bottoms on E
"4 inputs → 22 outputs"	**overstated**	honest cost ≈ 13–14 reals

The pattern is consistent and, properly read, honest: **the carrier-forcedness is strong and specific within a declared grammar; the absolute-minimality claim is open at four named, non-trivial seams.**

---

6. The honest bottom line

The minimality case for the 13D shape is **strong where it is strong and open where it is open**, and the program's discipline is to never let the two blur:

- **Strong (in-grammar, given E):** Once you commit to building gauge forces from internal isometries, the weak carrier (S^2, by F1), the hypercharge carrier ($S^1_Y/\mathbb{Z}_2$, by F2), and the *cleanness* of the color carrier ($K_6 = SU(3)/T^2$, by abelian-isotropy uniqueness — $C_{SU(3)}(T^2) = T^2$, with $\mathbb{CP}^2$ built end-to-end and broken at gauge recovery) are forced over whole shelves of competitors, anchored to real structure and a real measurement ($N_\nu = 2.984 \pm 0.008$). These are theorems or near-theorems, not fitted coincidences.
- **Open (the universal negative):** "No simpler competitor *anywhere*" is **not** established and, in its absolute form, **cannot** be — it is uncomputable (Kolmogorov), the dimension ladder is a first-pass survey rather than a certified classification, the description-length economy metric that makes 13D win is itself unproven against a dimension-first metric under which a 4D effective theory wins, and every chain ultimately **bottoms on the observed spectrum E** (given-E ≠ derivation of E). A "the 4D effective theory is actually cheaper" verdict remains a live, honest possibility.

So the precise, non-promotional statement is this: **the 13D shape is a declared, frozen axiom whose carriers are forced within an explicitly declared category and given the observed spectrum — a serious, specific minimality case — but its absolute minimality is OPEN. The shape is selected, not forced. This is not a proof of simplicity, and the program is a serious candidate / partial unification, not a validated theory.**

Sources (frozen corpus, this program only): the exact-geometry anti-drift anchor; the SG-1 geometry-specification closure-attack dossier (DECLARED-FROZEN); the SG-2 gauge-recovery closure-attack dossier (DERIVED-GIVEN-E). The full construction, certificates, and freeze records live in Paper I (GUT) on the published site; this explainer recaps only the minimality argument and carries the honest ceiling verbatim.