SG-7 — Threshold Unification (GUT Gate 7): CLOSURE RESULT — Fable_Version rendered package. Rendered from `SG7_THRESHOLD_CLOSURE_RESULT.md` ; frozen technical content unchanged by rendering.

SG-7 — Threshold Unification (GUT Gate 7): CLOSURE RESULT

> **Honest ceiling.** Serious candidate / diagnostic that recovers the threshold **signs** — **NOT** a validated unification. This result **does not** advance the gate: SG-7 stays **OPEN / Diagnostic**. What it adds is a *value-free, target-blind tool* and a falsifiable path — not a closure. **PROMOTIONS:0.** Frozen branch `dcc66f1b2685` / `a5b1e6f9d951` **READ-ONLY**. Updated 2026-06-25.

§1. What this result is (and is not)

SG-7's standing problem (the 2026-06-24 downgrade): the 8-row heat-kernel threshold ledger that drives $\alpha_1=\alpha_2=\alpha_3$ was **fitted-to-target** — the δ-triple magnitudes entered as constants tuned to hit unification, not derived from the frozen spectrum (deep check missed the printed rows by ~2×–18×). A fitted ledger is **diagnostic-only**.

This result delivers the **remedy infrastructure**, not the fix: a runnable `AXIOM-CANONICAL-ROW-TRACE-COMPILER` that regenerates the threshold rows *target-blind* from a frozen mode table, and **refuses** any input that smuggles in the answer. It has **not** been run on the real frozen spectrum. Therefore:

- **Reduced to Axiom (infrastructure): YES** — `AXIOM-CANONICAL-ROW-TRACE-COMPILER` is built, runnable, tested.
- **DERIVED-GIVEN-E: NOT CLAIMED** — the real frozen projected KK mode table has not been supplied/run.
- **Gate: OPEN / Diagnostic** (unchanged). **PROMOTIONS:0.**

§2. The axiom + the threshold functional

`AXIOM-CANONICAL-ROW-TRACE-COMPILER` — for the frozen admissible SG-7 spectral domain there is a unique *target-blind* row-label compiler ℓ: $\Omega_{adm} \to \{1,\ldots,8, \text{inert}\}$, induced **only** by compact support, actor type, gauge-charge insertion, orbifold/boundary status, chirality, bundle descent, Wilson/Hosotani data, and the frozen admissibility predicate. With row projectors $P_a = \mathbb{1}_{\{\ell^{-1}(a)\}}$, the threshold rows are the finite parts $R_{a,i} = FP_{s=0}\, Tr_{H_{KK}}[P_a\, Q_i^2\, \mu^{2s}\, D_{KK}^{-s}]$, and the threshold vector $\delta_i = (1/2\pi)\, \Sigma_{a=1..8}\, R_{a,i}$.

This axiom is value-free *only if* the row-label compiler and the finite-part scheme are frozen before any comparison with the printed SG-7 ledger. That freeze-before-compare discipline is what the tool mechanically enforces (§3). The eight rows: K6 gauge/ghost SU(3); K6 matter SU(3); S^2 gauge/ghost SU(2); S^2 matter SU(2); S^1_Y hypercharge zero-packet; S^1_Y hypercharge heavy-packet; Higgs Wilson-line; orbifold/boundary fixed points (+ an `inert` bin).

§3. What was independently verified (re-run here, not a captured log)

The runnable package (`SG7_row_trace_compiler_runnable` , published) was executed independently:

- **Runs end-to-end + reproducible:** on the bundled example mode table it produces row traces and column sums, reproducing the stored values to the printed digits (e.g. Delta1_sum = −6.86021259463). *These are TOY numbers from the example table — they are not SG-7 thresholds and must not be read as such.*
- **Tests pass:** 3/3 (`tests/test_compiler.py`).
- **Anti-target-loading genuinely works:** a well-formed mode table carrying a forbidden `delta` column is rejected with `TargetLoadingError: Target-loaded columns are forbidden …: delta` . The guard forbids any column whose name contains `delta, target, printed, expected, fit, fitted, ledger_value, row_value, answer, observed_output – in both the mode table **and** projectors.json` . So the compiler **cannot** be fed the answer; it can only compile rows from frozen primitives (eigenvalues, degeneracies, charge insertions, actor/factor/boundary/Wilson/admissibility labels).

This is the precise anti-target-loading remedy for the "fitted-to-target" downgrade — *as a mechanism*, verified.

§4. What stays OPEN — the falsifiable path to DERIVED-GIVEN-E

The gate stays OPEN because the load-bearing question is **unanswered**, not because a tool was missing:

> **Do not ask whether the rows sum. Ask where the rows come from.** If they come from frozen primitives → the ledger is derived; if they enter as constants → SG-7 stays diagnostic-only.

To lift SG-7 from *Reduced-to-Axiom (infrastructure)* toward **DERIVED-GIVEN-E**, an owner must: 1. Produce `real_frozen_modes.csv` — the real frozen, **target-blind** projected KK mode table (the columns the compiler requires), frozen *before* comparison to the published δ-triple. **(PENDING — not produced here.)** 2. Run the compiler and compare the regenerated rows to the printed ledger **after** the freeze. - rows reproduce within a pre-declared tolerance from primitives → **DERIVED-GIVEN-E**; - rows only match via a normalization reverse-engineered to hit the known δ → **RELABEL_FAIL** (relocation, not closure); - the real table cannot be produced target-blind → stays **OPEN / SCHEME-ANCHORED**.

Current baseline (honest): R1 (rows from primitives) BLOCKED; R2 (magnitudes) OPEN — the δ magnitudes remain fitted-to-target / not formula-derived. The single unfixed scheme-anchor is the per-factor KK ζ-normalization Z_X + overlap c_abi.

§5. Honest roll-up

- **SG-7: OPEN / Diagnostic** (unchanged). Recovered **signs** are real physics (gauge+ghost negative, matter/hypercharge positive, given E + measured α_i(M_Z)); **magnitudes are scheme-anchored, not derived.**

- **Reduced to Axiom:** `AXIOM-CANONICAL-ROW-TRACE-COMPILER` — built, runnable, target-blind, verified *as infrastructure*. This does **not** by itself derive the rows, prove the geometry unique, or derive the SM inputs (separate anchors / scope boundaries).
- **PROMOTIONS:0.** No status advanced. ρ-style target values (the printed δ-triple, the unification point) were inputs to nothing here — the compiler is mechanically forbidden from seeing them.

§6. The target-blind run was EXECUTED — FINDING: blind δ ≠ printed (the rows are fitted, demonstrated); 2026-06-25

The §4 owner-artifact step was carried out, not deferred. A real frozen mode table `real_frozen_modes.csv` (150 modes) was built **target-blind** — eigenvalues + degeneracies from the fixed $K_6/S^2/S^1\_Y$ KK spectra, gauge-charge insertions from the actual SM rep Dynkin indices, scheme freedoms set by **canonical, declared, un-tuned** rules ($Z\_X = 1$, $\mu = 1$, cutoff $(p,q) \le 12$ inherited from the frozen K_6 table). The compiler ran clean (anti-target-loading + schema passed). The result was then compared to the printed δ-triple **post-hoc**.

	δ_1 (U(1)_Y)	δ_2 (SU(2)_L)	δ_3 (SU(3)_c)
Target-blind (this run)	−63.897	+70.627	+227172.5
Printed (fitted) ledger	+4.8424	−3.1112	−1.7313
Match?	✗ sign+mag (13×)	✗ sign+mag (23×)	✗ sign+mag (1.3×10^5)

No match on any component — signs flip on all three, and δ_3 is off by five orders of magnitude. The δ_3 blow-up is **structural, not a bug** (independently confirmed): K_6 gauge-ghost and K_6 matter ride the *identical* KK tower, differing only by the fixed charge×coefficient ratio $q_3 \cdot \text{coef} = (3 \cdot -1)$ vs $(1.5 \cdot +1) = -2{:}1$, which can never cancel — so the canonical insertion cannot produce the small printed δ_3.

Disposition: FINDING — `blind ≠ printed`**.** The blind primitives **demonstrably do not generate** the printed δ-triple under the canonical scheme. This *demonstrates* (not merely asserts) the dossier's own R2/R3 concession: the printed δ are **injected / fitted reals; the invariant→row link is broken.** The post-hoc falsifier **fired** — and the printed values were an input to nothing. (Verified SAFE by an independent adversarial pass + an independent re-run of the column sums.)

Two honest caveats — this is a demonstration that the obvious reconstruction fails, NOT a clean falsification of the geometry: 1. The `log_mu` finite-part is a **self-admitted placeholder** (`finite_part.py`: *"runnable finite-part placeholder … should replace or freeze before comparison"*) — **not** the actual ζ-regularized KK finite part. The wild magnitudes (especially δ_3) are partly placeholder-scheme artifacts of the missing subtraction. 2. The frozen `spectrum_K6.csv` was not co-located in the run package, so its tower was internal-consistency- checked but not byte-traced to source (a flagged UNKNOWN; it does not change the disposition — even granting the towers, the triple fails on sign).

Net: SG-7 stays OPEN / Diagnostic — the gate does **not** move. What changed is the *evidence grade*: the "fitted-to-target" status is now **demonstrated by a target-blind run**, not just conceded. A genuine DERIVED-GIVEN-E attempt still requires the **actual** (non-placeholder) ζ-regularized KK finite-part co-packaged with **byte-anchored** spectra, then a fresh blind run. **PROMOTIONS:0.** *(Artifacts published:*

`real_frozen_modes.csv`, `outputs/real_run.csv`, the `build_real_modes.py` generator, and `FINDING.md`, inside the runnable package.)

The runnable package is published alongside this result so the verification is reproducible (a captured log is not a derivation). Serious candidate, NOT validated.