

SG-3 — Three Generations as a Topological Index $\chi(K_6, \mathcal{E}) = -3$: Per-Gate Closure — Fable_Version rendered package.
 Rendered from DOSSIER_SG3_THREE_GENERATIONS_CLOSURE_ATTACK.md ; frozen technical content unchanged by rendering.

SG-3 — Three Generations as a Topological Index $\chi(K_6, \mathcal{E}) = -3$: Per-Gate Closure-Attack Dossier

What this is. The closure-attack packet for **SG-3 (three chiral generations as a rigid topological index)** on the **frozen 13D K_6 branch ONLY**. It recaps how our geometry closes the family count (concise; the full derivation lives on the published site), verifies where the status actually stands, names every open residual, and lays out the attack plan to push the gate further. It is **not** a rival comparison (that is `BATTLE_GATES/`), **not** the one-line status ledger, **not** a reprint of the manuscript.

Binding discipline (carry verbatim). PROMOTIONS:0. Frozen branch `dcc66f1b2685 / a5b1e6f9d951` READ-ONLY. Honest throughout: the index $\chi(K_6, \mathcal{E}) = -3$ is a **rigid deformation-proof integer with no dial** (Atiyah–Singer), but it is computed **WITH $\mathcal{E}$ (the SM chiral content) as input** — it certifies *this geometry+bundle yields three families*, **not** that three is unique across all geometries; the count is read off a **chosen bundle/weight**, and the SHAPE finding is that **families are bundle-selected on every carrier** (the cheaper $\mathbb{CP}^2$ carrier produces "three by dial"; K_6 's forcedness is bought by **paying +2 dimensions** with a 3-generation–aligned tie-break — the suite's live smuggling surface); and a sharp technical finding the corpus already records: the **literal "spin- $\mathbb{C}$ " wording is REFUTED** for pure-SM content (spin- $\mathbb{C}$ is provably obstructed; Davighi–Gripaios–Lohitsiri) — the genuine forced object is the **twisted** $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$ bundle, which is **FORCED-GIVEN-E** and does **not** force $\mathcal{E}$. Closure paths must be **non-target-loaded** (the κ^3/π kill-test: a proposed axiom counts only if it would be written WITHOUT knowing the answer). **given-E $\neq$ derivation of E**. AXIOM-CLOSED $\neq$ proven; selection $\neq$ derivation; dissolved $\neq$ solved.

0. Header & Verdict

Field	Value
Gate id	SG-3 — Three generations as a topological index $\chi(K_6, \mathcal{E}) = -3$
GUT manuscript gate	Gate 4 — Chirality / no mirrors / family count (§6.4; narrative §5.3; certificate <code>certificates/G04_chirality/</code> ; authority Appendix E). Numbering note (read this): the SG-N scoped ledger labels this " GUT Gate 3 " for index-alignment, but the manuscript's own gate number for chirality/family-count is Gate 4 (§6.4); the manuscript's "Gate 3" (§6.3) is <i>hypercharge / electric charge</i> . This dossier follows the manuscript: every §6.4 / §5.3 / CR4 / Appendix-E citation is to manuscript Gate 4 . The status label travels with the <i>content</i> (the family-count index), not the digit.
Status label (binding)	DERIVED-GIVEN-E
Target anchor(s)	spectrum-E (the observed SM chiral content $\mathcal{E}$; the index $\chi = -3$ is a <i>facet of E</i> , read WITH $\mathcal{E}$ as input — measured-but-irreducible), pinned at the admissibility boundary by LEP/SLD $N_\nu = 2.984 \pm 0.008$. The deeper target " <i>3 across all geometries / bare-carrier-forced</i> " has no independent witness → OPEN ; in-category "no-dial" rigidity is SCHEME-ANCHORED . (Full statement: §A "🌀 Target anchor(s) for this gate".)
Manuscript card status	<i>Claimed certificate pass</i> (§6.4 gate card); recorded as <i>Certificate-complete under declared assumptions</i> (declared index values; frozen parity table) in the §5.3.7 closure block and Appendix E. DERIVED-GIVEN-E is the honest scoped-GUT roll-up of these.
Frozen hashes it rides	Branch <code>dcc66f1b2685</code> / manifest meta <code>a5b1e6f9d951</code> . Spin- C bundle data on K_6 (R1.4) <code>0fd19c9ae0c1</code> ; orbifold parity freeze (the $\mathbb{Z}_2$ fold + boundary parity ledger, R1.3) <code>ac4d2df3e708</code> . Carrier $K_6 = SU(3)/T^2$ and its Euler characteristic $\chi(K_6) = 6 = W(SU(3)) $ are primitives of the frozen branch. No new hash is introduced by this dossier; nothing is mutated.

0.1 Abstract — established / open / what would close it

Established (given- $\mathcal{E}$, given the selected & frozen geometry+bundle). On the frozen branch the family count is the **integer index** read by two complementary counting theorems (Appendix E.1; GP.3): the Borel–Weil–Bott (BWB) index of the frozen line/spinor bundle on the flag manifold $K_6 = SU(3)/T^2$ returns $\chi(K_6, \mathcal{E}) = -3$, and the Atiyah–Patodi–Singer (APS) boundary index on the orbifold $S_Y^1/\mathbb{Z}_2$ returns $(n_L, n_R) = (+3, 0)$. The **magnitude** $|\text{Index}| = 3$ is the number of generations; the **sign** -3 and the one-sidedness $n_R = 0$ are chirality (left-handed families, no surviving mirror). This is a **rigid, deformation-proof integer**: no continuous modulus can move it inside the declared search category — the genuine, defensible content. It is the kind of object a family count *must* be (a winding number, not a volume knob). The gate **creates** the surviving three-family multiplet list consumed downstream by Gates 5/7/9/10.

Open (the honest deductions). (1) **given- $\mathcal{E}$** — the index is computed **with** the SM chiral content $\mathcal{E}$ (the bundle whose Chern class returns -3) supplied as input; it certifies *this geometry+bundle yields three*, not that *three is unique across geometries*, and emphatically not that $\mathcal{E}$ itself is forced (the corpus T3 — "E is forced by anomaly-freedom + minimality" — is **REFUTED**: anomaly-freedom is a filter, not a determiner; infinitely many anomaly-free chiral $U(1)$ extensions exist; the generation number is unfixed; and $\chi(K_6, \mathcal{E}) = -3$ is **circular for forcing $\mathcal{E}$** because $\mathcal{E}$ is its input). (2) **the count is bundle-selected, not carrier-forced** — "3" is read off a **chosen first Chern class / weight**; the cheaper $\mathbb{CP}^2 = SU(3)/U(2)$ carrier produces the count as a **continuous bundle-moduli choice** ("three by

dial"), and K_6 's forcedness is purchased by **paying +2 dimensions** under a tie-break (anti-fitting / "adjustable counts as fail") that is itself **aligned with the 3-generation target** — the SHAPE suite's explicitly named **live smuggling surface**. (3) **bundle admissibility** — *which* spin- $\mathbb{C}$ /line-bundle data are admitted is set by no-fourth-generation + the hypercharge ledger + the $\mathbb{Z}_6$ centre; the index-changing deformation that would give 2 or 4 is excluded by the **LEP N_ν bound**, i.e. by *data*, not by an internal forcing of the bundle. (4) **the literal "spin- $\mathbb{C}$ " wording is REFUTED** for pure SM (Davighi–Gripaios–Lohitsiri arXiv:1910.11277): the genuine global object forced by $\mathcal{E}$'s charges is the **twisted** $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$ ($\mathbb{Z}_6$ via the Tong congruence $q \equiv 3z_2 - 2z_3 \pmod{6}$; $(-1)^F =$ the $SU(2)$ 2π rotation in the centre), **FORCED-GIVEN-E**, not spin- $\mathbb{C}$.

What would close it (the ladder). DERIVED-CLOSED would require a **target-blind** theorem that forces the family count to be **exactly 3 across the declared carrier-space without the bundle as input** — e.g. a bundle-uniqueness / " $\mathcal{E}$ forces the K_6 weight" theorem (the SHAPE color-rung **bundle-admissibility target**), or a proof that the carrier+admissibility data alone (no LEP input, no 3-aligned tie-break) return $|\text{Index}| = 3$ uniquely. The realistic ceiling per residual is **AXIOM-CLOSED**: name one explicit, target-blind posit (e.g. "the admissible bundle is the minimal-weight spin- $\mathbb{C}$ /twisted-spin lift compatible with the $\mathbb{Z}_6$ centre"; "the carrier is the unique purely-abelian-isotropy clean $SU(3)$ flag") that pays the debt in plain sight. The most likely honest outcome on residuals (b)/(c) is **sharper-OPEN**: the family count stays a **rigid index GIVEN a bundle that is selected, not forced**.

0.2 Website source-of-truth (link, don't duplicate)

The full construction, the worked index, the parity table, and the freeze records are the **published manuscript**, Paper I:

- **GUT.html §6.4** (Gate-4 card — chirality / no mirrors / family count; status verbatim) — <https://physics.magflowmeters.com/articles/GUT.html>
- **§5.3** narrative module (Gate 4); **§3.4** (topological-count machinery); the **Strongest-Objections** block (Objection 4: "the three-family count is asserted by choosing the geometry"); **§4.6 / §4.9** (the selector dimension-payment and the anti-fitting firewall); **Appendix E** (Chirality Closure — authoritative: E.1 mechanism, E.2 three generations, E.3 mirror ledger, E.6 certificate); **Appendix CR module CR4** (reader companion); **Appendices GP.3** (the counting theorems — Atiyah–Singer / APS / Borel–Weil–Bott), **B1/B2** (selector / category-relativity), **C2** (K_6 dossier), **C4** ($S_Y^1/\mathbb{Z}_2$ dossier), **A2.2** (bundle ledger); **Appendix Ro** freeze records. Same public article.
- Downstream cross-references (do not change this gate): TOE.html — <https://physics.magflowmeters.com/articles/TOE.html>.

This dossier recaps only what is needed to attack; the site controls all common material.

1. How OUR geometry closes the family count — the closure claim (concise)

Full derivation: **GUT.html §6.4 + §5.3 + Appendix E + GP.3 + CR4**. This section is the attack-grade recap, not the derivation.

1.1 The mechanism end-to-end

The family count is generated by **two index computations on the frozen geometry**, one per half of the gate (" K_6 counts the families; $S_Y^1/\mathbb{Z}_2$ makes them chiral by projecting out the mirror copy" — CR4.6). The pipeline is one forward chain (§5.3, Appendix E.1):

```
K6 = SU(3)/T2 + frozen spin-C/twisted-spin bundle --BWB index--> chi(K6,E) = -3 --> |Index| = 3 families
S1_Y/Z2 fold + frozen boundary parity ledger --APS index--> (n_L, n_R) = (+3, 0) --> mirrors removed
E_matter (chiral actors) + P_chi projector --zero modes--> three left-handed SM generations
```

The load-bearing factors, named so a reader can attack each:

- The carrier $K_6 = SU(3)/T^2$** (R1.4; dossier C2; hash of the bundle data `0fd19c9ae0c1`). The complete flag manifold of $\mathbb{C}^3$: $\dim = 8 - 2 = 6$, a compact homogeneous Kähler manifold with isometry group $SU(3)$ and Euler characteristic $\chi(K_6) = |W(SU(3))| = 6$. It is the **same** factor Gate 2 used for color recovery — the strongest reuse point in the manuscript (color carrier $\rightarrow$ family-count carrier, no new term). Line bundles on K_6 are classified by weights; their cohomology — hence the particle content they induce — is computed in **closed form** by Borel–Weil–Bott (GP.3). (Attack handle: the *carrier* is forced as the unique clean abelian-isotropy $SU(3)$ flag, but the *bundle/weight* read on it is **selected** — residuals (b)/(c).)
- The spin- $\mathbb{C}$ / twisted-spin bundle and its BWB index.** Kähler manifolds such as K_6 are canonically spin- $\mathbb{C}$; on a homogeneous Kähler space the relevant index is computable in closed form, and the frozen bundle returns $\chi(K_6, \mathcal{E}) = -3$, so $|\text{Index}| = 3$ (Appendix E.1; full bundle data A2.2; GP.3 Borel–Weil–Bott). **The sign records chirality; the magnitude records the family count.** No additional family appears (the index returns no further zero mode under the same bundle); no fewer (the integer is robust against continuous deformation **within the declared category**). (Attack handle: the *word* "spin- $\mathbb{C}$ " is technically refuted for pure SM — the genuine object is twisted ($\text{Spin} \times G_{\text{SM}}/\mathbb{Z}_6$; residual R4. The **value** -3 is robust; the **labelling of the structure** is the soft spot.)
- The no-mirror fold $S_Y^1/\mathbb{Z}_2$ and its APS index.** A *closed* factor is handedness-neutral (fact F2: left/right come in matched pairs, net handed count zero — a bare S_Y^1 mirrors every fermion, APS index $(+3, +3)$). The repair is an **edge**: the orbifold $\theta \mapsto -\theta$ with fixed points $\{0, \pi\}$ (freeze R1.3 `ac4d2df3e708`) admits one handedness and refuses its mirror. Under the frozen parities the APS boundary index returns $(n_L, n_R) = (+3, 0)$ — a one-sided count impossible on any closed factor. (Attack handle: a *no-go* leg; its force is structural — F2 plus the fold — and it is the cleaner half of the gate.)

4. **The chiral actors $\mathcal{E}_{\text{matter}}$ and the projector $P_\chi = \frac{1}{2}(1 + \gamma_5 \Gamma_8)$** (Appendix C7; P_χ derived from `ac4d2df3e708 + 0fd19c9ae0c1`). The matter bundle is the $\otimes$ -layer object on which P_χ and the index act; without it neither has a domain. The surviving content is exactly the SM set: the right-handed singlets (u_R, d_R, e_R , and the neutrino-sector mode) arise from the conjugate sector through the chamber projectors with the same generation count.
5. **The generation module $\mathcal{G}_{\text{gen}} = \text{span}\{g_1, g_2, g_3\}$, $\text{dim} = 3$** , matched to the family index -3 (A1.13; consumed by the F^+ flavor chamber, SG-8). SG-8 **inherits** this dimension; SG-3 is the gate that supplies "3" to flavor.

1.2 What "closed" means here, and its conditionality

"Closed" for SG-3 is **a rigid integer family-count mechanism under declared assumptions** — strictly **not** a from-nothing inevitability of three, **not** a proof that three is unique across all geometries, **not** a derivation of the SM chiral content $\mathcal{E}$. It is conditional on:

- **given- $\mathcal{E}$** — the SM chiral content (the bundle $\mathcal{E}$ whose Chern class returns -3 , and the hypercharge ledger that fixes admissibility) is supplied as input; SG-3 does **not** derive $\mathcal{E}$. (The index is the *output of a computation whose input is $\mathcal{E}$* — using it to "force" $\mathcal{E}$ would be circular; the corpus says so explicitly, REDUCTION_SPINTWIST_C §3(iii).)
- **given-the-selected-geometry+bundle** — the frozen carrier K_6 and the *selected* bundle/weight; the carrier is forced (abelian-isotropy uniqueness; CP^2 over-produces gauge at Gate 2), but the **bundle is selected inside a declared admissibility class**, not proven unique.
- **given-the-data-exclusion** — the deformation that would change the index to 2 or 4 is excluded by the **LEP N_ν bound** (2.984 ± 0.008 light species). The *count's robustness against fourth-generation rescue is partly a data input*, not a pure internal forcing.

The genuine, defensible content (the part that survives the honest accounting):

Genuine output (rigid-index prediction)	Mechanism
Family count = 3 as a deformation-proof integer (no dial)	$ \chi(K_6, \mathcal{E}) = 3$ from BWB; invariant under continuous moduli in-category
Chirality (left-handed families; the sign of -3)	orientation/chirality convention of the index + the one-sided APS count
No surviving mirror ($n_R = 0$)	APS one-sided boundary index on $\mathcal{S}_Y^1/\mathbb{Z}_2$ (F2 + the fold)
The integer is the right kind of number	topology, not a tuned bundle modulus ("three by dial" fails even at 3)

The conditional / non-genuine content (selection in disguise, or data-fed):

Claim	Honest reality
"Three families are derived "	given-$\mathcal{E}$ — the bundle that returns −3 is $\mathcal{E}$; not a from-nothing derivation
"Three is unique across geometries "	NOT claimed — only <i>this</i> geometry+bundle is certified; cross-geometry uniqueness is OPEN
" K_6 is forced to carry exactly −3 "	the carrier is forced; the bundle/weight is selected in a declared admissibility class (residuals b, c)
"the spin-$\mathbb{C}$ index"	REFUTED wording for pure SM — the forced object is twisted $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$ (R4)
"no fourth family, by topology"	the <i>robustness in-category</i> is topological; the <i>exclusion of the 2/4 deformation</i> leans on the LEP N_ν bound

So the precise statement of closure: **the geometry genuinely makes the family count a rigid deformation-proof integer of the correct kind (3, left-handed, no mirror) — but it does so by reading a bundle whose data is $\mathcal{E}$ itself; it certifies this geometry+bundle, not that three is forced across geometries, and not that the bundle is unique. given-E is the load-bearing qualifier of the whole gate.**

2. Verify the status — is DERIVED-GIVEN-E real?

This is the verification a skeptic would run. Each witness gets a grade — **hand-checkable / symbolic / machine-lane** — and an honest **reproduces?** flag.

2.1 The witness ledger

#	Witness	What it asserts	Grade	Reproduces?
W1	BWB index $\chi(K_6, \mathcal{E}) = -3$ (Appendix E.1; A2.2; GP.3)	family count $= -3 = 3$, a rigid integer	symbolic	Yes (given the bundle). Borel–Weil–Bott on G/T is a closed-form rep-theory computation; the value -3 is the frozen R1.4 datum (<code>0fd19c9ae0c1</code>). It reproduces as a function of the declared bundle ; it does not reproduce without the bundle as input (that is the given- $\mathcal{E}$ caveat made concrete).
W2	APS index $(n_L, n_R) = (+3, 0)$ on $S_Y^1/\mathbb{Z}_2$ (Appendix E.1/E.3; A1.8; GP.3)	mirror sector removed; one- sided count	hand- checkable (toy) / symbolic (full)	Yes. The one-line interval toy (Appendix T.3: even left-mode survives both walls, odd right-mode vanishes) is hand-checkable; the full count reads three rows of the frozen parity table <code>ac4d2df3e708</code> . Bare- S_Y^1 control returns $(+3, +3)$ — the fold is load-bearing.
W3	Euler characteristic $\chi(K_6) = 6 = W(SU(3)) $	the carrier is the $SU(3)$ flag (geometric sanity)	hand- checkable	Yes — $ W(SU(3)) = S_3 = 6$; dim $K_6 = 8 - 2 = 6$.
W4	Zero-weight multiplicity rule $m_0(p, q) = \min(p, q) + 1$ if $(p - q) \equiv 0 \bmod 3$, else 0	the rep-theory engine under the index is reproducible	machine-lane	Yes — reproduced target-blind (SHAPE_M0_DELTASPINC §1.1; closure-campaign batch1 "m ₀ (p,q) re-derived from scratch via Freudenthal"). All 8 table values regenerate.
W5	$\mathbb{Z}_6$ centre kernel via the Tong congruence $q \equiv 3z_2 - 2z_3 \bmod 6$	the centre acting trivially on $\mathcal{E}$ is exactly $\mathbb{Z}_6$	machine-lane	Yes — computed field-by-field from the actual SM hypercharges (REDUCTION_SPINTWIST_C §2; six elements listed; verified for $Q_L, u_R, d_R, L_L, e_R, H$). Not reverse-engineered (passes the κ^3/π anti-smuggle test).
W6	Spin-$\mathbb{C}$ obstruction for pure SM (Davighi–Gripaios–Lohitsiri arXiv:1910.11277 Sec.7)	no $U(1) \subset G_{\text{SM}}$ gives all-odd Weyl charges $\Rightarrow$ pure SM has no spin- $\mathbb{C}$ without an extra gauged $U(1)$	machine-lane	Yes — reproduced (empty all-odd- $U(1)$ scan; test-the-test: adding $B-L$ makes the scan succeed $\Rightarrow$ non-vacuous). This refutes the literal "spin-$\mathbb{C}$" label while leaving $ \chi = 3$ intact.
W7	Selector elimination of $\mathbb{CP}^2$ ("three by dial") (§5.3; §4.6; CR4.9)	the cheaper 4D $SU(3)$ carrier fails because its family count is a continuous bundle-moduli choice	symbolic/audit	Conditional. The <i>elimination</i> reproduces ($\mathbb{CP}^2$ family count is moduli-dependent); but it relies on the anti-fitting tie-break that "adjustable counts as fail," which is a programme criterion aligned with the 3-target (residual R2). The elimination is honest only if anti-fitting is a genuine <i>prior</i> pass/fail, not inserted to reach 3.
W8	Freeze hashes (<code>0fd19c9ae0c1</code> bundle; <code>ac4d2df3e708</code> parity) + meta <code>a5b1e6f9d951</code>	every Gate-4 object content-	machine-lane	Yes in principle (re-hash R1.3/R1.4 rows; recompute meta-hash); not re-run here.

#	Witness	What it asserts	Grade	Reproduces?
		addressed before comparison		
W9	Machine certificate certificates/G04_chirality/	parity-table lint + index-consistency against declared index values	machine-lane	AUDIT — the certificate checks consistency against the declared -3 and $(+3, 0)$, not an independent re-derivation of the index from scratch; it is a lint, not a from-nothing proof. Honest: " <i>under those declared index values</i> " (CR4.11).

2.2 What reproduces, plainly

- **The integer reproduces, and its rigidity is real.** $|\chi(K_6, \mathcal{E})| = 3$ is a Borel–Weil–Bott output — a closed-form rep-theory computation on $SU(3)/T^2$, not a numerical fit. The APS $(+3, 0)$ is hand-checkable via the one-line interval toy. The bare- S_Y^1 control $(+3, +3)$ confirms the fold is load-bearing. **This is the strongest, most defensible leg of the gate** — and it is genuinely stronger than SG-8's strongest leg, because an integer index is *more rigid* than a κ -power ratio.
- **The rep-theory engine reproduces target-blind.** The $m_0(p, q)$ multiplicity rule and the $\mathbb{Z}_6$ Tong congruence both regenerate from scratch (W4/W5), computed from the actual SM charges, not reverse-engineered.
- **The deformation-proof claim reproduces in-category.** No continuous modulus moves the index inside the declared search category — that is what "topological" means here (CR4.4).

2.3 What does NOT yet reproduce (honest gaps in the status)

- **The index does NOT reproduce without the bundle as input.** $\chi(K_6, \mathcal{E}) = -3$ is a function of $\mathcal{E}$ (the first Chern class / weight). Feed a different admissible weight and the index changes. So "three families" reproduces **given- $\mathcal{E}$** , and the gate's certificate is honest exactly to that conditioning — it is *not* a from-nothing reproduction of "3."
- **The literal "spin- $\mathbb{C}$ " structure does NOT reproduce — it is refuted.** W6 shows pure-SM content **provably forbids** a spin- $\mathbb{C}$ structure; the genuine forced object is the **twisted** $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$ (twist order $n = 2$, $(-1)^F$ = the $SU(2)$ 2π rotation in the centre; REDUCTION_SPINTWIST_C §3). The **value** $|\chi| = 3$ survives this correction (it is a count, not a structure label), but the manuscript's wording "spin- $\mathbb{C}$ index" is a structural overstatement that the corpus itself flags as a smuggle if asserted literally.
- **The "three is forced" reading does NOT reproduce across geometries.** The cross-geometry uniqueness of 3 is **OPEN**; the $\mathbb{CP}^2$ "three by dial" elimination depends on a 3-aligned tie-break (W7); the bundle is selected, not proven unique (residuals b, c).
- **The " $\mathcal{E}$ is forced" reading is REFUTED.** Corpus T3 ("E forced by anomaly-freedom + minimality") is explicitly refuted (REDUCTION_SPINTWIST_C §3(iii)): anomaly-freedom is a *filter*, the hypercharge family is 1-parameter, there are infinitely many anomaly-free chiral $U(1)$ extensions (Allanach et al.), the generation number is unfixed, and $\chi = -3$ is circular for forcing $\mathcal{E}$. **E remains a residual SHAPE primitive.**

2.4 Why DERIVED-GIVEN-E (not higher, not lower)

- **Not FORCED / DERIVED-CLOSED:** the count is a *rigid integer*, but it is read off a **selected bundle** with $\mathcal{E}$ as input; cross-geometry uniqueness is unproven; the carrier's forcedness over $\mathbb{CP}^2$ uses a 3-aligned tie-break; and $\mathcal{E}$ is not derived. Too much conditioning for an unconditional "derived."
- **Not PARTIAL (the SG-4..SG-9 tier):** unlike flavor (fitted sector scales) or thresholds (injected unreproduced δ 's), SG-3's central object **is** a rigid deformation-proof integer that **reproduces** as a closed-form index, with a clean no-go (APS one-sidedness) for the mirror half. There is no fitted continuous parameter inside the count itself. It sits **above** the PARTIAL gates.
- **DERIVED-GIVEN-E is exactly right:** a rigid index (the strong, genuine content) computed **with $\mathcal{E}$ as input** (the honest qualifier), tied with SG-2 as the two strongest gates in the ledger. This matches the SCOPED_GUT ledger SG-3 line verbatim: *"a rigid integer, no dial (Atiyah–Singer) — but computed WITH E as input: it certifies this geometry+bundle yields three families, not that three is unique across all geometries."*

3. The attack surface — what to attack (ranked by leverage)

Every open residual as a named object with its precise obstruction. **Leverage** = how much closing it moves the gate (and whether it removes the given- $\mathcal{E}$ conditioning).

3.1 The residual register

ID	Residual (named object)	Precise obstruction	Status	Leverage
R1	given- $\mathcal{E}$: "3" is computed WITH $\mathcal{E}$ as input	$\chi(K_6, \mathcal{E}) = -3$ takes the bundle $\mathcal{E}$ (the SM chiral content / first Chern class) as input; it certifies <i>this geometry+bundle</i> , not that 3 is unique across geometries , and cannot force $\mathcal{E}$ (circular). Corpus T3 REFUTED — anomaly-freedom is a filter, $\mathcal{E}$ stays a SHAPE primitive.	OPEN (the meta-residual; the binding qualifier of the whole gate)	HIGHEST — it conditions every "derived" reading; removing it would upgrade the gate to FORCED
R2	the count is bundle-selected, not bare-carrier-forced (the SHAPE color-rung finding)	"3" is read off a chosen weight ; $\mathbb{CP}^2 = SU(3)/U(2)$ produces the count as a continuous bundle-modulus ("three by dial"); K_6 's forcedness is bought by paying +2 dimensions under a tie-break aligned with the 3-generation target (REVIEW_SHAPE_SUITE defect 4 — the suite's live smuggling surface ; "families bundle-selected on every carrier").	OPEN	HIGH — if the carrier+admissibility data alone forced 3 (no 3-aligned tie-break), the gate would stop leaning on a target-shaped criterion
R3	bundle admissibility is data-/centre-fed, not internally forced	<i>Which</i> bundle is admitted is set by no-fourth-generation + the hypercharge ledger + the $\mathbb{Z}_6$ centre; the index-changing deformation to 2/4 is excluded by the LEP N_ν bound (a data input). No internal theorem says "the admissible weight is unique."	OPEN	HIGH — closing it (a bundle-uniqueness theorem) is the one path that could remove R1's conditioning
R4	"spin-C" wording REFUTED for pure SM	Davighi–Gripaios–Lohitsiri: no $U(1) \subset G_{\text{SM}}$ gives all-odd Weyl charges $\Rightarrow$ pure SM has no spin-C structure; the genuine forced object is twisted $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$ ($n = 2$). Asserting "spin-C" literally is a smuggle (claims a structure the content forbids).	DISCLOSED (refutation on-disk; manuscript wording not yet corrected)	MEDIUM — a wording/structure correction, not a value change ($ \chi = 3$ survives); honesty + technical-correctness gain
R5	the index-changing deformation excluded by data, not topology	The robustness of -3 is <i>in-category</i> ; the exclusion of the neighboring index values (2, 4) that <i>would</i> arise from an index-changing bundle deformation is the LEP N_ν bound (C2/C4/Appendix E). The "no dial" claim is therefore topological inside the admissibility class and empirical at the class boundary.	DISCLOSED	MEDIUM — keeps the "no dial" claim correctly scoped (deformation-proof in-category, data-pinned at the boundary)
R6	certificate G_{04} chirality is a consistency lint, not a from-scratch index re-derivation	The certificate checks survivors = SM set, mirrors projected, counts consistent against the declared index values $-3, (+3, 0)$ — <i>under those declared values</i> (CR4.11). It does not independently regenerate the index from the bundle data target-blind.	AUDIT	MEDIUM — converts "asserted index" to "machine-checked index" with no new physics
R7	sign/chirality convention bit ($\text{Pin}^\pm$) is unpinned	The magnitude 3 is robust; the sign (chirality orientation) rides the same $\text{Pin}^+/\text{Pin}^-$ / convention bit that leaves BG-10's $e^{\pm i\pi/4}$ unfixed (SHAPE_M0_DELTASPINC §4; HANDOFF_SPECIALIST_3). The sign of -3 as "left-handed not right-handed" is convention-	DISCLOSED (convention bit named)	LOW-MEDIUM — affects the chirality <i>label</i> , not the family count; shared with BG-10/SHAPE

ID	Residual (named object)	Precise obstruction	Status	Leverage
		stated, not independently pinned by the frozen record.		
R8	inheritance: SG-8 flavor and the threshold vector depend on "3"	SG-8's generation module $\dim \mathcal{G}_{\text{gen}} = 3$ is <i>inherited</i> from this index; the SG-7 threshold δ -vector is load-bearing on the family count -3 (Appendix E "if the count were -2 or -4 the column sums would not reproduce $(+4.8424, -3.1112, -1.7313)$ "). A revision of "3" cascades to Gates 5/7/9/10.	DISCLOSED (dependency wired)	LOW — not unique to SG-3; records the blast radius, no internal action

3.2 Leverage ranking (attack order)

1. **R3** (bundle admissibility / uniqueness) — the **single highest-value target**: a target-blind bundle-uniqueness theorem is the only route that could *remove* R1's given- $\mathcal{E}$ conditioning. This is the SHAPE color-rung **bundle-admissibility target**.
2. **R1** (given- $\mathcal{E}$) — conditions every "derived" reading; addressed *only* via R3 (or refuted as unreachable, which is the honest expectation).
3. **R2** (bundle-selected vs carrier-forced; the 3-aligned tie-break) — the live smuggling surface; closing it means forcing 3 without a target-aligned criterion.
4. **R4** (spin- $\mathbb{C}$ wording) — cheapest *correctness* win (refutation already on-disk; correct the structure label to twisted-spin/ $\mathbb{Z}_6$).
5. **R6** (certificate lint $\rightarrow$ machine-checked) — cheapest *reproducibility* win.
6. **R5 / R7 / R8** — disclosed scoping / convention / inheritance items.

The cardinal honest point. R1/R2/R3 are where the "given-E" lives. If a closure path *introduces a new target-aligned criterion* to force 3 (e.g. "pick the weight that gives three"), it has **relocated** the input, not removed it — the κ^3/π kill-test applies in full force here, and the SHAPE suite already flags the +2-dimension forcedness override as the live smuggle. R4/R5/R6/R7 are the *non-count* residuals (structure label + scoping + reproducibility + convention); closing them improves correctness and machine-reality but **does not remove the given- $\mathcal{E}$ conditioning**.

4. THE ATTACK PLAN — closure paths (the core)

For each residual: the **technique** (closure playbook T1 axiom-floor / T4 no-go / T5 firewall / T6 all-operator-conditional / T7 eliminative / T10 selector), the **named axiom it could reduce to** (stated so it would be written WITHOUT the answer — the κ^3/π kill-test), the **specialist target** (theorem to hand off) or the **owner artifact** (computation/ruling) needed, the **math to attempt**, and the **success ladder** (DERIVED-CLOSED rare $\rightarrow$ AXIOM-CLOSED likely $\rightarrow$ sharper-OPEN $\rightarrow$ REFUTED). Dispositions are kept conservative per the closure-campaign norm (default to the more conservative verdict; demote on skeptic-verify).

4.1 R3 — Bundle admissibility / "does $\mathcal{E}$ force the K_6 weight?" (the highest-value target)

Why first. Of all SG-3 residuals, **only R3** can in principle remove the given- $\mathcal{E}$ conditioning. If the admissible bundle on K_6 is **unique** (forced by the carrier + the $\mathbb{Z}_6$ centre + the hypercharge ledger, with no 3-aligned tie-break and no LEP input), then "3" is forced by the geometry+admissibility *alone*, and SG-3 moves from DERIVED-GIVEN-E toward FORCED. This is the SHAPE color-rung **bundle-admissibility target** named in the prompt's closure guidance.

Technique: T10 (selector) + T4 (no-go on alternatives). Run the admissibility selector **target-blind** over the full space of K_6 line/spinor weights compatible with the $\mathbb{Z}_6$ centre and the hypercharge ledger, and ask: is the weight returning $|\chi| = 3$ the **unique** admissible one — or merely one admissible one selected because it gives three?

Named axiom it could reduce to (κ^3/π -clean). The honest target-blind posit references only the carrier + centre + minimality, never the number 3:

AXIOM-MIN-WEIGHT-LIFT (target-blind form): *"the admissible bundle on K_6 is the minimal-weight twisted-spin ($(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$) lift compatible with the $\mathbb{Z}_6$ centre and the hypercharge ledger; the family count is its Borel–Weil–Bott index."*

This is writable with **no** flavor/generation number in sight — it is a statement about *which bundle is admitted*, not what its index equals. It **PASSES** the κ^3/π kill-test *as a statement*. The test is then whether it **yields** $|\chi| = 3$ without the 3-aligned tie-break. **Honest caution:** the corpus already records the nearby A5-actor (SHAPE Lemma 3) attempt — "is $\mathcal{E}_{\text{matter}} \oplus \dots$ the minimal actor realization?" — **DEMOTED to OPEN** because its success criterion is a **universal negative** ("no admissible competitor supplies a lower-cost realization") that is **UNMET** (competitor matrix all "Unknown"; NCG un-scored). R3 inherits this hazard: a uniqueness theorem requires ruling out *all* admissible weights, a universal negative.

Specialist target (to hand off). A **bundle-uniqueness theorem**: "Given $K_6 = SU(3)/T^2$, the $\mathbb{Z}_6$ centre $\ker(\text{centre} \rightarrow \text{Aut}(\mathcal{E}))$, and the hypercharge admissibility ledger — **and no 3-generation criterion** — the minimal-weight admissible twisted-spin lift is unique and its BWB index has magnitude 3." If proven, R3 closes, R2's tie-break dependence falls, and R1's conditioning lifts.

Math to attempt. (1) Enumerate the K_6 weights compatible with the $\mathbb{Z}_6$ centre (the Tong-congruence lattice $q \equiv 3z_2 - 2z_3 \pmod{6}$). (2) Apply the BWB index to each; record the family-count map weight $\mapsto \chi$. (3) Check whether minimality (lowest weight) selects a **unique** weight, and whether that weight's index is **3 without** invoking "three." (4) Sensitivity: does any neighboring admissible weight also give 3 (degenerate) or give 2/4 (excluded only by LEP)? If 2/4 are excluded only by LEP, R3 stays open and R5 is confirmed.

Success ladder. - DERIVED-CLOSED: minimal admissible weight proven unique **and** its index forced to 3 with no 3-aligned tie-break, no LEP input. **Unlikely** — the universal-negative obligation (rule out all weights) mirrors the demoted A5-actor Lemma 3. - **AXIOM-CLOSED (realistic ceiling):** AXIOM-MIN-WEIGHT-LIFT named; the carrier+centre forced, the *minimality* of the weight posited as a target-blind axiom that is **checked to yield 3**. Honest endpoint if the check passes. - sharper-OPEN:

minimality does not single out a unique weight, or excludes 2/4 only via LEP $\Rightarrow$ "3" stays **rigid-given-a-selected-bundle**, with the selection conceded. **The expected outcome.** - REFUTED: a strictly-lower-cost admissible weight gives $\chi \neq -3 \Rightarrow$ the minimal-weight axiom is false and the bundle was 3-selected (the good kind of negative — a structure-first datum departing from the target).

Honest disposition: OPEN, with one concrete bounded computation (the weight-enumeration + index map). The realistic ceiling is AXIOM-CLOSED (AXIOM-MIN-WEIGHT-LIFT) or sharper-OPEN; DERIVED-CLOSED is gated on an unreachable universal negative. Do **not** bank any "minimality" claim that secretly uses "it gives three" as the selector — that is the κ^3/π failure mode and the SHAPE-suite live smuggle.

4.2 R1 — given- $\mathcal{E}$ ("3" is computed with $\mathcal{E}$ as input)

The obstruction stated precisely. $\chi(K_6, \mathcal{E}) = -3$ is the index of a bundle $\mathcal{E}$ that **is** the SM chiral content. So the index **certifies** "this geometry+bundle yields three," but it **cannot force** $\mathcal{E}$ (using $\chi = -3$ to force $\mathcal{E}$ is circular — $\mathcal{E}$ is its input), and it does **not** establish that three is unique across geometries. This is the binding qualifier of the entire gate.

Technique: T1 (axiom-floor) — name the residual cleanly; closure is via R3 or not at all. R1 is not independently closeable: it dissolves **only** if R3 (bundle uniqueness) closes (then "3" is forced by geometry+admissibility, not read off $\mathcal{E}$). Absent R3, the honest move is to **name** the floor:

AXIOM-CONTENT-GIVEN (κ^3/π -clean): *"the SM chiral content $\mathcal{E}$ (the matter bundle and its hypercharge ledger) is a SHAPE primitive supplied as input; the geometry computes the family count of $\mathcal{E}$, it does not derive $\mathcal{E}$."*

This is writable with no generation number — it is a statement about *what is input*, not its value. It is essentially **already where the corpus stands** (REDUCTION_SPINTWIST_C §3(iii): "E remains a residual SHAPE primitive, irreducible-under-known-reductions"; REVIEW_SHAPE_SUITE: the suite **bottoms on E**, "the un-forced core," and the residue-mislabel-to-fix is precisely that the suite bottoms on "Shape" instead of E).

Specialist target / owner ruling. None new at the math level — R1's only genuine math route **is** R3. The owner artifact is a **labelling decision**: ensure every "three families derived" headline carries the given- $\mathcal{E}$ qualifier (the ledger already does; the manuscript §6.4 card says "Claimed certificate pass," and Objection 4 answers "the index is a topological integer ... routed through the selector, not hand-picked," with the honest residue named — "minimality holds only *inside* the declared search category").

Why this is NOT a relocation. Naming AXIOM-CONTENT-GIVEN does not turn one hard problem into three harder ones; it records the debt in plain sight (the debt = derive $\mathcal{E}$, which is **out of scope** of scoped-GUT, GAP-class, Gate 11). It is the conservative, honest floor.

Success ladder. - DERIVED-CLOSED: only if R3 closes (3 forced by geometry+admissibility).

Unlikely. - AXIOM-CLOSED (where it already sits): AXIOM-CONTENT-GIVEN named; $\mathcal{E}$ conceded as a SHAPE primitive; the gate honestly DERIVED-GIVEN-E. **This is the realistic and current endpoint.** - sharper-OPEN: if R3's enumeration shows the bundle is degenerate/data-pinned

⇒ the conditioning is *sharper* but unremoved. - REFUTED: not applicable (R1 is a scoping fact, not a falsifiable claim).

Honest disposition: AXIOM-CLOSED at AXIOM-CONTENT-GIVEN; the gate is correctly DERIVED-GIVEN-E and stays there unless R3 closes. given-E ≠ derivation of E — this is the load-bearing sentence of the whole dossier.

4.3 R2 — Bundle-selected vs bare-carrier-forced (the SHAPE color-rung finding / the live smuggle)

The obstruction. "3" is read off a **chosen weight**. On the cheaper carrier $\mathbb{CP}^2 = SU(3)/U(2)$ the family count is a **continuous bundle-moduli choice** — "three by dial." K_6 is preferred because its count is a rigid integer, but that preference is purchased by **paying +2 dimensions** ($\mathbb{CP}^2$ is 4D, K_6 is 6D) under a tie-break — anti-fitting / "adjustable counts as fail" — that is **target-aligned with the 3-generation count** (REVIEW_SHAPE_SUITE defect 4: this **overrides the declared lower-dimension Occam rule** with a 3-aligned criterion; the suite's **live smuggling surface to watch**).

Technique: T5 (no-target-loading firewall) + T10 (selector). The firewall question: is the anti-fitting tie-break a genuine **prior** pass/fail constraint (it was written into the selector *before* and *independently of* the 3-generation target), or was it inserted to make K_6 beat $\mathbb{CP}^2$ (i.e. to reach 3)?

The firewall, datum by datum:

Selection datum	Selection record	Target-blind?	Firewall verdict
Carrier $K_6 = SU(3)/T^2$ over $\mathbb{CP}^2$	abelian-isotropy uniqueness (T^2 unique purely-abelian $SU(3)$ isotropy; $\mathbb{CP}^2$ over-produces gauge at Gate 2)	yes at Gate 2 (gauge over-production is data-independent)	PASS — K_6 is the unique clean carrier <i>for gauge</i> , independently of family count
The +2-dimension payment "for forcedness" at Gate 4	§3.5 / §4.6: "selector pays two extra dimensions for K_6 's forced count"	SUSPECT — overrides the lower-dimension Occam rule with anti-fitting	FLAGGED — target-aligned criterion; the live smuggle
The anti-fitting firewall ("adjustable counts as fail", §4.9)	a declared <i>prior</i> pass/fail predicate (count must be a deformation-proof integer)	borderline — it is a <i>type</i> predicate (kind-of-number), not a value predicate	PASS-as-type / SUSPECT-as-tiebreak
The bundle/weight returning −3	selected in the admissibility class	NO at the weight level (R3)	FAIL — the weight is selected (this is R3)

The decisive point. The carrier choice K_6 -over- $\mathbb{CP}^2$ is **legitimately forced for gauge** (Gate 2, abelian-isotropy uniqueness — a data-independent result). The *family-count* preference, however, leans on the anti-fitting tie-break, and **the corpus itself flags this as the live smuggle**. The honest split:

AXIOM-COUNT-MUST-BE-INDEX (κ^3/π -clean): "an admissible family count must be a deformation-proof topological integer (an index), not a continuous bundle modulus."

This is writable with no value — it is a **type** predicate ("the count must be the *kind* of number an index is"), not "the count is 3." It **PASSES** the kill-test as a statement. It legitimately kills $\mathbb{CP}^2$'s tunable count **regardless of the value** — a $\mathbb{CP}^2$ tuned to *four* fails it just as a $\mathbb{CP}^2$ tuned to three does. **This is the genuine, non-target-loaded content of the tie-break**, and stating it this way *defuses* the smuggle: the criterion rejects "three by dial" because it is a *dial*, not because it is *three*.

Specialist target. Prove that **every** clean rigid-integer-count $SU(3)$ carrier of dimension < 6 either over-produces gauge (like $\mathbb{CP}^2$) or has a moduli-dependent count $\Rightarrow K_6$ is the minimal-dimension carrier with a deformation-proof family index, **independently of the value 3**. If proven, R2 closes at the type level (the +2-dimension payment is justified by "rigid-integer-count" not "three").

Math to attempt. (1) Compute the family-count map on $\mathbb{CP}^2 = SU(3)/U(2)$ explicitly and confirm it is a continuous bundle-modulus (not an index). (2) Verify AXIOM-COUNT-MUST-BE-INDEX rejects it *for being a dial*, value held abstract. (3) Survey the sub-6D $SU(3)$ carrier list for any rigid-integer-count survivor; confirm K_6 is the minimal one. (4) **Kill-test:** re-run the selector with the 3-generation target *masked* — does the +2-dimension payment still go through on "rigid-integer-count" alone? If yes, the smuggle is defused; if no, it is confirmed.

Success ladder. - DERIVED-CLOSED: K_6 proven the unique minimal-dim rigid-integer-count carrier, value-blind $\Rightarrow$ the +2-dimension payment is justified by *type*, not *target*; smuggle defused. **Possible (bounded computation).** - AXIOM-CLOSED (likely): AXIOM-COUNT-MUST-BE-INDEX named as the target-blind tie-break; the +2-dimension payment conceded as buying *rigidity* (kind-of-number), not *the value 3*. - sharper-OPEN: the masked-target selector run shows the payment only goes through *because* it yields 3 $\Rightarrow$ the smuggle is confirmed and R2 stays open (the honest negative the SHAPE suite warns about). - REFUTED: a sub-6D carrier has a rigid-integer count $\neq 3$ that the selector should have preferred $\Rightarrow$ the carrier choice was 3-selected.

Honest disposition: OPEN, with the live smuggle named and a clean target-blind type-axiom (AXIOM-COUNT-MUST-BE-INDEX) that would defuse it — pending the masked-target selector run. The carrier-over- $\mathbb{CP}^2$ choice is legitimately forced **for gauge** (Gate 2); the family-count tie-break must be re-run target-blind to confirm it is "rigid-integer" not "three."

4.4 R4 — The "spin- $\mathbb{C}$ " wording is REFUTED for pure SM (the cleanest correctness win)

The obstruction. The manuscript (§6.4, Appendix E.1, CR4.5, GP.3) calls the family-count object the **spin- $\mathbb{C}$ Borel–Weil–Bott index**. But pure SM content **provably forbids** a spin- $\mathbb{C}$ structure: no $U(1) \subset G_{SM}$ gives all-odd Weyl charges (Davighi–Gripaios–Lohitsiri arXiv:1910.11277 Sec.7; reproduced target-blind with a non-vacuous scanner — adding $B-L$ makes the all-odd scan succeed). The genuine global object forced by $\mathcal{E}$'s charges is the **twisted** $(Spin \times G_{SM})/\mathbb{Z}_6$ (twist order $n = 2$; $(-1)^F$ = the $SU(2) 2\pi$ rotation identified with the internal centre; $\mathbb{Z}_6$ via the Tong congruence). Asserting "spin- $\mathbb{C}$ " literally would itself be a **smuggle** (claiming a structure the content forbids).

Technique: T4 (formalize the computed no-go) — pure correctness, no new physics. The refutation is **already computed and on-disk** (REDUCTION_SPINTWIST_C; SHAPE_Mo_DELTASPINC §2). The closure is a **wording/structure correction**, not a value change.

Named result it reduces to (κ^3/π -clean):

AXIOM-TWISTED-SPIN-GIVEN-E (target-blind, in fact a computed reduction): *"the global spin structure on the active branch is the twisted $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$ forced by $\mathcal{E}$'s charges mod 2; the family count is the index of the associated twisted-Dirac operator on K_6 , $|\text{Index}| = 3$; spin- $\mathbb{C}$ is OBSTRUCTED for pure SM."*

This carries no generation number; it references only the SM charges, the centre, and Dai–Freed single-valuedness. It is **FORCED-GIVEN-E** (REDUCTION_SPINTWIST_C falsifier table: all three conditions PASS — spin-from-E, lands-on-E, E-stays-primitive). **Crucially: the value $|\chi| = 3$ is unchanged by this correction** — a count is not a structure label.

Owner artifact needed. A manuscript wording patch (owner-gated; PROMOTIONS:O, nothing applied here): replace "spin- $\mathbb{C}$ index" with "twisted-spin $((\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6)$ index" in §6.4 / Appendix E.1 / CR4.5 / GP.3, citing arXiv:1910.11277 (DGL) + arXiv:1705.01853 (Tong) + Hsieh–Tachikawa–Yonekura + arXiv:1808.00009 (García-Etxebarria–Montero). The §0.2 on-disk anchors that say "spin- $\mathbb{C}$ " are the ones to correct.

Math to attempt. None new — the no-go and the $\mathbb{Z}_6$ /twist are computed (scripts cited in REDUCTION_SPINTWIST_C §7). The remaining verification: confirm the **twisted-Dirac index** on K_6 returns the same magnitude 3 as the BWB-on-line-bundle computation did (consistency of the corrected structure with the frozen value). Note: the closure-campaign batch1 flagged the **twisted-Dirac spectrum CSV (`k6_dirac_spectrum.csv`) ABSENT** — so a full from-scratch twisted-Dirac re-index is **BLOCKED_INPUTS** pending that artifact.

Success ladder. - DERIVED-CLOSED: twisted-Dirac index on K_6 re-computed = 3, confirming the corrected structure reproduces the value. **BLOCKED_INPUTS** until `k6_dirac_spectrum.csv` is mounted. - **AXIOM-CLOSED / DISCLOSED-CORRECTED (the realistic, cheap win)**: wording corrected to twisted-spin/ $\mathbb{Z}_6$; spin- $\mathbb{C}$ obstruction disclosed; value $|\chi| = 3$ unchanged. **This is the cleanest correctness upgrade in the gate** and requires only an owner countersign. - sharper-OPEN: if the twisted-Dirac re-index is blocked, the structure is corrected on paper but the value-reproduction-under-the-corrected-structure stays AUDIT. - REFUTED: not applicable (the refutation is *of the wording*; it strengthens, not weakens, the gate's honesty).

Honest disposition: DISCLOSED-CORRECTED is reachable now (owner wording patch); DERIVED-CLOSED of the value-under-corrected-structure is BLOCKED_INPUTS on the twisted-Dirac spectrum CSV. The family count 3 survives the correction intact — this is a technical-honesty win, not a downgrade.

4.5 R5 — The index-changing deformation excluded by data, not topology

The obstruction. The index -3 is deformation-proof **inside** the declared admissibility class. The neighboring values (2, 4) that an **index-changing** bundle deformation would produce are excluded by the **LEP N_ν bound (2.984 ± 0.008)**, i.e. by *data* (C2, C4, Appendix E; Objection 4 answer). So "no dial" is **topological in-category + empirical at the class boundary** — two different kinds of robustness.

Technique: T6 (all-operator-conditional) — scope the claim correctly. No closure path "solves" this; the task is to **state the two robustnesses distinctly** so neither is overclaimed.

Named principle (κ^3/π -clean):

AXIOM-INDEX-RIGID-IN-CATEGORY *"the family count is invariant under every continuous deformation within the declared admissibility class; index-changing deformations across the class boundary are excluded by the LEP N_ν bound, an external datum."*

Carries no value; separates the topological claim from the empirical exclusion.

Action (no new physics). Confirm every "no dial / no fourth family by topology" statement is scoped: deformation- proof **in-category** (topology) vs **at the boundary** (LEP). The manuscript already routes the index-changing deformation exclusion to LEP (Objection 4); the residual is keeping that distinction visible wherever "topologically forced" appears.

Success ladder. DISCLOSED-CONSISTENT — already substantially in place (Objection 4 names LEP as the excluder); the residual is a consistency sweep. (Note: closing R3 would make the in-category rigidity *also* boundary-rigid — so R5 is the *dependent* of R3.)

Honest disposition: DISCLOSED-CONSISTENT; keep the topological-vs-empirical robustness split explicit. No new physics; keeping the "no dial" claim honestly scoped.

4.6 R6 — The `G04_chirality` certificate is a consistency lint, not a from-scratch index re-derivation

The obstruction. `certificates/G04_chirality/` lints the parity table (survivors = SM set, mirrors projected, counts consistent) **against the declared index values $-3, (+3, 0)$** — *under those declared values* (CR4.11). It is a consistency check, **not** an independent re-derivation of the index from the bundle data target-blind.

Technique: owner artifact (machine-lane), not an axiom. A fail-closed reproducibility task.

Owner artifact needed. (1) Mount the K_6 bundle data + the twisted-Dirac / BWB computation. (2) Re-derive the index **from the bundle data target-blind** (not read the declared -3). (3) Confirm it equals the frozen -3 and that the APS boundary count re-derives $(+3, 0)$ from the parity ledger `ac4d2df3e708`. (4) Re-hash R1.3/R1.4 and confirm the meta-hash recomputes to `a5b1e6f9d951`. **Caveat:** the from-scratch twisted-Dirac re-index shares R4's blocker (`k6_dirac_spectrum.csv` ABSENT) — so this is **BLOCKED_INPUTS** until that CSV is supplied; the BWB-on-line- bundle leg is symbolically re-derivable now.

Math to attempt. None new — execution + verification. Fail-closed: if the index cannot be regenerated from the bundle data, Gate 4 downgrades per the §5.3.7 falsifier (count $\neq 3$ / count moduli-dependent / mirror survives).

Success ladder. BLOCKED_INPUTS until the bundle-data + spectrum CSV are mounted → then VERIFIED (index regenerates target-blind; "Claimed certificate pass" becomes machine-real

beyond a lint) or **REFUTED** (the index fails to regenerate $\rightarrow$ downgrade). **High value-per-effort** because it converts an asserted index into a machine-checked one with no new physics — but gated on the same missing twisted-Dirac artifact as R4.

Honest disposition: BLOCKED_INPUTS (twisted-Dirac spectrum CSV absent) for the from-scratch re-index; the BWB-on-line-bundle symbolic leg + the APS parity-table leg are re-derivable now. Mount the CSV to reach VERIFIED.

4.7 R7 — The chirality/sign convention bit ($\text{Pin}^\pm$) is unpinned

The obstruction. The **magnitude 3** is robust; the **sign -3** (chirality orientation — "left-handed, not right-handed") rides the **same $\text{Pin}^+/\text{Pin}^-$ / convention bit** that leaves BG-10's $e^{\pm i\pi/4}$ unfixed (SHAPE_Mo_DELTASPINC §4; HANDOFF_SPECIALIST_3: $\sigma_\nu = +1 \rightarrow e^{i\pi/4}$ or $\sigma_\nu = -1, -3 \rightarrow e^{-i\pi/4}$ — not pinned by the frozen record). The family *count* is convention-independent; the chirality *label* is convention-stated.

Technique: T1 (axiom-floor), shared with BG-10/SHAPE. Name the bit:

AXIOM-CHIRALITY-ORIENTATION (κ^3/π -clean): *"the orientation/Pin bit fixing the sign of the index as left-handed (rather than right-handed) is a single discrete convention choice on the $S_Y^1/\mathbb{Z}_2$ reflection; the family-count magnitude is independent of it."*

No value; a discrete-bit statement. **The attack:** show the bit is fixed by the **same** spin-bordism / Dai–Freed single-valuedness datum that fixes the twisted-spin structure (R4), so chirality is structure-forced, not chosen. If the bit is the *default* of the twisted structure and that default is *correct* (left-handed), R7 strengthens; if the default is the *wrong* sign (as BG-10 Rule A's default departs from its target), that is a structure-first datum worth recording.

Specialist target. Hand the spin-bordism specialist: "Is the chirality-orientation Pin bit on $S_Y^1/\mathbb{Z}_2$ fixed by Dai–Freed single-valuedness of the chiral-fermion measure, or is it a free convention?" — the same handoff as HANDOFF_SPECIALIST_3.

Success ladder. AXIOM-CLOSED (AXIOM-CHIRALITY-ORIENTATION named; bit isolated) is the realistic endpoint; DERIVED-CLOSED if the bit is shown Dai–Freed-forced. sharper-OPEN if it stays a free convention. **Honest disposition: AXIOM-CLOSED; the count is unaffected — this is a chirality-label residual shared with BG-10 and SHAPE, three threads landing on one Pin/convention wall (convergence, not closure).**

4.8 R8 — Inheritance: SG-8 flavor and the SG-7 threshold vector depend on "3"

The obstruction. SG-8's generation module $\dim \mathcal{G}_{\text{gen}} = 3$ is **inherited** from this index (A1.13; SG-8 §1.1 attack handle "SG-8 inherits SG-3's conditionality"). The SG-7 threshold δ -vector is **load-bearing on the count -3** (Appendix E / G.3.2: "if the family count were -2 or -4 the column sums would not

reproduce (+4.8424, −3.1112, −1.7313)"). A revision of "3" cascades to Gates 5/7/9/10 (CR4.7; CR4.11 downgrade rule).

Technique: none internal — dependency record. This is not a closure target; it records the **blast radius** so a reviewer sees that SG-3 is the dependency graph's busiest node. The honest action: keep the cascade wired (a Gate-4 downgrade forces Gate-5 → AUDIT and propagates to 7/9/10).

Success ladder. Not applicable (dependency record). **Honest disposition: DISCLOSED; the inheritance is correctly wired. The leverage of SG-3 over the whole programme is high precisely because so many gates consume "3" — which is also why the given- $\mathcal{E}$ qualifier must travel with it everywhere (SG-8 already carries it).**

4.9 Attack-plan roll-up

Residual	Technique	Named axiom (κ^3/π -clean)	Specialist target / owner artifact	Realistic endpoint
R3 bundle admissibility/uniqueness	T10 + T4	AXIOM-MIN-WEIGHT-LIFT	bundle-uniqueness theorem; weight-enumeration + index map	sharper-OPEN → AXIOM-CLOSED if minimality yields 3 value-blind
R1 given- $\mathcal{E}$	T1	AXIOM-CONTENT-GIVEN	(closeable only via R3)	AXIOM-CLOSED (where it sits); gate stays DERIVED-GIVEN-E
R2 bundle-selected (live smuggle)	T5 + T10	AXIOM-COUNT-MUST-BE-INDEX	masked-target selector re-run; minimal-dim rigid-count carrier proof	OPEN → AXIOM-CLOSED (type-axiom defuses smuggle) if masked run passes
R4 spin- $\mathbb{C}$ wording	T4	AXIOM-TWISTED-SPIN-GIVEN-E	owner wording patch; twisted-Dirac re-index (CSV absent)	DISCLOSED-CORRECTED now; DERIVED-CLOSED BLOCKED_INPUTS
R5 data-excluded deformation	T6	AXIOM-INDEX-RIGID-IN-CATEGORY	consistency sweep	DISCLOSED-CONSISTENT (dependent on R3)
R6 certificate lint	machine-lane	(none)	re-derive index target-blind from bundle data (CSV absent)	BLOCKED_INPUTS → VERIFIED
R7 chirality Pin bit	T1	AXIOM-CHIRALITY-ORIENTATION	Dai–Freed bit-forcing (HANDOFF_SPECIALIST_3)	AXIOM-CLOSED (count unaffected)
R8 inheritance	dependency record	(none)	keep cascade wired	DISCLOSED

REDUCE-vs-RELOCATE verdict on the plan. The plan does **not** turn one hard problem into three harder ones. Each path either (a) names a single target-blind axiom that pays a debt in plain sight (R1, R2-type-axiom, R5, R7), (b) attempts a bounded, falsifiable computation that *could* remove conditioning (R3 weight-enumeration; R2 masked-target selector), (c) is a pure correctness/structure fix already computed

on-disk (R4), or (d) is a mechanical owner/CSV task (R6) or a dependency record (R8). The two genuine count-moving attempts (R3, R2) are **smaller** than the original gate (each is one enumeration/selector run) and **both carry an explicit κ^3/π kill-test** so a 3-aligned closure cannot be banked — the SHAPE suite's named live smuggle (the +2-dimension forcedness override) is exactly the trap the kill-test guards against. The honest expected outcome of a full campaign: **R4 DISCLOSED-CORRECTED (cleanest win), R1 AXIOM-CLOSED (already there), R5 DISCLOSED-CONSISTENT, R7 + R2-type-axiom AXIOM-CLOSED, R3 sharper-OPEN (universal-negative unmet, mirroring the demoted A5-actor Lemma 3), R6 BLOCKED_INPUTS → VERIFIED if the twisted-Dirac CSV is mounted, R8 DISCLOSED. No DERIVED-CLOSED is promised**; the gate would move from DERIVED-GIVEN-E-with-spin-C-wording to **DERIVED-GIVEN-E-with-corrected-twisted-spin-structure + a named axiom floor + a machine-verified index** — a real correctness/honesty gain, **not** a promotion. The given- $\mathcal{E}$ qualifier is **not removable** short of R3 closing, and R3's ceiling is AXIOM-CLOSED.

A. Anchoring & Hardening Map

This is SG-3 run through **our internal honesty methodology** — the same hardening method the live **Gaps & Walls Register** applies to the whole programme: for each residual ask **gap-vs-wall** (is the *route* known, or is the route itself the problem?), **hunt the implicit assumption** the residual smuggles in, **find the measured invariant** the residual must ultimately terminate on, then record the **honest disposition** and the **one concrete move** that would harden it. The framework and the disposition vocabulary used here (DERIVED / DERIVED-GIVEN-E / AXIOM-CLOSED / DISSOLVED / SCHEME-ANCHORED / OPEN / BLOCKED / measured-but-irreducible) are defined in, and consistent with, the live register — https://physics.magflowmeters.com/articles/GAPS_AND_WALLS_REGISTER.html. The existing anchor set this gate can terminate residuals on is $\{h, M_Pl, spectrum-E, \alpha_i(M_Z), y_t, |V_us|\}$; a residual hardens only when it reduces to one of those (or to a single *named, value-free, target-blind* posit), never when it reduces to "it gives three."

A.1 Per-residual anchoring (one row per §3.1 residual)

Residual (from §3.1)	GAP or WALL (+kind)	Measured invariant it must terminate on	Honest disposition	What would HARDEN it (concrete next step)
R1 — given- $\mathcal{E}$: "3" is computed WITH $\mathcal{E}$ as input	WALL (route-is-the-problem: the only anchor for " $\mathcal{E}$ is forced" would be $\mathcal{E}$ itself — circular)	spectrum-E directly: the index <i>reads</i> the observed SM chiral content; it cannot derive it. No deeper anchor exists (cross-geometry uniqueness has no witness).	AXIOM-CLOSED at AXIOM-CONTENT-GIVEN — $\mathcal{E}$ named as a SHAPE primitive supplied as input; the gate stays DERIVED-GIVEN-E . given-E $\neq$ derivation of E.	Closeable <i>only via</i> R3 (a bundle-uniqueness theorem). Absent that, keep the given- $\mathcal{E}$ qualifier on every "three families" headline (already done in §6.4 / Objection 4).
R2 — count is bundle-selected, not bare-carrier-forced (the SHAPE live smuggle)	WALL (the escape — the +2-dimension forcedness payment over $\mathbb{CP}^2$ — is target-aligned : it secretly leans on the 3-target; the suite's named smuggling surface)	spectrum-E at the <i>value</i> level; a target-blind TYPE invariant (the count must be a deformation-proof index, not a continuous modulus) at the kind-of-number level.	OPEN , with a clean target-blind type-axiom AXIOM-COUNT-MUST-BE-INDEX that <i>would</i> defuse the smuggle (it rejects "three by dial" for being a <i>dial</i> , not for being <i>three</i>). Carrier-over- $\mathbb{CP}^2$ is legitimately forced for gauge (SG-2).	Re-run the selector with the 3-generation target masked : does the +2-dimension payment still go through on "rigid-integer-count" alone? Plus: prove K_6 is the minimal-dim $SU(3)$ carrier with a rigid-integer count, value-blind.
R3 — bundle admissibility is data-/centre-fed, not internally forced	WALL (a bundle-uniqueness proof requires ruling out <i>all</i> admissible weights — an undischarged universal negative , mirroring the DEMOTED A5-actor / SHAPE Lemma 3)	spectrum-E (the hypercharge ledger + no-4th-generation) and the LEP/SLD $N_\nu = 2.984 \pm 0.008$ light-species count, which excludes the index-changing deformation to 2/4 — <i>data</i> , not internal forcing.	OPEN (highest-leverage); realistic ceiling AXIOM-CLOSED at AXIOM-MIN-WEIGHT-LIFT ; DERIVED-CLOSED gated on an unreachable universal negative. The single highest-leverage move for the whole gate — it is the <i>only</i> path that could lift R1's conditioning.	Bounded computation: enumerate K_6 weights compatible with the $\mathbb{Z}_6$ centre (Tong congruence), map weight $\mapsto \chi$, test whether <i>minimality</i> singles out a unique weight giving $ \chi = 3$ without invoking "three" and without LEP.
R4 — "spin- $\mathbb{C}$ " wording REFUTED for pure SM	GAP (route known: the no-go + the twisted structure are already computed on-disk; a wording/structure	spectrum-E (the SM charges mod 2 force the structure): the genuine forced object is twisted $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$ — FORCED-GIVEN-E , <i>not</i>	DISCLOSED-CORRECTED reachable now (owner wording patch); the from-scratch value-under-corrected-structure re-derivation is BLOCKED (twisted-Dirac spectrum CSV	Owner wording patch (§6.4 / App E.1 / CR4.5 / GP.3) $\rightarrow$ twisted-spin/ $\mathbb{Z}_6$, citing DGL 1910.11277 + Tong 1705.01853. To

Residual (from §3.1)	GAP or WALL (+kind)	Measured invariant it must terminate on	Honest disposition	What would HARDEN it (concrete next step)
	correction, not a value change)	spin- $\mathbb{C}$. The value $ \chi = 3$ is untouched.	k6_dirac_spectrum.csv ABSENT).	reach DERIVED-CLOSED, mount the twisted-Dirac spectrum CSV and confirm the index = 3.
R5 — index-changing deformation excluded by data, not topology	GAP (scoping: state the two robustnesses distinctly; already substantially in place)	LEP/SLD $N_\nu = 2.984 \pm 0.008$ (spectrum-E) at the class boundary; topology (in-category) inside it.	DISCLOSED-CONSISTENT — "no dial" is correctly scoped as deformation-proof <i>in-category</i> (topology) and data-pinned <i>at the boundary</i> (LEP); AXIOM-INDEX-RIGID-IN-CATEGORY names the split. Dependent on R3.	Consistency sweep: keep the topological-vs-empirical split visible wherever "topologically forced / no fourth family" appears. Closing R3 would make the in-category rigidity also boundary-rigid.
R6 — G04_chirality certificate is a consistency lint, not a from-scratch re-derivation	GAP → currently BLOCKED (route exists; the required twisted-Dirac spectrum file is absent — refuse to fabricate it)	spectrum-E indirectly (the certificate must regenerate $-3 / (+3, 0)$ from the bundle + parity ledger target-blind, not read the declared values).	BLOCKED for the from-scratch twisted-Dirac re-index (k6_dirac_spectrum.csv ABSENT); the BWB-on-line-bundle symbolic leg + the APS parity-table leg are re-derivable now.	Mount the bundle data + spectrum CSV; re-derive the index target-blind; confirm = frozen -3 and APS = $(+3, 0)$; re-hash R1.3/R1.4 → meta a5b1e6f9d951 . Reaches VERIFIED (or fail-closed REFUTED → downgrade).
R7 — sign/chirality convention bit (Pin $^\pm$) is unpinned	GAP (a single discrete bit; shares the Pin/convention wall with BG-10 and SHAPE — three threads, one wall)	A NEW named invariant: the Dai–Freed / spin-bordism single-valuedness datum that would fix the orientation bit (no <i>existing</i> anchor pins it). The family-count magnitude is bit-independent.	AXIOM-CLOSED at AXIOM-CHIRALITY-ORIENTATION ; the count is unaffected — a <i>chirality-label</i> residual, not a count residual.	Hand the spin-bordism specialist (same handoff as BG-10 Rule A): is the orientation Pin bit Dai–Freed-forced or a free convention? If the twisted-structure default is correct (left-handed), R7 strengthens.
R8 — inheritance: SG-8 flavor and the	GAP (a dependency record, not a closure target — the blast radius)	Inherited from this gate's spectrum-E anchor; SG-7's δ -vector is itself SCHEME-ANCHORED /	DISCLOSED — the cascade is correctly wired (a Gate-4 downgrade forces Gate-5 →	Keep the cascade wired and the given- $\mathcal{E}$ qualifier travelling with "3" everywhere

Residual (from §3.1)	GAP or WALL (+kind)	Measured invariant it must terminate on	Honest disposition	What would HARDEN it (concrete next step)
SG-7 δ -vector depend on "3"		Diagnostic-only (fitted-not-derived rows), so it adds no independent anchor for "3."	AUDIT, propagating to 7/9/10). No internal action.	it is consumed (SG-8 already carries it; SG-7 is diagnostic-only, so it cannot <i>certify</i> "3" downstream).

A.2 Gate-level rollup

- **Overall disposition.** SG-3 is **DERIVED-GIVEN-E** and stays there. The strong, genuine content — a **rigid, deformation-proof integer** $|\chi(K_6, \mathcal{E})| = 3$ (BWB, closed form) plus a clean **no-go** for the mirror half (APS $(n_L, n_R) = (+3, 0)$) — is real and reproduces; it ties SG-2 as the strongest gate. The honest floor is that **every count residual terminates on spectrum-E** (the index reads $\mathcal{E}$; the 2/4 deformation is excluded by the measured LEP N_ν), so "three" is a rigid facet *of the observed spectrum*, not a from-nothing derivation. Of the eight residuals: **R1/R2/R3 are WALLs** (route-is-the-problem / target-aligned escape / universal negative); **R4/R5/R6/R7/R8 are GAPs** (one BLOCKED on a missing file). No residual is DISSOLVED, and none promotes.
- **Anchors it depends on.** Primarily **spectrum-E** (the SM chiral content $\mathcal{E}$ and its hypercharge ledger), reinforced by the external measured invariant **LEP/SLD $N_\nu = 2.984 \pm 0.008$** at the admissibility-class boundary. It does **not** depend on $\hbar$, M_{Pl} , $\alpha_i(M_Z)$, y_t , or $|V_{us}|$ for the count itself (those enter downstream gates that *consume* "3"). R7 would need a **new** named invariant (the Dai–Freed orientation datum) that no existing anchor supplies.
- **Single highest-leverage hardening move. R3 — the target-blind bundle-uniqueness computation** (enumerate the $\mathbb{Z}_6$ -compatible K_6 weights, map weight $\mapsto \chi$, test whether minimality alone forces $|\chi| = 3$ with no 3-aligned tie-break and no LEP input). It is the *only* path that could lift R1's given- $\mathcal{E}$ conditioning; its realistic ceiling is **AXIOM-CLOSED** (**AXIOM-MIN-WEIGHT-LIFT**), and any "minimality" claim that secretly uses "it gives three" as the selector is the κ^3/π failure mode and must be killed, not banked. The cheapest *correctness* win remains **R4** (the spin- $\mathbb{C} \rightarrow$ twisted-spin wording patch, value 3 intact). **Ceiling: serious candidate, NOT validated.** PROMOTIONS:0; frozen branch dcc66f1b2685 / a5b1e6f9d951 READ-ONLY; nothing applied, nothing deployed.

 Target anchor(s) for this gate

In the **terminate-on** sense (every residual must reduce to a *measured* invariant or a single named, target-blind posit — never to "it gives three"), this gate's target anchor is **the observed spectrum E itself**: the family count "3" — and the rigid integer $\chi(K_6, \mathcal{E}) = -3$ that carries it — is a **facet of E** (χ is read *with* $\mathcal{E}$ as input), reinforced at the admissibility-class boundary by the external measured invariant **LEP/SLD $N_\nu = 2.984 \pm 0.008$** (which excludes the index-changing deformation to 2/4). **Honest status:** spectrum-E is **measured-but-irreducible** (the index reads it; it cannot derive it — given-E $\neq$ derivation of E); the would-be deeper target **"3 across all geometries / bare-carrier-forced"** has **no**

independent witness and stays **OPEN** ($R_1/R_2/R_3$); the in-category "no-dial" rigidity is **SCHEME-ANCHORED** to the declared search category (**AXIOM-INDEX-RIGID-IN-CATEGORY**), and the orientation/chirality-label bit awaits a **new** named invariant (Dai–Freed, **AXIOM-CHIRALITY-ORIENTATION**). The count's *magnitude* is genuine and reproduces; its *forcedness across geometries* is not anchored. PROMOTIONS:o.

5. References & source map

5.1 Website source-of-truth (common material — link, don't duplicate)

- **Paper I, GUT.html** — <https://physics.magflowmeters.com/articles/GUT.html>
- **§6.4** Gate-4 card (chirality / no mirrors / family count; binding status); **§6.12** falsification map (Gate-4 row); **§6.13** certificate summary.
- **§5.3** narrative module (Gate 4); **§3.4** topological-count machinery; **§4.6** selector dimension-payment ($\mathbb{CP}^2 \rightarrow K_6$, "+2 dimensions for forcedness"); **§4.9** anti-fitting firewall ("three by dial" fails).
- **The Strongest Objections** block — **Objection 4** ("the three-family count is asserted by choosing the geometry"; answer: topological integer routed through the selector, residue named — minimality holds only inside the declared search category; LEP N_ν excludes the index-changing deformation).
- **Appendix E** (Chirality Closure — authoritative): E.1 mechanism, E.2 three generations, E.3 mirror ledger, E.6 certificate; **Appendix CR module CR4** (reader companion, CR4.1–CR4.13).
- **Appendix GP.3** (the counting theorems — Atiyah–Singer, APS, Borel–Weil–Bott); **B1/B2** (selector, category-relativity B2.0.3.4); **C2** (K_6 dossier), **C4** ($S_Y^1/\mathbb{Z}_2$ dossier), **A2.2** (bundle ledger), **A1.8** (parity table), **A1.13** (generation module); **Appendix Ro** freeze records.
- **Paper IV, TOE.html** — <https://physics.magflowmeters.com/articles/TOE.html> (downstream cross-references only; does not touch SG-3).

This dossier recaps only what is needed to attack; the site controls all common material.

5.2 Corpus locations (authoritative inputs to this dossier)

Source	Path	Role
Per-gate dossier spec	.../rendered/TOE/PER_GATE_DOSSIER_SPEC.md	structure (sections 0–5)
Sibling exemplar	.../rendered/TOE/PER_GATE_DOSSIERS/DOSSIER_SG8_FLAVOR_CLOSURE_ATTACK.md	depth/format match
SG-3 status line	.../rendered/TOE/SCOPED_GUT_PER_GATE_STATUS_LEDGER.md	the DERIVED-GIVEN-E label + honest caveats (carried verbatim)
Spin-structure refutation (authoritative for R4)	.../rendered/TOE/REDUCTION_SPINTWIST_C_2026-06-23.md	spin- C OBSTRUCTED for pure SM (DGL); forced object = twisted $(\text{Spin} \times G_{\text{SM}})/\mathbb{Z}_6$; FORCED-GIVEN-E; T3 REFUTED — E un-forced
Sign/multiplicity convergence (R4/R7)	.../rendered/TOE/SHAPE_M0_DELTASPINC_QSIGN_2026-06-23.md	$m_0(p, q) = \min(p, q) + 1$ reproduced; $\chi = -3 + \times 2$ spinor-weight; Pin/convention bit unpinned
SHAPE review (authoritative for R2)	.../rendered/TOE/REVIEW_SHAPE_SUITE_2026-06-23.md	suite bottoms on E (un-forced core); defect 4 = live smuggle $(\mathbb{CP}^2 \rightarrow K_6 + 2$ -dimension forcedness override, 3-aligned); selection $\neq$ derivation
Closure campaign (κ^3/π discipline; A5-actor demotion)	.../rendered/TOE/CLOSURE_CAMPAIGN_RESULT_2026-06-24.md	κ^3/π kill-test; AXIOM-CLOSED $\neq$ proven; A5-actor/SHAPE Lemma 3 DEMOTED to OPEN (universal-negative unmet; E un-forced); batch1 $m_0/\mathbb{Z}_6$ reproduced target-blind; <code>k6_dirac_spectrum.csv</code> ABSENT
Closure campaign round 2	.../rendered/TOE/CLOSURE_CAMPAIGN_ROUND2_2026-06-24.md	demotion norm; BG-10 Pin-bit context (shared with R7)
GUT manuscript	.../rendered/GUT/GUT.md	§6.4 (Gate-4 card); §5.3 (narrative); Appendix E (E.1/E.2/E.3/E.6); CR4 (CR4.1–CR4.13); GP.3 (counting theorems); Objection 4; A2.2/A1.8/A1.13

5.3 Frozen-object hash index (load-bearing; READ-ONLY)

Branch `dcc66f1b2685` · manifest meta `a5b1e6f9d951` · spin-**C**/twisted-spin bundle data on K_6 (R1.4) `0fd19c9ae0c1` · $\mathbb{Z}_2$ orbifold + boundary parity ledger (R1.3) `ac4d2df3e708` . Carrier $K_6 = SU(3)/T^2$,

$\chi(K_6) = 6$, and the index value $\chi(K_6, \mathcal{E}) = -3$ / APS $(n_L, n_R) = (+3, 0)$ are frozen R1.3/R1.4 outputs. **No new hash is introduced by this dossier; nothing is mutated.** Missing artifact (blocks R4/R6 from-scratch re-index): `k6_dirac_spectrum.csv` — **ABSENT** (closure-campaign batch1; twisted-Dirac spectrum not on disk).

5.4 External literature (cited at the corpus's own citation level, for R4/R7)

- Davighi, Gripaio, Lohitsiri, "Global anomalies in the Standard Model(s) and Beyond," JHEP 07 (2020) 232, **arXiv:1910.11277**, Sec. 7 — pure SM has **no spin-C** without an extra gauged $U(1)$ (the R4 refutation).
- Tong, "Line Operators in the Standard Model," JHEP 07 (2017) 104, **arXiv:1705.01853** — $\mathbb{Z}_6 =$ maximal trivially-acting centre; congruence $q \equiv 3z_2 - 2z_3 \pmod{6}$ (the R2/R4 centre kernel).
- Hsieh, Tachikawa, Yonekura, "Anomaly of the Standard Model" — twisted **Spin** $\times G_{\text{SM}}/\mathbb{Z}_q$; $(-1)^F$ in the gauge centre.
- García-Etxebarria & Montero, "Dai–Freed anomalies in particle physics," JHEP 08 (2019) 003, **arXiv:1808.00009** — Dai–Freed / spin-bordism fermion-measure (the R4/R7 single-valuedness datum).
- Allanach et al., **arXiv:2111.04148** + anomaly-free atlas JHEP 02 (2019) 082 — infinitely many anomaly-free chiral $U(1)$ extensions $\Rightarrow$ anomaly-freedom fixes neither content nor generation number (the R1/T3-REFUTED basis).

Closing honest statement

SG-3 is **DERIVED-GIVEN-E** — and, among the ten gates, it ties SG-2 as the **strongest**, because its central object is a **rigid, deformation-proof integer** ($|\chi(K_6, \mathcal{E})| = 3$, computed in closed form by Borel–Weil–Bott) plus a clean **no-go** for the mirror half (APS one-sidedness $(n_L, n_R) = (+3, 0)$ on the $S_Y^1/\mathbb{Z}_2$ fold). The magnitude 3 is the family count; the sign and the one-sidedness are chirality; there is **no dial**. That is genuine, defensible, and reproduces (the rep-theory engine — $m_0(p, q)$, the $\mathbb{Z}_6$ Tong congruence — regenerates target-blind). Its honest open surface is equally clear: the index is computed **WITH $\mathcal{E}$ as input** (given-E — it certifies *this geometry+bundle yields three*, **not** that three is unique across geometries, and it **cannot force $\mathcal{E}$** , which stays a SHAPE primitive — T3 REFUTED); the count is **bundle-selected, not bare-carrier-forced** ($\mathbb{CP}^2$ gives "three by dial"; K_6 's forcedness is bought by paying +2 dimensions under a **3-aligned tie-break** — the SHAPE suite's named live smuggle); the **admissible bundle is data-/centre-fed** (LEP N_ν excludes the index-changing deformation), not internally proven unique; and the **literal "spin-C" wording is REFUTED** for pure SM (Davighi–Gripaio–Lohitsiri — the genuine forced object is the twisted **(Spin $\times G_{\text{SM}})/\mathbb{Z}_6$** , FORCED-GIVEN-E, with the value 3 intact). The attack plan reduces these to named, target-blind axioms (AXIOM-MIN-WEIGHT-LIFT, AXIOM-CONTENT-GIVEN, AXIOM-COUNT-MUST-BE-INDEX, AXIOM-TWISTED-SPIN-GIVEN-E, AXIOM-INDEX-RIGID-IN-CATEGORY, AXIOM-CHIRALITY-ORIENTATION) and bounded falsifiable computations (the weight-enumeration index map; the masked-target selector run) — with the κ^3/π kill-test guarding R2/R3 so the 3-aligned forcedness override cannot be dressed as a derivation. The realistic ceiling is a **corrected (twisted-spin) structure + a named axiom floor + a machine-verified index** — the given- $\mathcal{E}$ qualifier is **not removable** short of a bundle-uniqueness theorem (R3),

whose own ceiling is AXIOM-CLOSED. **PROMOTIONS: o; frozen branch** `dcc66f1b2685 / a5b1e6f9d951` **READ-ONLY; given-E $\neq$ derivation of E; selection $\neq$ derivation; nothing applied, nothing deployed.**

Dossier built 2026-06-24. Our geometry (13D K_6 branch) only. Common material referenced to the published website source-of-truth, not duplicated.