Survivor Field Theory (SFT)

Consolidated Master Document -- Full Derivation Record -- Version 48

Consolidated 2026-08-27

Contents

Survivor Field Theory (SFT)

formerly Grand Unified Emergence Theory (GUET)

0.1 Consolidated Master Document — Full Derivation Record

0.2 Version 48 — consolidated 2026-08-27 Documentation note, logged per the no-silent-corrections standard: this document's own front-matter version header read "Version 46" at the start of this consolidation, two versions behind the version-history text below it (which already documented v47's own G/AO–AX backfill in full) — the same stale-header defect class previously caught and fixed at v34 and v42. Corrected directly to "Version 48" rather than left to compound further. Structural change carried from v15: full addendum text lives in a companion file. This master document contains only the Normative Core (Parts I, II), macroscopic emergence and predictions (Part IV), external critique tracking (Part V), and the standing open-problems register (Part VI). Full session-addendum text — Addendum A through the present letter, chronological, append-only — lives in SFT_Addenda_Archive.docx. Every claim's derivation status and every cross-reference (e.g. "see BT.6") is unchanged; only the target file has moved, per the v15 split.

*Compiled from the retired original formulation, the revised eight-postulate normative core, and all addenda through Addendum HJ [Front-matter pointer updated at v48 from "**through Addendum HF**" following this consolidation's fold-in of Addenda HG–HJ.] (full addendum-level detail: SFT_Addenda_Archive.docx). [Front-matter pointer updated at v46 from "**through Addendum HC**" following this consolidation's fold-in of Addenda HD–HF.] [NOTE, v47: this pointer tracks the newest addendum LETTER produced, not documentation completeness — it is unchanged by this consolidation, since v47 folds in no new addenda letters, only the previously-lost G and AO–AX documentation gaps (see Part III below) and a companion-archive indexing fix. Still "**through Addendum HF.**"**] [Prior correction, retained: front-matter pointer updated at v45 from "**through Addendum GU**" following this consolidation's fold-in of Addenda GV–HC.] [Prior correction, retained: front-matter

pointer updated at v44 from "**through Addendum GQ**" following this consolidation's fold-in of Addenda GR–GU.] [Prior correction, retained: front-matter pointer updated at v43 from "**through Addendum GN**" following this consolidation's fold-in of Addenda GO–GQ.] [Prior correction, retained: front-matter pointer updated at v42 from "**through Addendum GI**" following this consolidation's fold-in of Addenda GJ–GN.] [Prior correction, retained: front-matter pointer updated at v41 from "**through Addendum GD**" following this consolidation's fold-in of Addenda GE–GI. Previously corrected post-v40 (documentation-accuracy pass, no version bump) from the stale "**through Addendum FW**" marker, unaudited since v35 and left 12 letters behind after Addenda FX–GD were folded in at v40 without updating this pointer. Front-matter pointer corrected at v35 from the v34 "**through Addendum FB**" marker. NOTE ON THIS CONSOLIDATION'S SOURCE TEXT: the v37 master document's own text states a full markdown-corruption cleanup pass was completed at v37 (67 stray-asterisk instances removed, verified zero remaining). The project-knowledge-mounted text used as this v38 consolidation's working source was found to still contain the corruption pattern extensively — 103 total instances (74 three-or-more-asterisk runs, 29 degenerate space-separated " " pairs) — including in this very bracketed note. This is read as confirmation of the project's standing finding that project-knowledge mounted copies of docx files do not reliably reflect the actual cleaned-up delivered binary, not as a fresh corruption of the genuine v37 file. A full sweep was re-applied to this session's working text before use, verified via before/after asterisk-run counts (103→0) and spot-checked against several previously-corrupted passages; logged here per the no-silent-corrections standard rather than fixed silently.] Three addenda folded in this consolidation, 2026-08-05: FS returns to H.17 via a holographic-tensor-network lens never previously applied to this project — a Ryu–Takayanagi (RT) min-cut diagnostic, validated against tree and random-4-regular controls before use — finding the substrate's min-cut scales linearly with boundary-region size (generic expander/area-law behavior) despite real local hyperbolic signatures found in prior addenda; a global-consistency-constrained variant (grow_consistent, enforcing a fixed Q=7 vertex valence, $\kappa=-1/6$) produces a genuine connected negatively-curved core, but the full-graph RT signal remains linear-dominant regardless, and every defect-suppression lever tested (search width, soft-cap tolerance, load-balancing) either converges toward the same finite-closure phenomenon found independently at Addenda AN and DS.6, or is actively counterproductive. FT tests directly whether isolating that weak (vertex-valence-only) core rescues the RT signal — it does not, the isolated core's own RT scan is if anything more linear-dominant than the full graph — then builds the structurally correct constraint (grow_manifold: genuine edge-level {3,Q}-manifold consistency, not merely vertex-valence capping) and finds true consistency under this stronger definition is two to three orders of magnitude rarer than FS's vertex-only core at fixed search width, though a proportional-search regime (K_existing_frac $\gtrsim 0.05$) sustains a genuinely stable, non-vanishing density across a 32× range in N (4,000–128,000) — the first such stable result anywhere in this arc. The resulting true-consistent vertices, however, are found to be completely disconnected from one another (42 isolated singleton components at N=128,000) and to show no spatial clustering whatsoever, statistically indistinguishable from random placement at three density regimes spanning a 25× range — local combinatorial consistency and spatial organization are fully decoupled properties of this construction. FU then tests, rather than merely proposes, the one candidate mechanism FT's own closing discussion named but did not test — that the underlying growth rule's complete lack of a feedback channel between spatially separated closure events explains the decoupling — by building an explicit feedback-biased variant and sweeping its strength directly: the result is a sharp collapse threshold (moderate-to-strong feedback destroys growth entirely, via a runaway preferential-attachment dynamic, well before any clustering could be assessed) with no elevated clustering detected even in the safe, growth-preserving regime below that threshold — a seventh independent, structurally distinct route to the same growth-vs-global-consistency tension this project has now documented via AN, DS.6, FS's three suppression levers, and FT's proportional-search thread. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PER-

MANENT ENTRY throughout. See SFT_Changelog_v38.md for the full diff. [Note retained from v34: three stale front-matter paragraphs, holdovers from the v28–v31 consolidations that had never been removed on subsequent version bumps, were found interposed here at v34 and deleted rather than left in place, per the no-silent-corrections standard.] Two addenda folded in this consolidation, 2026-08-08: FV builds the project**'s first positive-cone semigroup (N^8), confirming it genuinely escapes FR'**s negation-cancellation mechanism at the correct polynomial growth degree (8) — but three symmetric associator-based couplings show null localization, generalizing into a conjecture that any coupling built invariantly from the octonion**'s structure constants inherits some discrete symmetry of them; a genuinely position-dependent coupling (Fano-plane collinearity) breaks this with real, mechanistically-forced, dynamically-robust selection (0/70 collinear-support survivors, 5 seeds × 4 regimes) — but confined to the cone'**s boundary, since the bulk (support ≥5) is provably flat by Fano design regularity. FW answers the resulting open question directly: proves the boundary signal can never carry more than a single discrete bit regardless of magnitude (trilinearity theorem); proves the bulk wall exactly via the Fano plane**'s own 2-design axioms, a result that generalizes to any Steiner triple system rather than being octonion-specific; shows a richer vector-valued coupling escapes the bulk degeneracy but that escape is ~99.7% explained by generic vector-angle proximity rather than Fano-specific structure; and proves the exact symmetry floor (PSL(2,7), order 168) for the discrete coupling class. No formal status-tag changes; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. Ninth independent manifestation of the growth-vs-consistency tension documented by this project. Separately, this consolidation standardizes the aggregate thesis bar (poster and dashboard) as a computed simple average across all eight postulates' own segment data, equal weight per postulate — Derived 35.6% · Conditional 33.1% · Postulated 18.8% · Open 10.6% · Eliminated 1.9% — replacing the poster's prior hand-set editorial figure, which had drifted from the dashboard'**s own live-computed value. See SFT_Changelog_v39.md for the full diff. superseding all prior partial compilations, including SFT_Master_Document.docx (v2) through SFT_Master_Document_v22.docx, SFT_Revised_Paper.docx, SFT_Session_Addendum.docx, SFT_Session_Addendum_H.docx, and SFT_Session_Addendum_H2.docx. Every claim in this document is tagged with an explicit derivation status. Nothing is asserted without a tag.*

Version history: v1 through v19 – see SFT_Master_Document_v19.docx front matter for the full history, preserved there per the v15 archival-split principle. v20 – documentation-and-consistency consolidation only, no new addenda produced. The Addenda BM-BW archive-recovery task (flagged as the natural next application at v19, Part V S3) is complete: all thirteen addenda recovered and re-inserted into SFT_Addenda_Archive.docx, including a follow-up correction to an initial mislabeling of Addendum BN's sub-results (BN.6 was omitted, BN.8's content misnumbered; corrected against a second conversation_search pass and logged in place). Two further instances of the plain-text-saved-as-.docx bug were found and fixed (the archive itself, and this master document) – the ninth and tenth documentation-propagation-audit instances (Part V S3). v21 (this document) – regular consolidation, folding in Addenda CG–CN (Postulate 3 stabilization thread: single-site-invariant closure, R/S/T non-locality confirmed on the real H.17 substrate, a structural account of why R/S/T resist derivation, triangle-closure shown unnecessary for bare AY-escape but the source of a substantial discoverability advantage for genuine local attractors, replicated across four independent substrates). An eleventh documentation-propagation-audit instance was found and fixed during this consolidation: SFT_Master_Document_v20.docx itself had been saved as plain text with a .docx extension, the same bug fixed twice before (v16, v20) and evidently reintroduced a third time; rebuilt as genuine binary docx and verified by PDF render. Postulate 3's formal status is unchanged (POSTULATED, UNDERDETERMINED, four candidates) despite the substantial evidential advances. See SFT_Changelog_v21.md for the full diff. v22 (this document) – consolidation folding in an external-review audit

*(Gemini) and Addenda CO–CR: a band-touching test of Postulate 8's Route A/Route B fork (non-degenerate Clifford structure emerges generically from Route A's own machinery at 2D band-touching points; not in 3D under strict bipartiteness), a soliton-based sidestep of that fork tested and closed negatively, a 4-component resolution restoring 3D consistency without violating bipartite purity, and a first genuine dynamical candidate mechanism for AZ.13's direction-selection question. This master document itself was found, during reconstruction, to exist in three uploaded copies of v21 with different file sizes; two were confirmed byte-for-byte identical to the third after stripping a literal-asterisk text-corruption artifact (424 stray characters from unconverted markdown bold syntax) – a build-pipeline defect distinct from the plain-text-saved-as-.docx pattern, not yet formally logged as a numbered documentation-propagation instance pending the architect's confirmation of its cause. See SFT_Changelog_v22.md for the full diff. v23 (this document) – consolidation folding in Addenda CS–CV, a single-session thread on AZ.13's tachyonic-sign question, Postulate 3's stabilization adjudication, and H.17. Highlights: two routes to forcing AZ.13's sign closed and a third narrowed (CS); R/S/T shown to be patches by the project's own standard while the triangle term is structurally privileged (CT); the non-patching class fully characterized via the exact identity $|[a,b,c]|^2=4(\text{Gram\_det}-\varphi^2)$, showing it has exactly one new generator, φ, which matches or exceeds the associator-norm on every tested metric across three independent substrates (CU, correcting CT.3's uniqueness claim); and the same class shown structurally unable to reach H.17's honest pregeometric substrate, where triangle count is O(1) as $N\to\infty$ (CV). No formal status-tag truth-value changes this consolidation; all four addenda are new derivations, one synthesis correction (CT.3→CU.9), and negative results. The plain-text-saved-as-.docx checkpoint: this document is built and verified as genuine binary docx (PDF-render confirmed); the v22 master as mounted in project knowledge again surfaced as plain text, consistent with the mount-artifact reading logged this session rather than a fresh corruption of the stored file. See SFT_Changelog_v23.md for the full diff. v24 (this document) – consolidation folding in Addenda CW–CZ, a single-session thread characterizing a stable-spiral signature in Postulate 3's second-order dynamics across six independent fixed points and three mechanisms (triangle-term/associator, triangle-term/φ, R's quintic branch), reading as a likely-generic feature of the dynamics rather than a distinguishing property (CZ.6). One correction logged (CZ.1, retracting an incorrect claim in CY.8 about AY's scope). The plain-text-saved-as-.docx checkpoint, this time decisive: unlike v23's mount-artifact reading, both SFT_Master_Document_v23.docx and SFT_Addenda_Archive.docx, as stored in the project, failed a direct Python zipfile.is_zipfile() structural test on the actual file bytes at /mnt/project/ – not the ambiguous project-knowledge text-extraction signature. This is the twelfth confirmed instance of the specific plain-text-saved-as-.docx bug (distinct from the still-unconfirmed v22 literal-asterisk corruption noted above). Both rebuilt as genuine binary docx for v24 and verified by zipfile.is_zipfile()==True plus PDF-render before being finalized. See SFT_Changelog_v24.md for the full diff. v25 (this document) – consolidation folding in Addenda DA–DI, a single-session thread on Postulate 3's φ-vs-associator-norm literal-coefficient question. Highlights: φ proven the unique G2-invariant alternating trilinear form via direct $g2/\Lambda^3(\mathbb{R}^7)$ construction, not merely CU.9's search result (DB.3–DB.5); signed φ shown to fail the theory's own generic-IC protocol outright, narrowing the viable class to {associator-norm, |φ|} (DD); DB.5's reach and the AZ.13/CR "**economy of structure**" argument both found narrower than first framed and corrected in place (DC, DE); the associator-norm/|φ| differentiability distinction shown to have real dynamical consequence via exact zero-crossing tests (DH–DI), reopening a question DA had first found inert, though extensive multi-substrate testing found this does not discriminate between the two candidates (DG, DI) after an initial single-substrate result favoring one candidate did not replicate and was corrected in place (DF→DG). No formal status-tag changes. The standing Part V/Part VI pointer to Addendum BL.4 is implemented in Part II's Postulate 3 text this consolidation. Note on the recurring plain-text-saved-as-.docx checkpoint: per the architect's 2026-07-26 clarification, files rendering as plain text when read back from project knowledge storage is expected Anthropic-side import behavior, not a fresh corruption*

instance each time — this is the last version to carry a numbered instance count for this pattern; going forward it is noted only if a delivered artifact itself (not the project-knowledge copy) fails verification. See SFT_Changelog_v25.md for the full diff. v26 (this document) – consolidation folding in Addendum DJ (Sym³(ℝ⁷) classification: no G2-invariant symmetric trilinear form exists on Im(𝕆) at all, closing rather than paralleling DB.5's alternating-sector result) and Addendum DK (four threads applying the g2-classification technique to Postulate 8's Route A/Route B fork and both open AZ.13 questions for the first time). Highlights: the octonion-multiplication Clifford-generator realization is forced, unique up to gauge and one discrete chirality choice identified as octonion conjugation, closing AZ.12/AZ.13's realization gap (DK.1–DK.9); AZ.13's which-4-of-7 question is upgraded from an empirical search finding to a proof that no algebraic selection mechanism can ever exist there (DK.10–DK.12); Addendum CS.3's tachyonic-sign coupling proposal is shown to occupy the natural point in an exactly 2-dimensional space of G2-invariant options, tightening its form without resolving CS.4's legitimacy gate (DK.14–DK.17); and a flagged resonance between the found chirality sign and Postulate 6's chirality-parity leg is closed by dissolution — no standalone parity operator exists in the record to identify it with (DK.18). One scope correction logged (DK.14, to Addendum CS.5). No formal status-tag changes result from DJ or DK. v27 (this document) – consolidation folding in Addenda DL–DR, a single-day, seven-addendum thread extending Postulate 3's non-patching literal-coefficient question. Closed both follow-up items flagged at v25 on the associator-norm/|φ| kink-propagation finding, with the residual "natural near-degeneracy" puzzle traced fully to a mundane triangle-count/degree confound (DL–DM). Proved the associator is the unique G2-equivariant vector-valued trilinear map on Im(𝕆), a companion theorem to DB.5 (DN). Found, validated, stress-tested, and mechanistically explained a genuine third literal-coefficient candidate, φ^2, shown wider in discoverability window than both established candidates on four independent substrates and robust to seed noise by roughly an order of magnitude (DO–DP), with an analytic account building from a single-triangle self-consistency closure to a validated, near-fitted-constant-free cascade theory matching the true nonlinear dynamics at the node level, not merely in aggregate (DQ–DR). No formal status-tag changes; Postulate 3 remains POSTULATED, UNDERDETERMINED, now among a fully-characterized three-candidate literal-coefficient class. See SFT_Changelog_v27.md for the full diff. v28 (this document) – consolidation folding in Addenda DS–ED (twelve addenda), an extensive H.17/adjacency-before-geometry investigation: two purely-algebraic adjacency constructions confirmed log-diameter (DS), a dimension-parametrized generalization finds no escape at any tested dimension (DT), the log-diameter default's coarse-graining incapacity independently verified via Cheeger/Friedman (DU), a specific sequential-gluing construction found with a genuine hyperbolic-like double signature of unresolved geometric identity (DV–EC, with a methodological correction along the way at DW), and the signature confirmed independent of continuous background access under the strictest finite realization of Postulate 1, clarifying rather than resolving Addendum AJ.6's own standing discretization question (ED). No formal status-tag changes. See SFT_Changelog_v28.md for the full diff. v29 (this document) – consolidation folding in Addenda EE–EG. EE found CS.4's own internal pointer to CS.5 does not actually resolve its gating question, and surfaced a previously unflagged drift between Postulate 7's text ("algebraic triads of substrate elements") and BD/CR's realized construction (free internal fields); extracted the project's postulate-patching criterion (Addendum CT) and found it passes CS.3's coupling on a literal reading but was never tested against a cross-sector case before — CS.4 remains open, now for a precisely characterized reason. EF tested whether CR/BD's already-verified dynamical selection mechanism could supply Postulate 8's separate "which 4 of 7" generator question: verified a general algebraic fact (any orthonormal 4-frame works for DK.7–8's chirality construction, not only the canonical placeholder), raised and then corrected in place an initial counting obstruction, and sharpened the open question from "which 4 of 7" to "which orthonormal basis of L" — L being a specific, already-identified 3-dimensional flat direction, constrained neither by residual G2 symmetry (independently re-verified) nor by any dynamics currently in the theory. EG checked Postulate 4's early

Cheeger-type dichotomy (Addendum B.6, v3-era) against the v28 H.17 thread's newest construction for the first time, finding a genuine third regime outside B.6's two stated branches (logged as an extension, not a retraction); tested a candidate explanation for Addendum EB.2's own open geometric-identity puzzle in a controlled Watts-Strogatz model, finding the effect real but transient rather than stable; and made two honest, source-recovered (via conversation_search) attempts to test this directly on the real DS–ED construction, yielding one large-N result under parameters that don't match the record's own reported numbers (flagged, not treated as confirmatory) and one calibration attempt that correctly failed and was halted rather than extended past an unvalidated build. DZ/EA's ~ -1 conductance ceiling is left explicitly unconfirmed at the N range so far tested. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 4 remains DERIVED. See SFT_Changelog_v29.md for the full diff.*

v30 (this document) – consolidation folding in Addenda EH–EL, a single continuous thread returning to Postulate 1's own wording. EH found a genuine, N-independent, bimodal nucleation transition in grow_hybrid's growth process (glue vs. fresh competition), verified on literal recovered source and independently confirmed on a second construction, after reconnecting the DS–EG construction family to a previously-closed Trugenberger/Kelly mapping (Addenda K–N). EI tested this substrate against Postulate 3's own growth-vs-selection trade-off (AX.10): escapes Addendum CV's negative triangle-count mechanism but still reproduces Addendum BT.2's negative extinct/universal dynamical outcome, not BT.3's positive partial selection, after three isolation tests ruled out topology, coupling-value distributions, and two-point spatial correlation as the cause. EJ found the mechanism via direct basin-counting: grow_hybrid's landscape converges to exactly 2 attractors (confirmed across a 500-seed range) versus a genuinely rugged, still-growing count for an honest uniform substrate, with roughly half the gap traced to global S^6-coverage discrepancy via a controlled 24-realization causal test. EK formally closes, without resolving, an open contradiction from earlier in the session — a topology-invariance proof showing grow_contextual's recovered description cannot support Addendum DZ.3's own reported context_size effect — logged as a permanent open discrepancy. EL found the same bimodality a third time, in landscape type itself, on a second construction, tracing its severity to a critical-fluctuation pattern peaking at an internal phase boundary — sharpening the arc's synthesis into one claim: associator-minimizing constructions show critical-fluctuation-like bimodality at every level examined (growth, dynamics, landscape structure). No formal status-tag changes anywhere in the theory. See SFT_Changelog_v30.md for the full diff. v31 (this document) – consolidation folding in Addenda EM–EP, a continuous thread testing whether the near-critical decoupling finding from v30 generalizes (EM: partially — the conductance-weakening direction does, the δ-hyperbolicity sign-crossing does not, across two structurally distinct counter-examples), sharpening Addendum AW.5's lattice-vs-orbit tension into a provable no-go theorem and testing its constructive resolution against Postulate 3's stabilization question (EN: does not reopen it, more generally than previously shown), finding Postulate 4's bounded-degree requirement follows directly from Postulate 1's own lattice structure with zero free parameters and no dependence on Postulate 3 (EO), and testing whether this motivates a triadic rewrite of Postulate 2 itself — well-motivated, explaining two previously-unconnected historical results (Addendum BF's failure, Addendum BU–BV's normalization convention) for the first time, but not yet actionable, with a newly-raised origin-privileging concern tested directly against Addendum BV's own real substrate and dynamics and found not to threaten it (EP). No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 4 remains DERIVED, now on two independent grounds for its bounded-degree component. See SFT_Changelog_v31.md for the full diff. v32 (this document) – consolidation folding in Addenda EQ–EV (six addenda across two threads). The first thread (EQ–ET) tested whether Postulate 2 itself should be rewritten around a triadic (associator-based) coupling, motivated directly by EN's no-go theorem, EO's dynamics-free adjacency derivation, and EP's own triadic analysis: EQ drafted a candidate rewrite, gated

CONJECTURE/OPEN behind four explicit prerequisites; ER corrected a cross-reference error in EP.7's own framing (BW.6's coherence degradation is not actually unexplained, per the already-resolved Addendum BZ), retargeting the gate toward recovering BW's own literal h_scale parametrization instead; ES then did exactly that, bit-exact reconstructing the BW/BZ N=90,778 substrate under Addendum CF's exact recovered dynamics for the first time and finding BZ's imported h=0.033 is supercritical at this scale (~53% activation, near-random coherence), with the genuine selective/fragmented regime instead near h≈0.02–0.025; ET then tested EQ's own origin-privileging premise (first raised at EP.4, found dynamically inert on BV's substrate at EP.7) on a second, structurally different substrate (Gaussian scatter in $\mathbb{R}^7$ rather than a lattice) and overturned that inertness finding — the static bias replicates (r=−0.553) but this time propagates substantially into survivor dynamics (r=−0.088 to −0.368, depending on active fraction). EQ's draft rewrite is explicitly not recommended for adoption pending resolution of why BV's substrate was inert while lattice/Gaussian substrates are not — the gate remains open, now for a sharper and better-evidenced reason. The second thread (EU–EV) returned to H.17 with two independent diagnostics never previously applied to this project's adjacency constructions. EU computed spectral dimension (random-walk return probability / heat-kernel flow) directly on the DS.7/DS.8-family hybrid construction for the first time, verified against known reference cases (2D/3D lattice, random-regular expander) before being trusted, and found no dimension-lock at any scale tested up to N=50,000 — the finite-N spectral-dimension peak grows with N without converging (2.66→3.01→≥3.20→4.02 across N=6,000→50,000, passing through d_s=4 without stopping there), a second, independent line of evidence for Addendum DU's Cheeger-based no-go result. EV then tested a genuinely different question on the same substrate — whether physical dynamics riding on top of an expander-like graph can still be confined, independent of whether the bare graph itself has a fixed dimension — using Postulate 2's own already-derived complex coupling C_ij = W_ij + iΓ_ij (Addendum Z.7/AU, no new structure introduced) as a tight-binding-style Hamiltonian: the bare topology is delocalized (as EU/DU already established), but Postulate 2's real coupling alone produces genuine Anderson-style localization (IPR plateaus, N-independent, confirmed via seed-noise calibration and independently via real-space localization length), while the genuine octonion phase term Γ_ij specifically weakens rather than strengthens this confinement — a real, unexpected nuance reported precisely rather than smoothed toward the motivating hypothesis. No formal status-tag changes result from EQ, ER, ES, ET, EU, or EV; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 2 remains DERIVED (Hermiticity, uniformity) / POSTULATED (base algebraic form), its triadic-rewrite question now sharpened rather than resolved; Postulate 3 remains POSTULATED, UNDERDETERMINED. A related but separate gap was also flagged this session, not addressed: Addendum H's own Combes–Thomas/Wegner-estimate admissibility proof for H_W (H.11–H.14) depends on B.5/B.6's Lieb–Robinson bound, which requires polynomial substrate growth — a premise H.17's own subsequent work (DS through EU) has since shown does not hold for the honest pregeometric substrate. See SFT_Changelog_v32.md for the full diff. v33 (this document) – consolidation folding in Addenda EW–EY, a three-addendum thread continuing directly from v32's own flagged H_W/B.5-B.6 note. EW ran three structurally distinct constructions (the DS.7/DS.8 hybrid gluing family, DT's dimension-parametrized simplex family, and DS.2's static filter) through a new conductance-vs-confinement-scale diagnostic, finding bottleneck-bearing constructions show a significant, bootstrap-confirmed correlation between conductance strength and Postulate 2's confinement scale (ξ, from Addendum EV), while a construction independently known to lack a genuine bottleneck (DS.2) shows no such correlation — a triangulated, falsifiable link to Addendum DV's own dimension-dependent findings, cross-validated against DU's literal recovered conductance code. No status-tag changes result from EW. EX then re-examined v32's own flagged dependency directly rather than treating it as settled: found the Lieb-Robinson condition was never an absolute wall (it sharpens to a checkable inequality on the honest substrate, not an automatic pass), found B.5 likely used the wrong tool in the first place (Hernández-Jansen-Lüscher's own theorem proves overlap-operator locality directly from bounded degree and a spectral gap, no volume-growth

condition needed), but then found a second, independent, previously-unflagged polynomial-growth dependency inside H.14's own multiscale induction (an explicit L_k^d volume factor) — meaning the spectral-gap input the more direct route needs is itself unproven on the honest substrate. A literature check on Anderson localization for tree/expander-class graphs (DU/DV/EU's own characterization of the honest substrate) found this is a genuinely different, threshold-gated problem, not automatic at any disorder strength, and flagged a new, previously unlogged scope caveat on Addendum EV's own N-independence claim. EX recommended overlap admissibility be revisited as CONJECTURE/OPEN pending one concrete calculation, without unilaterally changing the tag itself. EY carried out that calculation. Measured directly on this project's own already-verified honest substrates (grow_hybrid, DS.2), using Postulate 2's real coupling and Addendum H.13's own disorder definition: H.11's deterministic Weyl certificate does not cover the honest substrate at any tested configuration (r_eff 5–9× past its own threshold); a connectivity-matched Anderson/Bethe-lattice critical-disorder threshold (calibrated against an independent literature reference point before use) is not cleared by SFT's actual measured disorder at any reliably-measurable configuration or connectivity, sparse or dense, with the shortfall worsening as connectivity increases. A same-session correction (EY.9) found the one apparent near-miss in the initial scan was a measurement artifact (disconnected graph fragments contaminating a whole-graph statistic); the corrected, giant-component-only sparse-end scan closes the one open item and finds the negative result more decisively supported, not less. **Status-tag change this consolidation: overlap admissibility (H_W spectral gap, Addendum H §H.11–H.14) is downgraded from SUBSTANTIALLY RESOLVED to CONJECTURE/OPEN and restored to the Part VI open-problems register, from which it was removed at v3.** This is an evidential, not forcing, negative result — see Part VI for full honest scope. H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY, unaffected. Postulate 8's own CONDITIONAL status (Route A/Route B fork) is unaffected — overlap admissibility is a downstream fermion-sector question, not a restatement of that fork. See SFT_Changelog_v33.md for the full diff. v34 (this document) – consolidation folding in Addenda EZ–FB, a three-addendum thread reopening AZ.13's tachyonic-sign question and then attacking AW.5's lattice-vs-orbit tension directly. EZ reframed AZ.13's unforced $\gamma<0$ sign as a physical-completeness argument (why the potential must eventually couple to substrate at all) and converted it into a threshold-coupling question via an effective mass $\gamma_{eff}(n)=\gamma-\kappa\cdot n^{T}Mn$, tested directly on both of the project's own real substrate constructions: the orbit substrate (by1, bit-exact) gives an exactly rank-5-of-7 second-moment tensor with maximal anisotropy, forced by D's own (1,4,2) eigenspace decomposition and robust across 8/8 seeds, while the lattice substrate (BU/BV, bit-exact) gives a near-isotropic, full-rank result, robust across 5/5 seeds — revealing the mechanism is substrate-type-dependent and, for the first time, linking AZ.13 directly to AW.5's own standing tension. A previously unnoticed structural fact surfaced in the same investigation: D's fixed axis $\hat{n}$ is an exact eigenvector of the orbit substrate's second-moment tensor (cos=1.00000), left unchased at the time (EZ.6). FA then attacked AW.5 directly: orbit density/ergodicity as a reconciliation route is ruled out for free from already-proven exact orbit periodicity (FA.1); a multi-orbit ensemble construction (K independent one-parameter-subgroup draws rather than one orbit sampled N times) achieves full rank at K=2 and anisotropy falling smoothly to 0.18 by K=500 (FA.2–FA.3) — but calibration against an IID-uniform-on-S^6 baseline at matched N found the construction's growth-exponent plateau is statistically indistinguishable from ideal finite-N sampling of a compact manifold, not a construction flaw (FA.4). The mechanistic diagnosis (FA.5): shell growth on any orbit-based substrate saturates to ~100% of N within log(N) hops regardless of construction — forced by G2's own compactness, the same compactness Addendum Z.15's Poincaré-recurrence argument already used, now shown to force H.17's own log-diameter fork-branch directly. Coupling/selection dynamics were confirmed to still work on the ensemble substrate after a too-coarse initial scan was corrected (FA.6). FB then tested a third hunch (a fixed-axis-generated substrate) and, in the course of doing so, found a genuine new closed-form identity — not previously stated anywhere in the record — that D's fixed axis $\hat{n}$ equals exactly $\pm\mathrm{Im}(a\cdot b)/|\mathrm{Im}(a\cdot b)|$, requiring no eigendecomposition (FB.1); combined

with Addendum AW.1/BU.1's already-established closure of $\mathbb{Z}^7$ under octonion multiplication, this shows lattice-valued generator pairs give exactly lattice-valued fixed axes (FB.2). An initially promising 'explosive growth' signal from products of large-radius lattice pools was caught and corrected before being reported positive — it is exactly the norm-multiplicativity identity |ab|=|a||b| acting on an already-large pool, not genuine generation (FB.3) — and the corrected, exhaustive (not sampled) test found pure multiplicative closure from the seven basis vectors alone is exactly finite: 14 points, zero growth after generation 1 (FB.4). Synthesis (EZ.9, FA.7, FB.5): AW.5 very likely reduces to H.17 rather than standing as an independently resolvable problem — a conclusion now supported by four structurally unrelated confirmations of the same compact/finite obstruction (Addendum AN's discrete-subgroup finiteness, Addendum Z.15's Poincaré recurrence, FA.5's graph-distance saturation, and FB.4's multiplicative closure). This is evidential, not a forcing theorem — it does not rule out non-compact orbit modifications beyond Addenda Z.16–Z.20's already-explored split-octonion route. No status-tag changes: H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 3's literal-coefficient candidates and CS.4's legitimacy gate are untouched. FB.1's closed-form $\hat{n}$ identity is promoted into Postulate 8's own text this consolidation (Part II), stated alongside (not replacing) the eigendecomposition-based description used since Addendum Z. One documentation defect caught and fixed this consolidation, per the no-silent-corrections standard: this document's own front-matter version header read 'Version 32' while the version-history text below it already documented v33 in full — a stale-header instance distinct from the previously-logged plain-text-saved-as-.docx pattern, not previously flagged. Corrected to read 'Version 34' and logged here rather than silently fixed. A related stale pointer ("through Addendum EV" in the front-matter summary, three versions behind) was also caught and corrected — see above. See SFT_Changelog_v34.md for the full diff. v35 (this document) – consolidation folding in Addenda FC–FH (six addenda across two threads). FC–FD close the DZ/EA conductance-ceiling thread standing open since v29: FC completed EM.14's three open items (d=3 pushed to its own collapse boundary with the δ-crossing still not completing, strengthening EM.13's conjecture; DS.2 pushed toward its percolation edge with a noisy sign-flip; d=4 tested for the first time and found distinct from d=3); FD delivered the decisive mechanism test (grow_hybrid's diameter stays flat across its stable range while DS.2 diverges sharply near percolation, confirming two genuinely different mechanisms), found DS.2's own M-ratio is an independent axis the project's own default value characterizes noisiest, and retracted FC.4's own tentative N-scaling reading after one further data point reversed it. FE–FH then form a single continuous arc testing the session's own opening hunch — whether some untested construction axis lets a non-compact algebraic substrate escape H.17's log-diameter default — to exhaustion: FE found the literal hunch already closed by Addendum AN's own general argument and redirected to a genuinely untested cell (lattice-ordered vertex set × non-metric algebraic adjacency), still log-diameter at a 900× N range; FF closed the independent-vs-correlated sampling axis on direct test, corroborated by the project's own prior grow_hybrid record, after catching a fixed-window construction artifact before it could be reported as a positive escape; FG found a real, mechanistically-explained, cross-validated slope effect from a properly N-scaled frontier — still log-diameter, not an escape; FH proved that effect algebra-independent via direct null-model substitution, supplied the correct (graph-theoretic) first-principles derivation with a rigorous lower bound and a Monte Carlo-verified upper bound, then caught a genuine indexing bug in its own follow-up test, corrected it, and identified for the first time the precise mechanism (an asymmetric recency boundary layer) behind this project's own multi-source diameter measure. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. A pre-existing literal-asterisk markdown-corruption artifact, surrounding the front matter's own "through Addendum X" pointer, was found and cleaned up in place while updating that same pointer this consolidation — logged here rather than silently fixed; the broader, previously-flagged instances of this pattern elsewhere in Part III's older addendum-summary text remain unaudited, per the standing note below. See SFT_Changelog_v35.md for the full diff. v36 (this document) – consolidation folding in Addenda FI–FN (six addenda), a single con-

tinuous thread on CS.4's legitimacy gate and, separately, its physical consequences if adopted. FI executes Addendum EE.6's own recommended precedent audit — a systematic sweep of the addenda history for any other cross-sector potential coupling (Postulate-to-Postulate, not within-sector) — and finds none: the closest candidate (Addendum BU.2/EP.4's reuse of Postulate 2/3's associator machinery on lattice substrate points) is disqualified as same-sector under Addendum EE.5's own test, and the one genuine cross-postulate attempt (Addendum EF, testing whether Postulate 7's CR/BD mechanism could supply Postulate 8's generator question) found the opposite of precedent — the theory's existing potentials are cross-sector-blind by construction. CS.3 stands confirmed as the theory's first proposed cross-sector coupling, with no prior case to judge it against. FJ–FN then took a different, direct approach: rather than continuing to argue whether the coupling is licensed, they built it and tested what it does. Reconstructing the full toolchain bit-exact (octonion algebra, g2 = Der(O) from scratch, both real substrate constructions, verified against Addenda EZ.4–EZ.8's recorded figures), CR.3's exact pairwise-collapse mechanism was coupled to each substrate's own second-moment tensor M across a full κ sweep: the mechanism survives completely intact at every tested coupling strength on both substrates ($\cos(v_1,v_2)=\pm1.000000$ throughout, FJ), and BD's triad sector was confirmed decoupled under genuine joint 35-parameter optimization rather than assumed by construction (FK.1). What changes is the previously flat, arbitrary S^6 direction CR.3 left undetermined: substrate-type-dependent resolution, full determination on the generic lattice substrate, and — via direct finite-difference Hessian computation, not inference — an exactly massless, permanently protected mode on the orbit substrate at every tested κ, traced to M's own doubly-degenerate eigenvalue and a nonzero Stab_G2(M) (FK.3–FK.4). One correction was logged along the way (FK.2, an apparent lattice-substrate threshold found to be a BFGS local-minima artifact rather than real physics, once checked via the Hessian directly rather than an alignment metric alone) and one existing claim was checked and corrected (FL: EZ.6's n̂/M-eigenvector finding confirmed real but mislabeled — n̂ aligns with M's isolated, non-degenerate middle eigenvalue, not its top eigenvalue as previously stated, ruling out a hoped-for connection to the massless mode). FM then derived the orbit-substrate finding from first principles rather than leaving it as an observed pattern: a full closed-form theorem (any bounded rotation plane, generated by a compact-group element, gives an exactly isotropic time-averaged block by Haar-averaging) predicts M's entire seven-eigenvalue spectrum from D's own eigenspaces alone, verified to 7+ decimal places — identifying the massless mode as a fourth independent manifestation of the identical G2-compactness obstruction already documented in Addenda AN, Z.15, and FA.5, and recovering its exact manifestation rate (91.73%, confirmed via a direct 20,000-trial Monte Carlo) for free from Addendum AF's own decade-old Dirichlet(1/2,2,1) distribution — with k=1 and k=2 harmonic dominance shown to produce the identical signature, confirming k=0-vs-k≠0 is the complete classification (FM.3.1–FM.3.2). FN closed the one remaining open question — whether an AG-style spectral-gap conditionality applies — with a direct stress test at 30× tighter neighboring-eigenvalue gap than any previously-tested case, finding zero degradation and confirming the two claims are different in kind: AG's threshold is intrinsic to a linear approximation standing in for nonlinear dynamics, while FK/FM's claim is an exact symmetry of the coupled potential itself, gap-independent by construction. v37 (this document) – consolidation folding in Addenda FO–FR, a continuous single-session thread returning to H.17 directly and to CS.4's legitimacy gate. FO closed Addendum FE.10a's last flagged gap in the H.17 log-diameter arc — DS.7/DS.8's sequential hybrid-gluing construction re-run on a lattice-ordered vertex set instead of its native Haar-random points — extending the range to 1024× (N=200–204,800) after an initial 256× pass proved genuinely ambiguous; the log fit pulled decisively ahead (R^2=0.9764 vs. 0.9656), corroborated by a shrinking free power-law exponent and decelerating diameter increments — the fifth structurally independent construction axis tested in this arc, and the fifth to return log-diameter. FP then closed CS.4's legitimacy gate directly, by the architect's explicit decision: the v-M coupling proposed in Addendum CS.3 and tested unconditionally in Addenda FJ–FN is LICENSED, per Addendum EE.2's reading that Postulate 7's own text ("algebraic triads of substrate elements") was al-

ways the operative statement of intent, with BD/CR's free-field construction an unreconciled stand-in. This promotes FJ–FN's results from conditional to unconditional theory content, and identifies a substrate-dependent path to discharging AZ.13's previously-unforced tachyonic sign via Addendum EZ.2's effective-mass mechanism — close to forced on the orbit substrate, still fine-tuned on the lattice substrate — tying AZ.13's resolution directly to the outstanding AW.5/H.17 lattice-vs-orbit question. FQ then revisited H.17 directly, attacking a horn of its own circularity fork — polynomial growth presupposes an imposed lattice — that the entire prior arc (Addenda DS through FO) had left untouched, having instead attacked the opposite horn throughout. Building Coxeter's canonical maximal order in the octonions from scratch (independently reproducing, and correcting, the historical Kirmse's-mistake construction error along the way), FQ showed growth degree is exactly 8 by theorem regardless of which order is used — closing a gap Addendum EO.5 had flagged — and that Postulate 3's own non-patching dynamics forces the canonical choice, since the non-canonical integer lattice's minimal-vector graph is exactly triangle-free and cannot carry the triangle term at all. On a genuine bulk region of this substrate, polynomial growth and real algebraic selection were shown to coexist for the first time in the project's history, but the selected sets proved to be exactly centrally symmetric — no spatial localization — restating H.17's second horn as a symmetry obstruction (selection can only return globally spread group orbits) in place of the original circularity charge. FR then tested this new restatement directly and, over the course of a single extended addendum, corrected its own proposed mechanism twice before landing on the operative one: the true cause is not orbit structure under some specific automorphism group (refuted directly, since an alternative un-normalized coupling gives a different partition of nodes yet identical non-localization) but exact invariance under global negation, $x \to -x$, forced by -1 being central in the octonions and fixed by every automorphism — a fact verified to hold, exactly, for every literal-coefficient candidate the theory currently admits as non-patching (associator-norm, $|\varphi|$, φ^2, all independently built and tested), and shown to survive two structurally distinct attempts at an explicit chirality escape (dropping the magnitude entirely gives a second, independently-arising instance of the identical parity obstruction, for an unrelated combinatorial reason tied to alternating-form sign conventions). A closing test of whether abandoning additive-lattice substrates entirely might escape the obstruction found it coherent in principle — the project's own orbit substrate construction is a real, verified example of a forced, non-patched substrate that is not negation-closed — but not rescuing, since every substrate construction the project has built, across four structurally distinct paradigms, fails to deliver both genuine polynomial growth and freedom from this symmetry, each for its own independently-established reason. No formal status-tag changes to H.17 itself, which remains CONJECTURE/OPEN, PERMANENT ENTRY; CS.4 is recorded as DECIDED (licensed), the first use of that status tag in the project's history. See SFT_Changelog_v37.md for the full diff. v40 (this document) – consolidation folding in Addenda FX–GD (seven addenda). FX tests real Postulate-3 dynamics on one concrete random realization of Addendum FV's positive-cone semigroup rather than a closed-form invariant coupling: genuinely escapes the closed-form degeneracy (a converged, IC-independent PARTIAL attractor, replicated across 10 realizations), but the mechanism behind it (local triangle-associator mass predicting survival) reduces almost entirely to generic vector-angle geometry once decomposed, extending FW.9's honest negative to the dynamical case. FY tests the positive-cone semigroup directly as a candidate answer to Addendum AJ.6's discretization question: builds a genuine bounded-degree adjacency on it for the first time (correcting two failed attempts — raw-magnitude hub concentration, then over-corrected fragmentation — both traced to a magnitude bias, fixed via normalized compatibility), finding log-diameter growth with an elevated constant, an eighth independent construction family joining the same wall; states explicitly, per AN.3–4's original argument, that AJ.6 and H.17 are one problem, not two. FZ proves a genuine classification theorem strengthening Addendum FV.9's conjecture and FW.10's single-instance proof: the algebra-automorphism group (order 21) is exactly the normalizer of a Sylow-7 subgroup within the Fano-collinearity group (order 168, PSL(2,7)), and is maximal there — no intermediate discrete symmetry level exists between them — while a second, previously

unidentified maximal-subgroup family (order 24, point-stabilizers ≅ S4) governs couplings that privilege one fixed basis direction; a hypothesized empirical connection to Addendum EZ's n̂-based coupling machinery was tested directly and retracted in the same session once n̂ was confirmed generic (no basis alignment) rather than reported as confirmed. GA–GD close a four-addendum thread on Addendum EF.7's flat direction (L's basis) and Addendum BI.7's seed-point arbitrariness, showing they are the same underlying gap and then closing both routes that could have resolved it: GA finds M's (the substrate second-moment tensor, licensed at Addendum FP) own kernel structure depends on the orbit's seed point e_0 rather than supplying a canonical answer; GB finds D's high-frequency eigenspace is exactly span{a,b} (the orbit's own generators), ruling out reusing them as e_0; GC proves via Artin's theorem that no algebraic combination of a,b can ever reach the remaining 4-dimensional eigenspace at all; GD confirms by direct construction (not by label) that the residual su(2) stabilizer acts irreducibly on that eigenspace, with no invariant direction — closing both the algebraic and representation-theoretic routes and confirming e_0 is a genuine, irreducible free parameter rather than an unresolved gap. [Corrected at v43, logged rather than silently reconciled: this synthesis's extension of e_0's irreducibility to also closing Addendum EF.7 rested on an unverified premise (GA.5) never independently re-checked by GB, GC, or GD themselves — Addendum GQ (v43) resolves EF.7 directly via a construction referencing neither e_0 nor the D-eigenvector machinery this thread concerns; see Addendum GQ.16 for the full account. GB–GD's own results about e_0's irreducibility are unaffected by this correction.] No formal status-tag changes anywhere in the theory. See SFT_Changelog_v40.md for the full diff. v39 – consolidation folding in Addenda FV–FW. FV builds the project's first positive-cone semigroup, confirming genuine escape from FR's negation-cancellation mechanism at the correct polynomial growth degree, but three symmetric couplings show null localization; a Fano-plane-collinearity coupling breaks this with real, dynamically-robust selection, confined to the cone's boundary since the bulk is provably flat by Fano design regularity. FW answers the resulting open question with three exact theorems (trilinearity of the boundary signal; the bulk wall via the Fano plane's 2-design axioms, generalizing to any Steiner triple system; the PSL(2,7)-order-168 symmetry floor for the discrete coupling class) plus an honest negative (a richer vector coupling's escape from the bulk wall is ~99.7% generic geometry, not Fano-specific). No formal status-tag change; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. Aggregate-bar methodology standardized this consolidation: both poster and dashboard now use a computed simple average across all eight postulates' own segment data (Derived 35.6% / Conditional 33.1% / Postulated 18.8% / Open 10.6% / Eliminated 1.9%), replacing the poster's prior hand-set figure. See SFT_Changelog_v39.md for the full diff. v38 – consolidation folding in Addenda FS–FU, a continuous three-addendum thread returning to H.17 through a holographic-tensor-network lens never previously applied in this project. FS ran a Ryu–Takayanagi (RT) min-cut diagnostic on the project's own established substrate constructions, validated against a regular tree (positive control, near-flat cut scaling) and a random 4-regular graph (negative control, strongly linear) before being trusted: the honest substrate's min-cut scales linearly with boundary-region size — generic expander/area-law behavior, not the logarithmic signature genuine holographic structure requires — despite real local hyperbolic signatures (bounded Gromov δ-hyperbolicity, a shrinking conductance exponent) established in earlier addenda. A constructive response — grow_consistent, enforcing a hard Q=7 vertex-valence cap (the minimal hyperbolic {3,Q} triangulation, $\kappa = -1/6$) — succeeds at producing a genuine connected, negatively-curved, proportionally-growing consistent core for the first time in the project's history, but the full-graph RT signal is unmoved: still linear-dominant, the defect-laden majority (roughly 80% of the graph at baseline search width) swamping any holographic contribution from the core. Three independent defect-suppression levers were tested against this: search width is the only one that helps at all, but only by converging toward the identical finite-closure phenomenon Addenda AN and DS.6 already found by unrelated methods, now via a third, structurally independent route; soft-cap tolerance is actively counterproductive (causes disclination pile-up rather than spreading); load-balancing selection bias is neutral-to-counterproductive. FT then tests the one hypothesis FS's own re-

sult structure left open — that the core itself is holographic but the full-graph test is too noisy to see it — by isolating the core directly and re-running the identical RT scan on it alone: the isolated core is, if anything, more linear-dominant than the full graph (steeper linear slope), decisively closing that hypothesis. FT then builds the structurally correct constraint FS's own diagnosis had identified as missing but not yet implemented: grow_manifold enforces genuine edge-level consistency (every edge capped at exactly 2 incident faces, the literal definition of a discrete simplicial 2-manifold) in addition to the vertex-valence cap, and finds true consistency under this stronger, correct definition is two to three orders of magnitude rarer than FS's vertex-only core at any fixed search width — but that a proportional-search regime (search width scaling as a fixed fraction of the growing candidate pool, K_existing_frac $\gtrsim$ 0.05) sustains a genuinely stable, non-vanishing density across a 32× range in N, 4,000 to 128,000 — the first stable (non-vanishing-in-the-limit) consistency density found anywhere in this project's entire H.17 arc. This result is immediately qualified by two further tests: the resulting true-consistent vertices form no connected structure at all (42 isolated singleton components at N=128,000, each fully surrounded by inconsistent neighbors) and show no spatial clustering whatsoever under a rigorous pairwise-distance test against a random-placement baseline, confirmed not just at this one density but at two further, independently-tested density regimes spanning a 25× range — an in-session speculative explanation for the isolation (a possible repulsion mechanism between closure events) was proposed, tested, found unsupported, and retracted per the project's no-silent-corrections standard once the appropriate control calculation was run. FU takes the one remaining named-but-untested candidate explanation from FT's own closing discussion — that the growth rule's complete absence of any feedback channel between spatially separated closure events is what forces this decoupling — and tests it directly, rather than leaving it as a plausible but unverified account: building a feedback-biased variant of the same construction, where achieving true-closed status locally discounts nearby candidates' selection cost, and sweeping the feedback strength systematically. The result is a sharp collapse threshold, not a rescue: feedback strong enough to plausibly bias local structure (roughly bias_strength $\gtrsim$ 0.3) triggers an early, complete growth collapse via a runaway preferential-attachment dynamic (the frontier empties entirely at N in the low hundreds, far short of any target scale), while feedback weak enough to preserve normal growth ($\leq$ 0.2) shows no clustering benefit over the zero-feedback baseline at all. This closes FT's own open mechanistic question in the negative, and constitutes a seventh independent, structurally distinct manifestation of the growth-vs-global-consistency tension this project has now documented across Addenda AN (discrete-subgroup finiteness), DS.6 (deterministic-iteration closure), FS's three suppression levers, and FT's proportional-search thread — now via engineered local positive feedback rather than any form of search-width, tolerance, or load-balancing tuning. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY throughout FS–FU, as it has since the entry was first sharpened into its present form. See SFT_Changelog_v38.md for the full diff.
v42 (this document) – consolidation folding in Addenda GJ–GN (five addenda, two threads). GJ extends the Heisenberg-group k-NN construction (Addenda GG–GI) with a systematic seven-scale study (radius 9–28, n=2,845–261,815): confirms the GG/GI nucleation mechanism is a genuine bulk phenomenon across a 92× range in N, not a finite-size artifact, resolving an initially-apparent radius-16 outlier as a scan-depth artifact once corrected; closes GG.6's own flagged replication requirement more strongly than expected — two fresh non-Fano embedding triples reproduce bit-identical results, revealing an automorphism-level invariance across all 28 non-Fano triples; and finds the Postulate-3 literal-coefficient tie-back structurally inaccessible on this construction, since φ is identically zero for every triangle, a direct consequence of every substrate point sharing one fixed 3-dimensional subspace. GK redesigns the embedding (a generic orthonormal 3-frame rather than coordinate axes) to escape that degeneracy, enabling the three-candidate literal-coefficient comparison on this construction for the first time: φ^2 is widest in all four tested instances, matching the pattern already established on the by1 orbit substrate; associator-norm and $|\varphi|$ are tied in all four, and the tie is proven structural (forced by any single-3-plane embedding, not a property specific to this con-

*struction); the smaller margin here versus by1 is traced end-to-end to the node-aggregation step and quantitatively explained by measured within-node triangle correlation. GL tests AZ.13's discharge condition on this same substrate (M has rank exactly 3/7, anisotropy exactly 1.0, confirmed at all 7 scales and independent of embedding frame — though identified as a near-tautological consequence of the same single-3-plane-embedding property behind GJ and GK's findings, not an independent discovery) and composes it successfully with genuine Postulate-3 dynamics, the cleanest of three substrate types tested to date — meaning five separate theory requirements (growth, negation-closure escape, nonzero associator, bulk-scale competitive dynamics, and AZ.13's discharge condition) are cleared simultaneously on one construction for the first time in the project's recorded history. GM returns to Postulate 3's literal-coefficient question directly, pursuing both routes previously flagged as the only paths to genuinely new evidence. A sharpened postulate-patching argument, built on the CU.9 identity, is closed negative via a rigorous SO(7)/G2 orbit test — confirming Addendum DN.1–2's earlier finding by an independent method — though a genuine, if modest, two-tier sharpening of the patching criterion itself survives the negative result. Route two delivers, for the first time, an exact closed-form probability distribution for all three literal-coefficient candidates (a Dirichlet/Beta decomposition in the style of Addendum AF's own harmonic-weight result, verified against direct computation), which alone reproduces the known skewness ordering from pure algebra; confirmed directly on a bit-exact reconstructed by1 substrate; and validated by reimplementing Addendum DR's cascade theory from its recovered algorithm, run cold on real triangle data with zero fitted constants, correctly predicting the real discoverability windows' ordering, near-tie structure, and order of magnitude, with a diagnosed (not mechanism-level) residual gap traced to an unrecovered classification-threshold constant. A self-caught correction is logged in place (GM.14–17): an earlier claim within this same addendum, that the derived distribution explains a previously-reported real-substrate correlation (Addendum CU.1's −0.33 figure), is retracted — that figure was always a synthetic calculation over random triples, never a substrate measurement. GN closes the resulting open item: the real-substrate correlation is measured for the first time (−0.13), confirmed general and essentially N-independent across six independent substrates spanning a 10× range, and its actual mechanism — genuine local k=7-NN graph combinatorics, not the isotropic clustering approximation GM tested and found wrong-signed — is built directly and validated, reproducing the correct sign and order of magnitude down to local populations as small as 30–50 points. No formal status-tag changes anywhere in the theory; Postulate 3 remains POSTULATED, UNDERDETERMINED among {associator-norm, |φ|, φ²}; AZ.13 remains tied to substrate type; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. Documentation note, logged per the no-silent-corrections standard rather than silently fixed: this document's own front-matter version header read "**Version 40**" at the start of this consolidation, one version behind the version-history text below it (which already documented v41's own GE–GI fold-in in full) — the same stale-header defect class previously caught and fixed at v34, evidently reintroduced since. Corrected to read "**Version 42**" directly, skipping a separate "**Version 41**" intermediate header update since v41's own consolidation is already fully documented in the text; a similarly corresponding "**v41 (this document) –**" narrative sentence appears to be genuinely absent from this run-on version-history paragraph (v41's own updates were recorded only in the front-matter bracketed pointer notes, not as a new trailing sentence here, unlike every consolidation from v20 through v40) — flagged here as an observation for a future audit pass, not backfilled unilaterally in this consolidation. See SFT_Changelog_v42.md for the full diff.*
New at v43 (Addenda GO, GP, GQ): three connected threads, one session. GO extends the DH/DI kink-crossing dynamical protocol to φ² for the first time, running DH's own long-flagged multi-triangle generalization test after two methodological dead ends are caught and logged in place (a wrong dynamical-object substitution; an undetectably coarse whole-substrate observable) — a two-stage joint-criticality screening finally isolates a valid protocol, finding φ²'s local response tightly bounded (3.6–4.7%) across seven properly-targeted triangles while |φ|'s spans three orders of magnitude including two genuine derivative discontinuities — suggestive, not forcing. GP extends the DB.5/DJ invariant-classification technique to the full tensor

algebra for the first time: φ is proven the unique G2-invariant trilinear function at degree 3, full stop, a genuine strengthening of DB.5; at degree 6, where φ², associator-norm², and Gram_det all live, the full invariant ring is classified exactly (6-dimensional, explicit basis), with φ² and associator-norm² both proven independently necessary — a complete, negative result for the original forcing-theorem hope, alongside a structural identity (Gram_det = φ² + ¼·associator-norm²) independently re-derived and then correctly re-attributed to Addendum CU.9 rather than reported as new. GQ resolves Addendum EF.7's long-standing flat direction (Postulate 6/8's "which basis of L" question): for BD's dynamically-selected associative triad and n̂ orthogonal to it, the associator [triad_i,triad_j,n̂] over the three triad pairs gives an exact orthonormal basis of L (Gram=4I_3), proven directly via Cayley-Dickson doubling rather than left as a numerical observation, and replicated across four independent full dynamical realizations using the real CR.3/BD potentials. A chain of corrections is logged in place: Addenda GA–GD's own claim to have already closed EF.7 via an unrelated argument about the seed point e_0 rested on an unverified premise (GA.5) never independently re-checked, and does not survive contact with this addendum's direct construction — GB–GD's own results about e_0's irreducibility remain valid, answering a genuinely different question. STATUS-TAG CHANGE: Postulate 6/8's "which 4 of 7" question upgraded from CONJECTURE/OPEN to DERIVED. See SFT_Changelog_v43.md for the full diff. v44 (this document) — consolidation folding in Addenda GR, GS, GT, and GU (four addenda, two sessions, 2026-08-13 through 2026-08-16). GR builds $H_3(Z)$, the rank-3 Heisenberg group, dimension exactly 7, embedding natively into all seven octonion imaginary directions with no coordinate compression — avoiding the shared root cause behind GJ/GK/GL's near-tautological findings — confirming the GG/GI nucleation mechanism replicates in full and composing AZ.13's discharge condition cleanly via the actual 35-parameter joint dynamics on a genuinely non-degenerate M-tensor (full rank 7/7, anisotropy 0.386). Extending to six constructions in total to isolate the variable behind a genuine decay-vs-plateau distinction, a tentative dimension-threshold hypothesis is proposed, directly falsified by a further construction, and retracted in place (GR.19→GR.22) rather than removed — the honest net position is that only $H_1(Z)$ plateaus among six independent constructions now tested, for a cause not yet identified. GS tests four hunches proposed by two different external LLM collaborators against H.17 and AZ.13 across two sessions — a correlation-structure locality diagnostic, growth measured on the dynamically-selected surviving sector rather than the raw substrate, composing the M-tensor mechanism with FV's polynomial-growth cone, and a claimed non-associativity bounding mechanism — all four tested directly and found negative, each via an independent method with its own finite-size or artifact control, extending the project's honest-branch negative case rather than opening a route through it. GT builds G2 = Der(O) directly from the project's own octonion structure constants and proves, rather than continues to observe empirically, that G2 has no nonzero fixed vector in Im(O) and acts transitively on S^6 — so no G2-invariant construction, of any kind, can select a preferred direction — and that Λ^3(Im O) has an exactly 1-dimensional G2-invariant subspace, supported on the Fano lines, the classical associative form φ recovered independently rather than assumed. This unifies Addenda FR, FV.7–FV.9, GL.4, and GS's own cone result as one fact viewed four times, and supplies a predictive filter for future proposed mechanisms: check the representation content before simulating. GU, opened as a direct follow-on to GT, catches and resolves a substrate-naming collision: Addendum EZ's "orbit substrate," cited since Addendum FP as the evidence that AZ.13's discharge mechanism is close to forced on orbit-type substrates, is traced through the record and found to be a different, genuinely interesting but differently-named construction — a literal orbit under a fixed rotation generator, not by1 as used consistently throughout Addenda BY–CB, CD, DO–DP, and GK since well before EZ was written. Recomputed on the substrate the project actually and consistently calls by1 (8 seeds × 3 weighting schemes), AZ.13's discharge mechanism is weak, not forced; BU/BV's lattice figures, independently reverified for parity, are found to sit in the same weak regime rather than the opposite one EZ.9's dichotomy described. Addendum FP.0's licensing decision for CS.4 itself rests on independent textual grounds (Postulate 7's own wording) and is unaffected; Addendum EZ is

not edited or retracted, standing as a correct result about a differently-named construction. No formal status-tag changes this consolidation. See SFT_Changelog_v44.md for the full diff. v45 (this document) — consolidation folding in Addenda GV, GW, GX, GY, GZ, HA, HB, and HC (eight addenda, three sessions, 2026-08-17 through 2026-08-25). GV tests whether the Bombelli–Henson–Sorkin no-go theorem for discrete Lorentz invariance bears on H.17's own adjacency-before-geometry fork; found not to apply directly (the theorem presupposes a continuum being sprinkled, which this project's substrate never has), but its underlying mechanism identifies a second, currently-untestable symmetry-based obstruction axis alongside the counting-based one GF.2 already proved — a redirection toward order-relation primitives this suggests is tested and closed: the natural algebraic route (split octonions) is already closed twice in this project's own record, and causal set theory's own manifoldlikeness problem shows the redirection would not evade the tension regardless. GW tests two nilpotency-invariant hypotheses (center-dimension and derived-subalgebra ratios) against GR.24's own open question — why $H_1(Z)$ uniquely plateaus among six tested constructions — and falsifies both by direct computation, with $\mathbb{Z}^3$ (identical embedding to $H_1(Z)$, ratio 1.0, yet decays) as the cleanest counter-example; separately finds that $H_3(Z)$'s own AZ.13 M-tensor anisotropy figure (0.386, GR) is not scale-invariant under its own reported convention, climbing steadily across every tested radius with no sign of convergence — a genuine, previously-unflagged gap in that result, not a reversal of it (see the AZ.13 entry, above, New at v45). GX builds and tests the sharpest available counter-example to GW's own hypotheses — N(3,2), the free 2-step nilpotent group on 3 generators, matching $H_1(Z)$'s nilpotency class exactly while exceeding its center ratio — and finds it decays regardless, extending the tested-construction count to seven with $H_1(Z)$ still the sole plateau. GY and HC together complete GR.23's own long-flagged, previously unresolved question: whether GJ's original $H_1(Z)$ plateau claim survives independent reconstruction at full rigor out to its own largest tested scales. GY reaches radius 20 and 24 and finds the trend still declining, below GJ's claimed floor; HC completes the arc at radius 28 (N=261,815, exactly matching GJ's own largest documented point, after a single continuous computation of approximately 121 hours — the longest in this project's history) and finds the local decay exponent collapsing monotonically across the five largest radii (0.462→0.428→0.324→0.174→0.055), the genuine mathematical signature of an approaching asymptote rather than continued decay: $H_1(Z)$ does plateau, upholding GJ's qualitative claim, but at a materially lower value (an estimated 1.5–1.7%, against GJ's originally reported 2.05–2.39%) than previously recorded — explaining rather than merely disputing GR.23's own partially-reconciled discrepancy. GY's own interim speculative reframing (that $H_1(Z)$ might be merely the slowest-decaying construction rather than a qualitative exception) is retracted in light of HC's completed result — the exponent-collapse signature is not consistent with ordinary slow decay. GR.24's own open question is thereby sharpened rather than resolved: the plateau is now confirmed on firm empirical footing, but no mechanism explaining it has yet been found, and GW/GX's own tested hypotheses (nilpotency class, center ratio, derived-subalgebra ratio) are now decisively excluded as candidates. Separately, GZ (parallel session) tests a genuinely new construction paradigm — growth admission gated by Postulate 3's own dynamics directly, rather than by a fixed geometric rule — and finds it asymptotically self-defeating on a compact-manifold substrate via k-NN coupling saturation, a further independent manifestation of this project's own growth-vs-selection tension (cf. AX.10). HA (parallel session) tests nonlocal Jordan–Wigner fermionization, motivated by an external Rydberg-array CFT-spectrum result, as a candidate matter-emergence mechanism independent of Postulate 8; after two further self-correction passes logged in place, the honest surviving finding is that $H_1(Z)$ has the lowest fermion-string bandwidth exponent of the tested constructions by an ordinary growth-degree effect, not (as an intermediate pass had claimed) a special interaction between the octonion-embedding pipeline and this specific group. HB (relettered from an independent GV assignment caught via this project's own letter-collision check) tests BFSS Matrix Theory, motivated by the division-algebra classification of minimal super-Yang-Mills (D=3,4,6,10, the octonion case among four), as a further candidate escape from H.17's own graph-based dichotomy; an apparently exciting dimensional-collapse signal is caught and cor-

rectly retracted via calibration against synthetic point clouds of known dimension before being trusted — no genuine collapse is found on the tested construction. No formal status-tag changes anywhere in the theory this consolidation; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY throughout, on firmer footing than before. See SFT_Changelog_v45.md for the full diff. v46 (this document) — consolidation folding in Addenda HD, HE, and HF (three addenda, one continuous arc, 2026-08-24 through 2026-08-26 — HD's own external-hardware run completed 2026-08-25/26). HD runs a genuinely thorough sweep on GR.24's own open question (why $H_1(\mathbb{Z})$ uniquely plateaus): builds F(2,3) (free nilpotent group, rank 2, class 3) to complete the rank/class structural design GX began — catching and fixing a real construction bug along the way (a naive Lie-bracket lift is not the actual group commutator) — and finds it decays cleanly across a 72× N-range, run to radius 8 on the architect's own external hardware after radii 9–10 proved impractical within available time (a genuine hardware ceiling, reported honestly rather than extrapolated past). Redesigns the M-tensor-anisotropy diagnostic to remove an ambient zero-padding artifact that had forced every sub-7-dimensional construction to read identically; the corrected version still finds no discrimination, and identifies why more decisively than a flat negative would — raw-covariance anisotropy is mechanically forced toward 1 with radius for any non-abelian construction, so the metric was only ever tracking abelian-vs-non-abelian status, already known not to split plateau from decay. Two dynamics-level hypotheses (h_ijk distribution width; h_scale coverage) are also tested for the first time and both falsified, the second only after a seed-count confound is caught and resolved via a seed-matched external rerun. Eight constructions and six candidate discriminators are now excluded; $H_1(\mathbb{Z})$ remains the sole plateau, with the negative space around GR.24 now unusually complete rather than merely extended. Separately, HE catches a real instance of this project's own working-memory failure mode: an external literature-survey session had proposed four candidate H.17-resolution mechanisms as though fresh, but direct project-knowledge search finds two (causal sets, Trugenberger/Kelly-style symmetry breaking) were already closed weeks earlier, unknown to the survey since it worked from web search alone. Family D (entanglement-derived adjacency) is then read from primary sources and tested in two parts on the real octonion substrate: correlation tracks distance cleanly when a graph is already presupposed, but is indistinguishable from noise once the presupposed graph is genuinely removed and every triple becomes a valid coupling triangle — Family D does not escape H.17's circularity. HF closes the remaining family (C, on-shell symmetry) by structural argument rather than simulation, following the precedent Addendum AN set: after reading its primary source in full, finds a category-level type mismatch (the mechanism presupposes an already-existing continuum spacetime, exactly the kind of prior geometric existence H.17 asks about) plus, even charitably reframed, no content beyond what SFT has already tested (Addendum AX.10; Addenda CR/BC/FJ for the neighboring AZ.13 question). With HE and HF, all four of the external survey's proposed H.17-resolution families are now closed. Separately, this consolidation also completes a documentation backfill unrelated to HD–HF's own research: the previously-lost Addenda AA–AH thread (session date 2026-07-17, recovered 2026-08-26) is restored to the record in full — see Part III for the complete account of what it contains and why three already-consolidated addenda (FB.1–FB.2, FM, GE.2) had been citing it by letter with nothing behind the citation. **v47, immediately following, completes the matching second half of that same recovery effort:** the previously-lost Part III §G snapshot and the AO–AX thread (nine further addenda, session-dated 2026-07-18/19, cited at length in Part V §2 and the Part VI register but never folded into any consolidation) are restored in full, and the companion archive's own long-standing FV–GU Contents-index gap (flagged and left open at every consolidation from v34 through v46) is finally closed — see Part III below for the full account. No new physics results and no formal status-tag changes result from either backfill; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY throughout. See SFT_Changelog_v47.md for the full diff. v48 (this document) — consolidation folding in Addenda HG, HH, HI, and HJ (four addenda, one continuous session, 2026-08-27). HG builds and tests this project's first genuine attempt at FR's own flagged "fifth, unconstructed paradigm" for H.17 — an associator-directed accretion walk,

using the associator generatively (as a displacement direction) rather than as a static filter for the first time. The natural linearizations of this idea are proven exactly skew-symmetric (a fifth independent confirmation of the project's compactness/rotation obstruction, joining AN, Z.15, FA.5, FB.5), and a normalized-step variant escapes that specific trap, checked and found genuinely negation-odd per realization — a real, non-hand-waved instance of ensemble-symmetric/realization-asymmetric structure distinct from FR's own closed mechanism. A dynamical collapse onto a lower-dimensional subspace is caught and fixed (periodic generator resampling) before being used for a growth-class test. Once fixed, the resulting vertex set still gives log-diameter under both independent-triple-sampling and sequential/correlated adjacency (an eighth structurally independent construction landing in this class), sharpening Addendum FE.10a's own hypothesis: the log-diameter default now looks driven by the associator-norm-threshold adjacency rule itself, not by vertex-set provenance or sampling architecture. HH is a documentation-only session, catching two stale items in the same external session-priorities summary before treating it as a reliable basis for further work: n=4 uniqueness's flagged "one identification still owed" was actually closed at v26 (Addendum DK) and again at v43 (Addendum GQ, unflagged by the summary at all), and Postulate 8's rotation-only covariance argument, flagged as "never completed," was actually completed at v13 (Addendum BB) and found the opposite of the summary's framing — independent of H.17, resolving into a permanent non-forcing result between Route A and Route B, not an incomplete task. No content in this document's own body was stale; both corrections are scoped entirely to the external summary. HI returns to Postulate 3's literal-coefficient question and, checking the specific gap Addendum U's own recommended next step names, finds AL/AM/AR (cited as having executed it) actually address a different, superseded candidate framework (the pre-CT R/S/T mechanism) and never touch the current three candidates at all. Building the G2-Casimir decomposition of $\Lambda^3(\mathrm{Im}\,\mathbb{O})$ from scratch, HI finds and verifies an exact new identity ($\varphi^2+¼|\text{associator}|^2=\text{Gram-determinant}$), catches and corrects two of its own errors in the course of interpreting it, and proves via Schur's lemma (alternativity forces the associator to be alternating, hence G2-equivariant as a map to the 7-dimensional representation alone) that the theory's two existing candidates already exhaust the full invariant-theory content of a triple's G2-decomposition — there is no fourth independent invariant to find, closing off rather than opening a candidate escape route, while leaving Postulate 3's own three-way underdetermination exactly as open as before. HJ returns to H.17 with a diagnostic never applied to HG's own constructions: Addendum DU's finding that conductance and diameter are independent properties, with DS.7/DS.8's sequential gluing showing a genuine bottleneck (exponent -0.61, $R^2=0.977$) despite the same log-diameter verdict as DS.2's flat-expander case. Checked directly (pipeline validated against DU's own published control figures first): HG's independent-sampling construction reproduces DS.2's flat-expander signature almost exactly on a completely different vertex-generation mechanism, while its sequential construction shows a real but much weaker decay (exponent -0.10, $R^2=0.59$) than DS.7/DS.8's own clean result — sharpening rather than merely failing to reproduce DU.2's finding: the earlier bottleneck-preserving property looks specific to AW's lattice-ordered vertex set, not general to sequential-vs-independent adjacency architecture. Before proposing this test, HJ also confirmed the Trugenberger/Kelly dynamical-rewiring route (Addenda K–N) remains correctly closed and was not re-opened without new motivation. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 3 remains POSTULATED, UNDERDETERMINED. See SFT_Changelog_v48.md for the full diff.
Legend: DERIVED (proven from postulates) · PARTIAL (mechanism identified, proof incomplete) · POSTULATED (assumed, not derived) · CONJECTURE/OPEN (unresolved) · DECIDED (a choice was made between live options, with reasons recorded) · RESOLVED (a surfaced tension was closed) · RETIRED/ELIMINATED (formally abandoned, kept for the record).

1 Part I — The Original GUET Formulation (Retired)

This section preserves the original 2026 formulation of the Grand Unified Emergence Theory in summary, exactly as first proposed, together with the three structural defects that forced its replacement. It is retained for historical completeness and honest accounting — not as a currently valid description of the theory.

1.1 1. Original Core Postulates

- The Pre-Geometric Substrate S: a discrete, non-spatial set of N algebraic configurations, independent of physical coordinates or a predefined Hilbert space.
- The Substrate Interaction Matrix W: a real, symmetric tensor W_ij governing pre-geometric compatibility between states s_i and s_j.
- The Primordial Transition: a global phase transition setting non-zero initial amplitudes A_i(0) ≠ 0 for all states, replacing the classical Big Bang singularity.

1.2 2. Original Master Field Equation

$dA_i/dt = (\Sigma_j W_{ij} A_j^2 - \gamma_i) A_i - \lambda A_i^3$

A first-order-in-time, non-linear pitchfork-bifurcation filter. Stable configurations ($\sigma_i > 0$) settle to $A_i = \pm\sqrt{\sigma_i/\lambda}$; *unstable configurations ($\sigma_i<0$) decay to zero. The surviving set $M = \{s_i \in S \mid \sigma_i>0\}$ was proposed as the physical universe.*

1.3 3. Original Claims (as first stated)

- Coordinate-free flow: $dA/d\tau = A \odot (W[A \odot A] - \gamma - \lambda[A \odot A])$, with $\tau = -\ln \Lambda$ an RG-like algorithmic scale.
- Emergent unitarity: linearizing around the attractor was claimed to yield $d\xi/dt = iH\xi$ from a Hermitian H derived from W.
- $E = mc^2$: energy as free substrate perturbations obeying the d'Alembert wave equation; mass as trapped informational vortices; c as the network transduction constant.
- A Grand Unification effective action combining Einstein–Hilbert gravity, Yang–Mills gauge fields (including a claimed SU(3) sector from algebraic triads), and a Dirac fermion term.
- Falsifiable predictions: linear Lorentz-invariance violation (LIV) in gamma-ray timing, a decoherence floor, and CMB polarization anomalies from the primordial transition.

1.4 4. Structural Defects That Forced Retirement

1.5 Defect 1 — Unitarity failure

The Jacobian of the linearized master equation is real and symmetric ($J = J^T \in \mathbb{R}$), not anti-Hermitian. A real-symmetric generator produces pure exponential growth/decay, never rotation — it algebraically cannot generate the claimed iH unitary evolution. This is not a derivation gap; it is a direct contradiction between the postulated real-symmetric W and the claimed emergent unitarity.

STATUS: ELIMINATED — real-symmetric W cannot generate unitary dynamics under any linearization

1.6 Defect 2 — Wave equation failure

The master equation is a strict gradient flow ($dA/dt = -\nabla V$ for a Lyapunov functional V built from W and γ). Gradient flows are globally non-oscillatory by construction — coarse-graining the spatial structure yields a reaction-diffusion equation, not the d'Alembert wave equation. The original $E = mc^2$ derivation via 'free substrate perturbations obeying the wave equation' is therefore invalid as stated.

STATUS: ELIMINATED — first-order gradient flow forbids wave propagation; same root cause as Defect 1

1.7 Defect 3 — SU(3) unification unsupported

Neither the scalar coupling path nor a tested representation-valued generalization of W can simultaneously satisfy the photon-sector decoherence bound ($\sim 10^{-16}$) and the strong-force coupling magnitude ($\sim 10^{-1}$ to 1) without contradiction. The representation-valued generalization independently reintroduces a multiple-limiting-speeds problem excluded by GW170817-class multimessenger bounds. The 'Grand Unified' framing — a single coupling structure producing all gauge forces — is not supportable.

STATUS: RETIRED — SU(3) gauge unification impossible under either tested coupling path; theory renamed SFT

Both Defect 1 and Defect 2 trace to the same root cause: the master equation is first order in time. Both were repaired by the same postulate change — promoting the dynamics to a second-order, damped-oscillator system with a conjugate momentum sector (see Part II, Postulate 3). Defect 3 is independent and structural; it is why the theory was renamed from 'Grand Unified Emergence Theory' to 'Survivor Field Theory.'

2 Part II — The Revised Normative Core (Eight Postulates)

The current, load-bearing formulation of SFT. Each postulate below reflects all corrections accumulated through every derivation session, including the very latest (Part III).

2.1 Postulate 1 — The Pre-Geometric Substrate S

Unchanged in its core wording from the original formulation: a discrete, non-spatial set of N algebraic configurations, prior to any notion of physical space, time, or Hilbert space.

Major revision, carried from v10 (Addenda Z–AJ, AN, AV–AX). A concrete candidate was proposed and substantially developed: S as (a discretization of) the unit imaginary octonions, with $\mathrm{Aut}(\mathbb{O}) = G_2$. The competing free-algebra reading is ELIMINATED by proof (Addendum AJ). The discrete-subgroup-of-Aut(𝕆) discretization route is CLOSED (Addendum AN). Two independent, genuinely polynomial-growth lattice discretizations were found instead (Addenda AV, AW) — discrete subrings of the algebra itself, one of which (AW) coincides exactly with the already-established $\mathbb{Z}7 \times \mathbb{Z}3$ symmetry. AW.5 identified a sharp tension between this lattice picture and the bounded-orbit picture underlying every dynamical construction since Addendum Z. Addendum AX found the orbit side can achieve genuine polynomial growth on its own (the full 6-dimensional orbit, not a single curve) but that connecting this back to the actual octonion coupling kernel forces uniformly-positive coupling, precluding genuine selection — the growth-vs-selection trade-off.

Substantially advanced at v16 (Addenda BT–BW). Both of AW.5/AX.10's demonstrated trade-offs have at least one working construction, in this project's own record, where they

do not bind — using the triangle-term mechanism (Addenda BE–BQ), unavailable when either trade-off was first identified. Addendum BT reconstructed AX's own orbit-based substrate directly and found the triangle term (unlike AX's pairwise coupling) produces genuine, dynamically-confirmed, spatially coherent partial selection, at zero pairwise coupling. Addenda BU–BV built a genuine, verified octonion lattice ($\mathbb{Z}^7$, closed under this session's own independently-constructed structure constants) and found growth (classical/structural, not merely sample-verified) and triangle-term selection coexist there too — first via normalized point-directions (BU), then confirmed to hold with the full, unnormalized octonion product directly (BV), closing an honestly-flagged gap in BU's own construction. Addendum BW extended both findings toward Addendum AX's own N=100,000 scale, properly controlling for a sampling-density confound found along the way: growth-class discrimination (power-law vs. exponential fit) remains genuinely ambiguous even with the confound removed, mirroring the same ambiguity Addendum AQ needed Addendum AS's much larger effort to resolve on the unrelated random-graph substrate — not yet undertaken here. The selectivity phenomenon itself persists robustly across a 30× increase in N, but its spatial coherence (the more physically striking part of BT–BV's results) substantially degrades over that range (72%→14% graph-contiguity). Neither AW.5 nor AX.10 is fully dissolved by this work — both remain real, demonstrated tensions for the mechanisms and constructions they were originally shown on — but the practical question each raised (can genuine growth and genuine selection coexist on one substrate) now has real, working, if not yet scale- or instance-general, positive examples.

Substantially advanced this consolidation (Addendum CF): BT.4's own reported generic-IC result — genuine partial selection on the orbit substrate at zero pairwise coupling, triangle term alone — is now reproduced, not merely characterized as a working construction. Following CE.6's recommendation, conversation search recovered BT.4's and BV.3's literal dynamics code, revealing an initial-condition convention (generic ICs drawn from Uniform(0.3,0.8), not the Uniform(−0.3,0.3) every session since Addendum BY.1 had used) as the sole cause of a discrepancy this project carried for many sessions. Reproduced qualitatively on a reconstructed orbit substrate (4/4 seeds partial, matching BT.4's outcome class) and near-exactly on a substrate recovered and verified to an exact statistical match (5/5 seeds, matching BV.3's own reported figures at two of five tested points). A clean isolation test found the IC convention alone necessary and sufficient. This does not change AW.5's or AX.10's own status — both remain real, demonstrated tensions — but it does mean the working counterexample this postulate has relied on since v16 (Addenda BT–BV) is now independently reproduced and mechanistically understood, not merely a once-reported result this project's own later reconstructions had been unable to confirm.

Substantially advanced at v28 (Addenda DS–ED). A twelve-addendum thread tested this section's own standing discretization question from a new angle — whether adjacency itself, not just coupling values on a pre-given graph, can be derived purely from the octonion algebra's non-metric structure without presupposing an ambient metric — and found a genuine, if not yet geometrically understood, hyperbolic-like signature on a specific sequential-gluing construction, confirmed independent of continuous sampling under the most literal possible finite reading of this postulate's own wording (a fixed pool of exactly N pre-given configurations, no further continuous sampling at any point). Full technical account in Part III's H.17 entry, since Addendum AN already established these are the same underlying problem. AJ.6's own question — how the continuous candidate algebra discretizes into literal N elements — remains exactly as open as before this thread.

Further advanced at v30 (Addenda EH–EL). A return to this postulate's own literal wording reconnected the DS–EG construction family to a previously-closed mapping onto Trugenberger/Kelly combinatorial quantum gravity and found a genuine, verified, N-independent bimodal phase transition in the growth process itself, later shown to recur at two further

levels (Postulate-3 dynamics and landscape structure) and to be sharpest specifically near internal critical boundaries. Full account in Part III's H.17 entry. AJ.6's own discretization question remains exactly as open as before.

Further advanced at v31 (Addendum EM). The near-critical decoupling finding from v30's own EH–EL thread (a specific sequential-gluing construction's conductance signature weakening and δ-hyperbolicity signature strengthening as its own internal glue_bias parameter approaches a collapse boundary) was tested for generality across two structurally distinct counter-examples: a dimension-parametrized simplex-gluing family at d=3, and DS.2's static associator-filtered 2-complex (tuned toward its own percolation threshold instead of a collapse boundary). Both replicate the conductance-weakening direction; neither completes the δ-hyperbolicity sign-crossing that made the original d=2 finding notable, within the range each could safely be tested. Full technical account in Part III's H.17 entry. AJ.6's own discretization question remains exactly as open as before.

Further advanced at v35 (Addenda FE–FH). A direct, exhaustive test of whether some untested construction axis lets this postulate's own honest substrate escape H.17's log-diameter default: compactness is closed on general argument, independent-vs-correlated edge-sampling is closed on direct test, and a third axis (frontier breadth) yields a real, cross-validated slope effect proven — via direct null-model substitution — to be a property of sequential random-graph construction, not this postulate's own algebra. Full technical account in Part III's H.17 entry, including this thread's identification, for the first time, of the precise mechanism (an asymmetric recency boundary layer) behind the project's own diameter measure. AJ.6's own discretization question remains exactly as open as before.

Further advanced at v42 (Addenda GJ–GL). A systematic seven-scale study on the Heisenberg-group k-NN construction (Addenda GG–GI) confirms its nucleation mechanism is a genuine bulk phenomenon across a 92× range in N, and closes that construction's own flagged replication requirement more strongly than expected — an automorphism-level invariance across all 28 non-Fano embedding triples (GJ). A redesigned generic-frame embedding escapes a structural degeneracy found along the way and, composed with AZ.13's discharge condition and genuine Postulate-3 dynamics, clears five separate theory requirements simultaneously on one construction for the first time in the project's recorded history — growth, negation-closure escape, nonzero associator, bulk-scale competitive dynamics, and AZ.13's discharge condition (GL). Full technical account in Part III's H.17 entry. AJ.6's own discretization question remains exactly as open as before.

STATUS: POSTULATED — base algebraic form remains an assumption, per the postulate's founding wording; not derived from anything deeper. The free-algebra reading is ELIMINATED (Addendum AJ); the discrete-subgroup discretization route is CLOSED (Addendum AN); two independent lattice discretizations exist (Addenda AV, AW) in real tension with the bounded-orbit picture used throughout the theory's dynamical work (AW.5); the orbit side can achieve genuine growth but at the cost of selection dynamics when connected to the actual coupling kernel (AX.10). Both of these demonstrated trade-offs now have at least one working counterexample using the triangle-term mechanism (Addenda BT–BV) — on specific, unscaled, single-instance substrates (Addendum BW: growth-class scaling remains genuinely ambiguous at the tested depth, and spatial coherence of the selective mechanism does not survive a 30× increase in N cleanly). New this consolidation: BT–BV's own generic-IC claim, previously unreproduced by any later session (BY–CE), is now independently reproduced and its cause (an IC convention mismatch, not a substrate, coupling, or algebra property) fully identified (Addendum CF). The discretization question remains open, narrowed to a sharper, better-characterized set of subquestions than at any prior consolidation — and, for the first time, resting on a working example this project can independently reproduce rather than only cite. At v28 (Addenda DS–ED): a specific sequential-gluing adjacency construction found with a

genuine hyperbolic-like signature on two independent diagnostics, confirmed independent of continuous sampling under this postulate's own strictest finite reading — narrows but does not resolve AJ.6's standing discretization question, unchanged in status. At v35 (Addenda FE–FH): compactness and edge-sampling architecture both closed as escape routes from the log-diameter default, on general argument and direct test; a real, algebra-independent slope effect found and explained; the project's own diameter measure mechanistically identified for the first time (an asymmetric recency boundary layer). Formal status unchanged. NEW at v39 (Addenda FV–FW): eighth and ninth independent manifestations of the growth-vs-consistency tension, tested via Fano-plane-collinearity structure rather than graph-growth diagnostics. FV's positive-cone semigroup escapes FR's negation mechanism but shows null localization for symmetric couplings; a Fano-collinearity coupling produces genuine, dynamically-robust, boundary-confined selection. FW proves this exhaustively for the invariant-coupling class (exact trilinearity theorem; exact bulk-wall theorem, generalizing beyond octonions to any Steiner triple system; exact PSL(2,7)-order-168 symmetry floor) and shows a richer coupling's apparent escape from the bulk wall is generic geometry, not Fano-specific structure. See Part VI, H.17 register entry, for full detail. Formal status unchanged.

2.2 Postulate 2 — The Substrate Interaction Tensor

Upgraded from real-symmetric W_ij to a complex-Hermitian tensor C_ij = W_ij + i Γ_ij, where W is symmetric (real part) and Γ is antisymmetric (imaginary part). The antisymmetric structure is DERIVED, not assumed: Hermiticity of the linearized generator forces it once complex amplitudes are admitted. Translation invariance across the substrate forces Γ to be uniform, eliminating what would otherwise be a separate postulate.

Latest revision (Part III, §4): locality is generalized from strictly compact support to an exponentially-decaying coupling kernel, $|C_{ij}| \leq C \cdot \exp(-\gamma|i-j|)$, to accommodate the overlap-fermion route to fermion doubling. This is a wording generalization, not a new physical assumption — bounded-degree sparsity (Postulate 4) is preserved.

Sharpened at v31 (Addenda EN, EP). Addendum EN proved this postulate's own "translation invariance forces Γ uniform" claim (Part I §3.2's own historical construction) is specifically a statement about invariance under a continuous one-parameter rotation, not additive translation in the sense this postulate's exponential-locality clause actually needs — and that no nonzero antisymmetric coupling built as a scalar function of a single difference vector can have both properties simultaneously, a consequence of G_2's transitivity on spheres in $\mathrm{Im}(\mathbb{O})$. A constructive lattice-native alternative Γ, satisfying the weaker discrete $Z_7 \rtimes Z_3$ symmetry instead of full G_2, was built and verified — genuine, nonzero, and provably distinct from the orbit-based recipe — but shown not to reopen Postulate 3's stabilization question (see Postulate 3 below). Addendum EP tested whether this motivates replacing this postulate's pairwise coupling with a triadic (associator-based) one as primary: the associator's antisymmetry is forced by the algebra alone, a cleaner derivation than this postulate's own Hermiticity-derived route, but the associator inherits EN's translation-invariance obstruction exactly — newly explaining Addendum BF's historical elimination of the associator route as a special case of this general fact rather than a standalone empirical quirk. Addendum BU–BV's own normalization convention, previously unmotivated in this project's record, is now understood as the necessary (not arbitrary) fix — though it introduces a new, tested, and non-threatening origin-privileging effect (real in the static coupling field, r=0.52 on BV's own reconstructed substrate; negligible in BV's own actual dynamics, r=0.04). Neither addendum proposes or adopts a formal rewrite of this postulate's text; both sharpen the scope of its existing DERIVED tag.

STATUS: DERIVED (Hermiticity, uniformity) / POSTULATED (base algebraic form).

Scope sharpened at v31 (Addenda EN, EP): uniformity is DERIVED specifically for continuous-rotation translation, and DERIVED to be structurally incompatible with additive-translation invariance for any nonzero antisymmetric single-difference coupling — not previously distinguished.

Further tested at v32 (Addenda EQ–ET). The triadic-rewrite hypothesis EP left well-motivated but not actionable was drafted formally (EQ) and gated CONJECTURE/OPEN behind four prerequisites, then tested against a chain of corrections and new evidence. ER found EP.7's own framing of Addendum BW.6's coherence degradation as "unexplained" was itself incorrect — Addendum BZ had already resolved it — and retargeted the gate toward recovering BW's own literal h_scale parametrization. ES executed that recovery directly: a bit-exact reconstruction of the BW/BZ N=90,778 substrate under Addendum CF's own recovered dynamics, finding BZ's imported h=0.033 supercritical at this scale, with the genuine selective regime sitting lower, near h≈0.02–0.025. ET then re-tested EP.4's origin-privileging concern — found dynamically inert on Addendum BV's own substrate at EP.7 — on a second, structurally different substrate, and found the inertness does not generalize: the effect propagates substantially into survivor dynamics there (r=−0.088 to −0.368). Taken together, EQ's draft rewrite remains unadopted and its gate unresolved, now for a specific, well-evidenced reason (why one substrate is immune to origin-privileging and others are not) rather than an untested concern. This postulate's own DERIVED status (Hermiticity, uniformity) and POSTULATED status (base algebraic form) are unchanged by EQ–ET.

*Cross-reference: the downstream "**emergent unitarity**" consequence historically associated with this postulate's Hermiticity result (ξ=iHξ, Part I §3.2) is tracked under Postulate 3 and Addendum O below, not under this postulate's own tag.*

2.3 Postulate 3 — Second-Order Damped Dynamics

The master equation is promoted from first-order gradient flow to a second-order damped-oscillator system, introducing a conjugate momentum field and one new free parameter η:

$$d^2A_i/dt^2 + \eta\, dA_i/dt = (\Sigma_j\, C_{ij}\, A_j\, \bar{A}_j - \gamma_i)\, A_i - \lambda |A_i|^2 A_i$$

This single change simultaneously repairs both Defect 1 (a genuine Hopf bifurcation now exists, yielding a Stuart–Landau limit cycle and real oscillatory dynamics near threshold) and Defect 2 (the continuum limit now yields the d'Alembert wave equation rather than reaction-diffusion). The overdamped singular-perturbation limit recovers the original first-order equation as an approximation, with computable first-order corrections.

Major revision, carried from prior consolidations (Addenda O–U, AL, AM, AR, AU, AY, BE, BF). The standard $\sigma_i>0$ fixed point is a structural saddle (Addendum O, $\kappa=2\gamma_i>0$ exactly, general). Full nonlinear dynamics has no bounded nonzero attractor under the bare cubic equation (Addenda P–Q). A stabilizing correction beyond the bare cubic term is NECESSARY (forced by the theory's own stated design goal) but its SPECIFIC FORM is underdetermined — three mechanisms (R/S/T) each work but none is singled out by Postulates 1–4. Addendum AY proved, generally and exactly, that no coupling matrix W, however derived from Postulates 1–2's algebra, can ever stabilize the bare pairwise/same-degree cubic form — the associator (AM/AR) and Γ (AU) both fail for this same underlying reason. Addendum BE found a fourth candidate structurally outside AY's scope — a three-different-site triangle term, cubic rather than quartic in the potential — which stabilizes fixed points in regimes the bare pairwise mechanism cannot reach at all. Addendum BF tested and ELIMINATED this term's first candidate coupling-strength origin (a broad-inclusion-radius octonion associator, which grows with separation and structurally excludes locality).

Major revision, carried from v15 (Addenda BG–BL). The triangle-term mechanism, with h left as a free parameter after BF's elimination, was tested far beyond its original toy case and substantially advanced. Addendum BH found it generalizes to sparse, heterogeneous

networks and produces genuine, dynamically-confirmed, spatially LOCAL survivor sets — the first Postulate 3 candidate of any kind to do so — scale-stable N=14–55, with an exact mechanical account of its parameter-space boundary. Addendum BI reopened the associator route to h's origin, restricted this time to genuine sparse-graph triangles (local by construction, unlike BF's broad-inclusion-radius sum), and found it reproduces BH's independently-found attractors exactly, once scaled. Addendum BJ traced the resulting scale factor to a parametrization artifact and closed it with no fudge factor, using the network's own natural O(1) angular scale. Addendum BK found partial (7/10) robustness against the one remaining free choice (which g2 generator defines the orbit direction). Addendum BL audited this whole arc externally per standing policy; two of the audit's three central claims did not survive direct checking, and one legitimate open item was confirmed: the natural-scale argument's first-principles status.

Major revision at v16 (Addenda BM–BQ): the flagged open item is closed, R/S/T are shown not to share the triangle term's locality property, and BK.4's generator-selection question is closed and validated out-of-sample.

Addendum BM closed BJ.6/BK.4's flagged open item — the natural-scale argument's first-principles status — deriving the O(1) angular scale from Postulate 1's own substrate geometry rather than empirical calibration. Two results: an exact algebraic theorem (the octonion associator, a totally antisymmetric trilinear form, vanishes identically on linearly dependent triples — confirmed numerically to machine precision on planar orbits at every tested angle) explaining why non-planarity is essential to the mechanism and automatic for G2/octonions specifically (G2 acts irreducibly on Im(𝕆), admitting no invariant 2-plane); and a closed-form covering-radius formula, using only Im(𝕆)'s fixed dimension (6) and the network's own (N,k), that independently matches BJ.2's empirically-calibrated scale to within 4% and lands within the dynamically-required associator-magnitude window with zero circularity.

Addendum BN tested whether R, S, and T — the theory's original three Postulate 3 stabilization candidates — share the triangle term's locality property, on the same network class and validated methodology. They do not: R and T produce only spatially uniform/dense stable states near their own existence boundary; S produces no stable state at all in the tested window; none produce anything resembling BH's local survivor sets. This validation required a logged mid-session correction: the first attempt used guessed parameters and an unvalidated search methodology, and failed to reproduce even the triangle term's own known result as a sanity check — flagged PROVISIONAL rather than reported as final. The architect supplied the original Addendum BH document directly; cross-referencing it against the source session transcript recovered the actual parameters (γ=λ=1.0, a custom network construction, not the values initially guessed), and the corrected reconstruction exactly reproduced Addendum BH's own reported profile. R/S/T's non-locality was then reconfirmed on this fully validated substrate and search method, with a mechanical explanation: R/S/T are single-site self-limiting corrections with no spatial coupling structure, unlike the triangle term's genuinely multi-site, non-pairwise form — locality requires exactly that structure.

Addenda BO–BQ closed Addendum BK.4's flagged generator-selection question. Addendum BO built an independent 14-dimensional g2 matrix representation from scratch (verified two ways: antisymmetry and the exact derivation property, both to machine precision) and reproduced genuine locality end-to-end using entirely independent code — an independent replication of the BI–BK mechanism, not a re-run. Addendum BP executed BO's own recommended follow-up (per-generator w-optimization rather than one shared test point) and found the true success rate is 9/12 (75%), matching Addendum BK's own 7/10 figure, with a real, cleanly-separated selection criterion — but its direction corrected an earlier, confounded preliminary signal from BO, logged explicitly as a correction rather than silently revised. Addendum BQ then validated this criterion properly out-of-sample: predictions locked in from the minimum-triangle-associator statistic alone, before any dynamics were run, scored 9/10 correct with zero false positives against fresh generator draws — upgrading the criterion's status from CONJECTURE to DERIVED, physics-level validated.

Major revision at v16 (Addenda BR–BS): the generator-selection criterion, honestly, does not generalize across topology instances. Addendum BR tested the validated criterion on three fresh networks (a new topology instance at the same N, a smaller N matching Addendum BH's own original scale, and a larger N matching Addendum BH's own later scale) and found the absolute threshold (0.6) does not transfer — refuted by a clear counterexample (a high-magnitude generator failing outright on a fresh instance) and by both real successes on a third instance falling below that threshold entirely. Addendum BS then tested a network-relative (percentile-rank) version of the criterion, proposed as more promising, and found — after explicitly catching and correcting an initially-misleading pooled statistic that was dominated by the original derivation network's own strong in-sample effect — that this relative form does not survive testing on genuinely fresh networks alone either (r=0.04, p=0.87, no effect). BK.4's original question is genuinely closed for the network it was asked on; the general form, across topology instances, is reframed as a separate, larger open problem, likely requiring a per-triangle rather than per-generator criterion.

Major revision at v17 (Addenda BX–CB): the triangle term's locality is now mechanistically understood, not merely empirically confirmed — and the mechanism connects Postulate 3 directly to H.17. Addendum BX tested whether Addendum L's shared-neighbour resource-competition bias (a Postulate 4 combinatorial argument) and the triangle term's own locality (a Postulate 8 algebraic argument) are secretly one phenomenon rather than coincidentally similar. On constructed toy substrates with triangle count, degree sequence, and edge count held fixed (isolating overlap concentration as the only free variable, after an initial confounded design was caught and corrected), overlap concentration measurably and reproducibly increases the triangle term's propensity to produce local survivor sets — replicated at two independent (N, triangle-count) configurations, with the primary signature a widening of the supporting (w,h) parameter region rather than a change in the mechanism's own capability ceiling. Addendum BY confirmed this transfers to the theory's actual H.17 substrate: overlap concentration is both an existence condition for seeded partial selection and a containment mechanism keeping the resulting activity spatially bounded, with distant overlap-anchored clusters shown to coexist independently and superpose exactly. Addendum BZ then executed a specific prediction BY.3 registered before testing it, at Addendum BW's own N≈91,000 scale: a mechanism-constructed multi-cluster state reproduces BW's reported 14% global contiguity almost exactly (0.171) while per-component direction-coherence is not merely preserved but exceeds every prior reference level in the record (0.893 mean, vs. BT's 0.53) — evidence that BW.6's coherence "degradation" is the global contiguity statistic failing to see a multi-cluster state, not locality genuinely failing at scale, with full adjudication against BW's own literal state still pending recovery of its exact parametrization. Addendum CA tested and eliminated the leading candidate explanation for a separate, long-standing discrepancy (generic random initial conditions failing to reproduce the BT/BU/BV record's reported partial-selection states) — a signed rather than magnitude associator reduction, tested via the natural G2-invariant 3-form on both substrates — while logging a genuine new phenomenon as a side-finding: under the signed coupling, ignition threshold from generic initial conditions is seed-dependent, a disorder signature with no counterpart under the magnitude coupling. Addendum CB explained that side-finding mechanistically: every run passes through a sharp activity bottleneck, and the surviving nucleation sites are significantly overlap-enriched (z=+7.29 against a degree-matched null, n=1,867, replicated across 10 seeds, after an initial global-cubic-form hypothesis was tested and cleanly disconfirmed) — unifying CA's disorder finding with BX–BZ's overlap mechanism as one phenomenon seen from the initial-conditions side.

Major revision at v18 (Addenda CC–CE): a refined predictor for the per-seed nucleation variance, and the first verified substrate/coupling-scale recovery progress on the BT/BU/BV discrepancy. Addendum CC introduced R_i (pairwise triangle-overlap — the count of triangle-pairs at a node sharing an additional common node, distinct from the raw triangle-count metric T_i) and found it predicts individual-seed ignition threshold more

strongly than T_i ($R^2 \approx 0.55$–0.64 vs. 0.38–0.60), with a topology-vs-algebra decorrelation test finding the octonion associator's specific magnitude contributes little beyond topology, given fixed topology. Addendum CD demonstrated that a prior session's exact generating code — not only its prose description — can be recovered via conversation search, and confirmed every CC finding directly on the real by1 substrate. Addendum CE applied the same technique to BU/BV, recovering their exact lattice substrate and coupling-scale convention to a verified match, and eliminated three further candidate explanations for the BT/BU/BV discrepancy, one (octonion basis convention) by a general Gram-matrix-invariance proof rather than only an empirical test.

Major revision this consolidation (Addendum CF): the standing BT/BU/BV generic-IC discrepancy is RESOLVED. Following CE.6's recommendation, conversation search located BT.4's and BV.3's own literal dynamics/integration code. The cause was an initial-condition convention mismatch: generic ICs drawn from Uniform(0.3,0.8) in the original code, versus the Uniform(−0.3,0.3) every session since Addendum BY.1 had used, including CC–CE earlier in this same session. Reproducing BT.4's and BV.3's exact configurations with the recovered code gives genuine partial selection from generic ICs — qualitatively on a reconstructed orbit substrate (4/4 seeds, matching BT.4's outcome class) and near-exactly on the verified-recovered lattice substrate (5/5 seeds, matching BV.3's own reported node counts almost precisely at two of five tested h_scale points). A clean two-factor isolation test found the IC convention alone necessary and sufficient — the second-order-vs-first-order dynamics question that looked live going in turned out not to matter. This is not a narrowing but a genuine resolution: not a property of the coupling, substrate, algebra, or dynamics order, all correctly cleared of blame across CA–CE's five addenda of honest elimination. Those addenda are not retracted — they characterize a real, reproducible IC regime, now understood to be distinct from BT/BU/BV's own rather than a failed reconstruction of it.

Major revision this consolidation (Addenda CW–CZ): the second-order dynamics' own stable-spiral signature is characterized for the first time, and reads as generic rather than distinguishing. Addendum O (superseding Addendum J's original complex-κ reading) established the correct signature for genuine decaying-oscillation behavior at a fixed point: a real eigenvalue κ of the conservative-force Jacobian satisfying $\kappa < -\eta^2/4$, with every other mode ≤ 0. O.3–O.4 searched for it on the bare-cubic/pairwise branch at scale and found it nowhere (0/197, 0/3485), explained structurally by Addendum AY's exact $\kappa = 2\gamma_i > 0$ result for that branch's dominant direction. Addendum CW found the same signature on the triangle-term mechanism's original toy 3-site case instead — present in ≈99% of a 44,579-point stable-fixed-point sample, robust across a wide η range, nonlinearly confirmed to 4 decimal places. Addendum CX extended this to a real, independently-validated Watts-Strogatz network (Addendum BH/BN/CK's own published fixed point, exactly reproduced via source-code recovery): complete (25/25) coverage. Addendum CY found the identical complete coverage on a second, structurally independent attractor on the same network, and — the theory's actual substrate — on the real N=3,000 octonion-orbit substrate itself (3000/3000), under the associator-norm coupling. Addendum CZ extended this further to φ (Addendum CU's second non-patching generator) on the same real substrate (3000/3000, matching the associator-norm's coverage exactly) and to R's (quintic) genuinely-stable branch on an exactly-regular network (25/25) — correcting, in the process, an error in Addendum CY.8's own framing (CY.8 had claimed Addendum AY forecloses R/S/T's dominant direction from ever being stable; this is false — AY covers only the bare cubic with no correction term, and R/S/T's own correction terms, per Addendum S's two-branch theorem, do genuinely stabilize a branch; R/S/T's actual disqualification is the postulate-patching argument of CG/CI/CT, not instability). Across six independent fixed points spanning three structurally distinct mechanisms (triangle-term/associator-norm, triangle-term/φ, R's quintic correction) and both toy and real substrates, spiral coverage is complete or near-complete in every single case. Addendum CZ.6's honest synthesis: this now reads as a broadly generic feature of $\eta = 0.8$ second-order dynamics wherever a genuinely deep, stable fixed point exists at all, not a property that favors

one Postulate 3 candidate over another — the genuinely informative result would be a stable fixed point that fails the criterion, not yet found anywhere it has been checked.

Major revision this consolidation (Addenda DA–DI): the non-patching class's literal-coefficient question (φ vs. associator-norm) was tested from five independent angles; two results stood up, two proposed arguments were downgraded on their own follow-up scrutiny, and one thread reversed an intermediate finding before it could stand. DB.3–DB.5 constructed g2 = Der(Im 𝕆) directly (dimension 14, confirmed) and computed its induced action on $\Lambda^3(\mathbb{R}^7)$ (35-dimensional): the G2-invariant subspace is exactly 1-dimensional and matches φ's own coefficients at cosine 1.00000000 — φ is provably the unique (up to scale) G2-invariant alternating trilinear form on Im(𝕆), upgrading CU.9's search-based finding to a classification theorem. Addendum DD tested signed φ directly against the theory's own recovered generic-IC dynamics (CF.1): 0/3,000 active nodes across 5 seeds and multiple magnitude-matched h_scale values, while |φ| (magnitude only, matching CU.3's own established convention) succeeds once calibrated — sign, not magnitude, is decisive, narrowing the viable literal-coefficient class from {associator-norm, φ} to {associator-norm, |φ|}. Two proposed arguments did not survive their own follow-up: Addendum DC found DB.5's uniqueness theorem lives entirely in Λ^3 (alternating forms) and does not reach the associator-norm, which is fully symmetric (Sym^3) — confirmed directly (all six permutations of a random triple leave |[a,b,c]| unchanged while φ's sign alternates exactly as expected); and Addendum DE found the "economy of structure" argument for φ's reuse in AZ.13/CR mostly restates DB.5 rather than adding independent evidence, since DB.5 already forces any alternating-trilinear-form use onto φ regardless of choice, and the two roles (a single-triad potential vs. a large-ensemble force sum) are not dynamically analogous in any case. A fifth thread tested whether the two candidates' confirmed non-differentiability (a genuine V-kink at each one's own zero locus, present regardless of which is used) has any physical consequence: Addenda DH–DI located exact zero-crossings via constrained optimization on real substrate triangles and found the kink propagates into large, unambiguous discontinuities in the full nonlinear dynamics' own aggregate output (signal-to-background ratios up to ~59,000x in the cleanest verified instances) — a real, load-bearing physical phenomenon, reopening a question Addendum DA had first found inert (no effect on Jacobian-in-amplitude stability analysis, generic-substrate proximity, or the theory's current continuum-limit derivations, all of which remain correct as scoped). An initial finding that this favored one candidate's differentiability profile over the other's (Addendum DF, single substrate) did not replicate under a properly powered 20-substrate test and was corrected in place (Addendum DG, Wilcoxon p=0.013 in the *opposite* direction); the same caution was applied to a follow-up propagated-severity comparison, which likewise found no statistically significant difference between candidates (Addendum DI, Mann-Whitney p=0.351, verified against a possible metric-contamination artifact before being accepted). None of DA–DI changes Postulate 3's formal status.

Major revision this consolidation (Addenda DL–DR): the non-patching literal-coefficient class is extended from two candidates to three, with the new candidate substantially advantaged and mechanistically explained. DL–DM closed out v25's kink-propagation thread: the by1 substrate was reconstructed bit-exact from recovered primary-source code, a full N=3,000 scan confirmed genuine kink propagation on the theory's real canonical substrate, an extended 20-substrate batch confirmed DI's null with much greater power (no discriminating advantage between associator-norm and |φ|), and a striking natural near-degeneracy flagged in |φ|'s survivor set was traced, on full scrutiny, to a mundane triangle-count/degree coupon-collector confound rather than any deeper physics. Addendum DN proved the associator is itself the unique (up to scale) G2-equivariant vector-valued trilinear map $\mathrm{Im}(\mathbb{O})^3 \to \mathrm{Im}(\mathbb{O})$ — a companion theorem to Addendum DB.5, which covers only the scalar (φ) case — closing the residual question of whether some other vector-valued candidate might exist; a proposed postulate-patching argument against |φ| (that its absolute-value operation is an ad hoc dynamical fix where the associator-norm's

non-negativity is "free") was tested and found to fail: both quantities are literally the same canonical octonion norm map, applied to different-rank objects. Addenda DO–DP found, validated, and stress-tested a genuine third literal-coefficient candidate, φ^2 — smooth everywhere (avoiding the associator-norm/$|\varphi|$ differentiability question by construction, not dynamical accident), non-negative by construction, and shown wider in relative discoverability window than both established candidates on every one of four independent substrates spanning a 10× range in N (roughly 1.5–1.9× wider each time), with the gap shown to exceed seed-to-seed noise by roughly an order of magnitude. Addenda DQ–DR then supplied the analytic account DP flagged as the only route to qualitatively new evidence: a single-triangle self-consistency closure gives an exact ignition threshold, first calibrated empirically to two fixed percentiles and validated out-of-sample on three fresh substrates (DQ), then superseded by a genuine zero-fitted-constant cascade theory — a self-consistent fixed-point iteration on the real per-triangle graph structure — matching the true nonlinear dynamics to within a few percentage points instead of tens, and validated at the node level (not merely aggregate counts), with amplitudes agreeing with the true dynamics to 3–4 significant figures (DR). φ^2's advantage is traced to its higher skewness relative to the other two candidates, itself a direct consequence of squaring an approximately sign-symmetric quantity. None of DL–DR changes Postulate 3's formal status or forces a selection among the three surviving candidates — this remains evidential weight of the same kind CU.4/CU.5 established as the project's working standard, now applied more thoroughly than to any candidate in the theory's history, not a forcing argument.

Implementing the standing Part V/Part VI pointer to Addendum BL.4 (queued since v22, re-flagged at CS.1). BL.4 refuted a claim that using the associator's vector norm instead of the raw signed associator "destroys G2 invariance" — correctly refuted, since the norm is exactly G2-invariant (G2 acts orthogonally on $\mathrm{Im}(\mathbb{O})$), and this stays refuted; external reviews that re-raise it are re-raising a settled point. Addendum DA.2's finding is a separate, later, and narrower claim: BL.4's own text concedes the norm is smooth only "away from zero," and DA.2 characterizes exactly what happens at that zero locus — a genuine, confirmed non-differentiability, distinct from any symmetry question. Addenda DH–DI subsequently showed this specific, narrower fact has real dynamical consequences. Any future external review raising the differentiability point should be checked against DA.2/DH–DI, not against BL.4, which addressed a different claim.

Extended at v31 (Addenda EN, EP). Addendum EN tested whether the constructive lattice-native Γ motivated by AW.5's sharpened tension (see Postulate 2 above) changes this postulate's standing stabilization verdict. It does not, and more generally than previously shown: Addendum AU's own chiral-rotating-wave instability result was demonstrated only for its own specific octonion-derived Γ; this consolidation's wide numerical sweep (32/32 tested cases across baseline, lattice-native, and randomized-parameter regimes) plus a partial symbolic confirmation (the radial self-restoring Jacobian entry collapses to exactly $-2\lambda r^2$ after self-consistency substitution, independent of Γ and W) shows the instability was never actually sensitive to Γ's algebraic origin — a generalization of Addendum AY's own W-independent cancellation mechanism to the chiral-state case. Separately, Addendum EP tested a concern this consolidation's own Postulate 2 work raised — whether the origin-privileging effect in a normalized-associator coupling field could bias survivor selection by position rather than genuine algebraic compatibility — directly against Addendum BV's own exact recovered substrate and dynamics (independently reconstructed and verified to match every one of BV's own reported figures). The static coupling-field bias is real and substantial ($r=0.52$) but does not propagate into BV's own actual survivor dynamics ($r=0.041$ for survivor status, $r=0.047$ for final amplitude) — very unlikely to explain Addendum BW's own reported spatial-coherence degradation at larger N, which remains open for a different, still-unidentified reason.

STATUS: DERIVED (the second-order promotion itself, and its repair of Defects 1–2) / DERIVED (the necessity of an additional stabilizing correction) / DERIVED, GEN-

ERAL (Addendum AY: no coupling matrix W can ever stabilize the bare cubic form, scoped to pairwise/same-degree corrections) / POSTULATED, UNDERDETERMINED among FOUR candidates — the triangle term remains the best-characterized: DERIVED, dynamically confirmed, spatially local, scale-stable N=14–55 (BH), with coupling strength h fully DERIVED including its natural angular scale from first principles (BM), R/S/T independently confirmed NOT to share its locality property on the same validated substrate (BN), a generator-selection rule DERIVED and out-of-sample validated for the specific network it was derived on (BO–BQ) but explicitly NOT general across topology instances (BR–BS), its locality mechanistically DERIVED as an overlap-concentration effect at v17 (BX–BY), with a falsifiable prediction for a related open coherence question executed and SUPPORTED (BZ), and its own flagship generic-IC result (BT.4, BV.3) independently REPRODUCED, with the prior discrepancy's cause fully identified as an IC-convention mismatch rather than any property of the mechanism itself (Addendum CF). Uniqueness among the four candidates remains as underdetermined as ever — mechanistic understanding and independent reproduction are not the same as a forcing argument, and none exists. At v23 (Addenda CT–CU): R/S/T are permanently disqualified from non-patching status (they cannot carry Postulate-1 content, CG/CI) and are, by the project's own postulate-patching standard, patches rather than forced mechanisms (CT); the non-patching class is fully characterized as having exactly one new generator, φ, with the associator-norm an exact derived combination of φ and pairwise data (CU.9). At v24 (Addenda CW–CZ): a stable-spiral dynamical signature (real $\kappa < -\eta^2/4$, otherwise stable) is characterized across six independent fixed points and reads as generic to the dynamics rather than distinguishing among candidates. NEW this consolidation (Addenda DA–DI): φ proven the unique G2-invariant alternating trilinear form (DB.5, a classification theorem superseding CU.9's search-based framing); signed φ shown dynamically nonviable, narrowing the non-patching class's literal-coefficient choice to exactly TWO candidates, {associator-norm, |φ|} (DD); the differentiability distinction between them shown to be real and dynamically consequential in general (DH–DI) but non-discriminating between the two candidates specifically, after due multi-substrate testing (DG, DI). None of this changes this entry's formal status. The four-candidate formal underdetermination is unchanged in kind, now with three candidates (R/S/T) disqualified in principle, a fully-mapped and provably-closed survivor class, that class narrowed to two viable literal coefficients, and a characterized (but non-distinguishing) dynamical signature and differentiability profile for both / NEW at v27 (Addenda DL–DR): the two-candidate class is extended to THREE. φ^2 — smooth everywhere, avoiding the associator-norm/|φ| differentiability question entirely rather than resolving it — found, validated across four independent substrates (N=600–6,000), and shown wider in relative discoverability window than both established candidates every time, by a margin (roughly 1.5–1.9×) shown to exceed seed-to-seed noise by roughly an order of magnitude (DO–DP). A companion theorem to DB.5 proves the associator is separately, independently the unique G2-equivariant vector-valued trilinear map on Im(𝕆) (DN), and a proposed postulate-patching argument against |φ| was tested and found not to hold — both |φ| and the associator-norm reduce to the same canonical octonion norm map (DN). φ^2's discoverability advantage is now analytically understood, not merely observed: traced through an exact single-triangle ignition threshold to a validated, near-fitted-constant-free cascade theory matching the true dynamics at the node level (DQ–DR). Still POSTULATED, UNDERDETERMINED — evidential weight, not a forcing argument; the literal-coefficient class is {associator-norm, |φ|, φ^2}, three ways / NEW at v42 (Addenda GK, GM, GN): φ^2's discoverability advantage replicated on a second, structurally independent construction (Heisenberg-substrate generic-frame embedding), with associator-norm/|φ|'s tie now proven structural rather than merely observed — forced by any single-3-plane embedding (GK). A sharpened postulate-patching argu-

ment (via the CU.9 identity, tested through a rigorous SO(7)/G2 orbit test) confirms Addendum DN.1–2's negative finding by an independent method, still not separating the three candidates; but an exact closed-form probability distribution for all three candidates is derived for the first time, explaining the skewness ordering behind φ²'s advantage from pure algebra and validated, with zero fitted constants, against the real by1 substrate's measured discoverability windows (GM). A same-session self-correction is logged in place: an earlier claim to explain a previously-reported real-substrate correlation is retracted (the figure was always synthetic), the true real-substrate correlation is measured for the first time, and its actual mechanism — local k=7-NN graph combinatorics — is identified and validated (GM/GN).

NEW at v43 (Addenda GO, GP). Route 3B (GO) extends the DH/DI kink-crossing sensitivity protocol to φ² for the first time, using the recovered F_clean transient dynamics rather than a mismatched mean-field substitute (an initial substitution error caught and logged before being reported). Properly targeted (joint local-criticality screening, not blind triangle selection), φ²'s response to crossing a kink locus is tightly bounded (3.6–4.7% across five live instances) while |φ|'s is highly direction-dependent, spanning three orders of magnitude and including two genuine derivative discontinuities — the expected statistical signature of a smooth function versus a kinked one, though not itself a forcing argument. Route 3A (GP) extends the DB.5/DJ classification technique beyond where it had ever been taken: the full space of G2-invariant polynomials on a triple (a,b,c), not only the alternating trilinear part. At degree 3, φ is shown to be the unique G2-invariant trilinear function of (a,b,c), full stop — not merely the unique alternating one — strengthening DB.5. At degree 6, where φ², associator-norm², and Gram_det all live, the full invariant ring is shown to be exactly 6-dimensional, with an explicit basis: $\{(a\cdot a)(b\cdot b)(c\cdot c),\ (a\cdot a)(b\cdot c)^2,\ (b\cdot b)(a\cdot c)^2,\ (c\cdot c)(a\cdot b)^2,\ \varphi^2,\ \|[a,b,c]\|^2\}$. φ² and associator-norm² are proven both independently necessary — the classification is complete, but the result is negative for the original forcing-theorem hope: no DB.5-style uniqueness exists at this degree, and the two candidates are shown co-equal rather than one forced over the other. A genuine structural fact surfaces en route — Gram_det = φ² + ¼·associator-norm² exactly — independently re-derived here via invariant-theory methods, though it is Addendum CU.9's own identity, not a new discovery as first drafted within GP before being checked against the archive and corrected in place.

None of GM, GN, GO, or GP changes this entry's formal status / ELIMINATED AS STATED (emergent unitarity on the standard $\sigma_i>0$ fixed point) / CONJECTURE, OPEN (whether it holds on some other attractor class) / POSTULATED ($\gamma_i>0$, standing assumption)

2.4 Postulate 4 — Sparsity / Bounded Degree (RMT constraint)

May's stability criterion, combined with random matrix theory, forces sparsity and bounded degree on C_ij: dense, unstructured coupling produces chaotic dynamics rather than a clean survivor manifold. This is DERIVED as a consequence of demanding a non-chaotic survivor set, not independently postulated.

New at v29 (Addendum EG.1–EG.3). Addendum B.6's early (v3-era) Cheeger-type dichotomy for this postulate — bounded-degree + polynomial growth forces a vanishing spectral gap; O(1) gap forces non-polynomial growth, framed as "any future substrate proposal must respect" — was checked directly against the v28 H.17 thread's own contextual-gluing construction (Addenda DZ/EA) for the first time. The construction satisfies bounded degree (≈4) and shows non-polynomial (log-diameter) growth, yet its spectral gap shrinks at a clean, measured rate rather than staying O(1) — a genuine third regime outside both of B.6's stated branches. Not a violation (both of B.6's implications remain individually true), but a real incompleteness in how the "permanent structural constraint" was framed, logged as an extension rather than a retraction. This postulate's own DERIVED status is unaffected.

New at v31 (Addendum EO): the bounded-degree half of this postulate now has a

second, independent, dynamics-free justification. Rather than relying on May's criterion (which implicitly presupposes some dynamical notion of stability already exists), Addendum EO showed bounded degree follows directly from Postulate 1's own lattice structure (Addendum AW): "minimal-vector adjacency" — differ by a shortest nonzero lattice vector, using the norm the octonion algebra already supplies — is basis-independent and requires no free parameter, unlike every prior adjacency construction in this project's history (Addenda BU–BV's own k-nearest-neighbor graph included). Applied to Addendum AW's lattice, this gives exactly bounded degree (=16, matching AW's own reported kissing number exactly, independently reconfirmed this session) and, via the same adjacency structure used as the actual graph, the correct polynomial growth degree (8, matching AW's own theta-series-derived figure via a completely independent computational method). May's criterion's role is correspondingly narrowed: from "the reason sparsity is required" to a weaker, still-true consistency check on whatever dynamics the theory eventually adopts. This is not a uniqueness proof — minimal-vector adjacency has not been shown to be the only possible reading of Postulate 2's algebraic-relation wording — and addresses only the bounded-degree half of this postulate; the "dense coupling ⇒ chaos" half remains inherently dynamical.

STATUS: DERIVED — forced by May's criterion + RMT universality, given the requirement of a stable survivor manifold. Addendum B.6's companion Cheeger-type dichotomy extended at v29 (Addendum EG) to cover a third regime found in the H.17 thread; this postulate's own status unchanged. NEW at v31 (Addendum EO): the bounded-degree component now has a second, independent, dynamics-free justification directly from Postulate 1's own lattice structure — see above.

2.5 Postulate 5 — Lorentzian Signature (3,1)

Signature is DERIVED, not postulated, via three combined arguments: (i) the real Clifford algebra classification (Bott periodicity), (ii) the Majorana reality constraint on the Clifford generators, and (iii) a gradient-flow argument restricting the system to a single timelike direction. This demotes Wick rotation from an assumption to a consequence.

Latest revision (Part III addendum, session of 2026-07-13): a genuine contradiction was found between the Majorana reality condition and spontaneous U(1) breaking of the ground state. Resolved by dropping the Majorana condition and rebasing the argument on Hessian Hermiticity (q = 1) instead. The signature-(3,1) conclusion survives this repair.

STATUS: CORRECTED AT V13 (Addendum BB.3) — CONDITIONAL, not DERIVED as previously stated. The repair described above (Hessian Hermiticity replacing the Majorana condition) is genuine but does not make this postulate independent of Postulate 8′s Route A/Route B fork: H.19′s own formalization of legs (i)-(ii) takes Postulate 8′s Euclidean Clifford algebra $\{\Gamma^a,\Gamma^b\}=2\delta^{ab}$ as a starting premise, and leg (iii) is only well-posed once that generator set already exists. Two Route-A-native alternative formulations of leg (iii) were checked directly and ruled out (BB.4). Under Route A alone, this postulate′s signature-selection mechanism as written has no derivation. See Postulate 8 below and Addendum BB. Confirmed unchanged and correctly reflected through v13–v18 and this consolidation — verified directly against this document′s own text before this consolidation′s session began.

2.6 Postulate 6 — Dimensional Floor (n ≥ 4)

Substantially revised at v11 (Addendum AZ), refined at v13 (Addendum BB.5). The dimensional floor's two legs do not share a status. The vortex/homotopy leg ($n\geq 3$, $\pi_1(S^1)=\mathbb{Z}$ codimension-2 defect classification) rests only on Postulate 2/3's own U(1) order parameter and is Route-B-independent — DERIVED, survives Route A intact. The chirality-parity leg (tightening to $n\geq 4$, built on the Clifford pseudoscalar) is Route-B-dependent by construction,

since chirality is defined in this theory through Postulate 8's Clifford structure. Under Route A alone, the floor is n≥3, not n≥4.

New at v11: Postulate 1's octonion substrate supplies at most 7 genuinely independent anticommuting generators (the Hurwitz–Radon limit, proven via a named 1962 theorem). Combined with n≡2,4 (mod 8) and n≥3, n=4 is the unique value satisfying every constraint derived, conditional on one identification (Postulate 8's generators realized via octonion multiplication) that is well-motivated but not yet independently derived (AZ.13).

New at v26 (Addendum DK.1–DK.9, DK.18). AZ.13's octonion-generator identification (the "not independently derived" clause below) is closed, conditionally: given Route B and G2-equivariance, the octonion-multiplication realization is forced, unique up to gauge and one discrete chirality choice — shown to be exactly octonion conjugation. A flagged resonance with this leg's own "chirality-parity" naming was checked directly and closed by dissolution (DK.18): no standalone parity operator exists in the theory's record, distinct from the pseudoscalar ω, for the found chirality sign to be identified with.

STATUS: PARTIAL. Vortex/homotopy leg (n≥3) DERIVED, Route-B-independent. Chirality-parity leg (tightening to n≥4) CONDITIONAL, Route-B-dependent by construction. Uniqueness (n=4 among all constraints) CONDITIONALLY CLOSED — unique given AZ.11's octonion-generator identification, now closed at v26 conditional on Route B and G2-equivariance (Addendum DK.1–DK.9), and additionally conditional on Route B being adopted at all (BB.5). See Addendum AZ, Addendum BB, and Addendum DK.

2.7 Postulate 7 — The Triadic Substrate Postulate (conditional, SU(3) sector)

Any SU(3) gauge content is conditional on this postulate — algebraic triads of substrate elements rotating through an internal 8-dimensional space. The original unconditional SU(3)-unification claim was formally retired at v1 (Part I, Defect 3); this postulate was written to wall that failure off, not to resolve it.

Major revision, carried from prior consolidations (Addenda V, W, Z, AJ). The discrete spectral precondition (exact 3-fold degeneracy with nonzero Γ) is DERIVED via two independent, verified mechanisms (A4 and Z7×Z3 tensor-square fusion). The continuum structure group is forced to SU(3) among all compact simple Lie groups admitting a complex 3-dimensional representation, reusing the theory's own already-established single-speed bound. Representation type (complex, not real) is DERIVED CONDITIONALLY on Postulate 1's octonion candidate (Addendum AJ), via direct intertwiner-equation solution rather than analogy.

Major revision, carried from v14 (Addenda BC, BD). Triad-sector stabilization (item ii) is now DERIVED: a non-radial, octonion-motivated three-body correction (a triadic determinant term) produces a genuine, dynamically-confirmed, maximally SU(3)-structured stable attractor, with a closed-form stability boundary and exactly 8 Goldstone-type zero modes. This configuration is fully SU(3)-broken, and was shown (via Schur's lemma) structurally incapable of supplying AZ.13's needed unbroken SO(3) subspace — closing that connection negatively but precisely, not leaving it an open gap.

Major revision at v22 (Addendum CP, Gate 1). Postulate 7's own construction (BD.1–BD.4) is a single, spatially-uniform triad — no field theory, nothing for a soliton or any other spatially-extended structure to be built on. Addendum CP established, for the first time, that Postulate 7 admits a genuine site-dependent generalization: BD's on-site potential reconstructed exactly, extended with a minimal nearest-neighbor coupling in the same style Postulate 2/3 already use, and shown — via full Hessian computation across regular, irregular, and disconnected substrates — to support BD's attractor as a stable local vacuum (48 zero modes at zero coupling, collapsing to exactly 8 global zero modes once coupled, zero negative modes in

every case tested). This does not change item (ii)'s DERIVED status; it establishes a genuine prerequisite (a field theory to build on) that had never been checked.

New at v22 (Addendum CP, Gate 2). Whether solitons on this newly-established field theory could supply fermionic statistics via the Finkelstein-Rubinstein/Skyrme mechanism — an attempt to sidestep Postulate 8's Route A/Route B fork rather than force it — was tested directly and found NEGATIVE. A hedgehog-type soliton (Witten's SU(2)⊂SU(3) embedding, the correct target space given BD's fully-broken vacuum manifold ≅ SU(3)) shows no finite-size energy minimum within the physically valid (boundary-respecting) regime: a standard Derrick-instability signature. No non-patching stabilization mechanism was found; the theory's existing quartic term is amplitude-type (stabilizes the radial/density sector, already reflected in fixing A_0) and cannot substitute for the derivative-quartic term genuine Skyrme stabilization requires. This route is closed, cleanly, without narrowing Postulate 7's own DERIVED status for item (ii).

New at v22 (Addendum CR). A separate, dynamical (spontaneous-symmetry-breaking) approach to AZ.13's standing "which direction, then which 3 of 6" question — distinct from the soliton route — found a genuine candidate mechanism. See the AZ.13 entry below and Postulate 8's entry for the full account; the short version: a 2-vector collapse mechanism, built from nothing but Postulate 1's inner-product structure, dynamically selects a preferred direction (breaking G2→SU(3) exactly), and composes without interference with Postulate 7's own already-established triad mechanism (BD) for a second stage. This is a genuine advance on a previously-open question, precisely scoped against Addenda BD.5–BD.6's prior, permanent closure of a stronger version of the same question — not a reopening of that result.

New at v36 (Addenda FI–FN). CS.4's legitimacy gate — whether CR's direction-selecting pair may couple to the substrate's own configuration data — was attacked from two directions. FI closed Addendum EE.6's own recommended precedent audit: no other cross-sector potential coupling exists anywhere in the addenda history, and the one genuine attempt at a cross-postulate connection (Addendum EF) found the theory's existing mechanisms actively resist it — CS.3 is confirmed the theory's first such proposal, with nothing to judge it against. FJ–FN then tested the coupling directly on both real substrate constructions rather than continuing to argue about it in the abstract: CR.3's exact SU(3)-selection mechanism survives fully intact at every tested coupling strength, and BD's own triad sector was directly confirmed decoupled under genuine joint optimization, not merely assumed separable by construction. The previously flat, arbitrary direction CR.3 leaves undetermined is resolved, not damaged, by the coupling — fully on the generic (lattice) substrate, down to a single forced, exactly massless mode on the (structurally disfavored, per Addendum FA.7) orbit substrate, which is itself traced to an exact closed-form spectral theorem identifying it as a fourth manifestation of the same G2-compactness obstruction already documented in Addenda AN, Z.15, and FA.5. This is real evidentiary progress on what the coupling would do if adopted; CS.4's own legitimacy gate — whether it may be adopted at all — remains untouched and open.

New at v37 (Addendum FP). CS.4's legitimacy gate is now DECIDED: the architect's explicit call, per Addendum EE.2's reading (b), is that this postulate's own text ("algebraic triads of substrate elements") was always the operative statement of intent, and BD/CR's free-field construction was a tractable stand-in never reconciled with it. The v-M coupling is licensed. FJ–FN's results above are therefore promoted from conditional characterization to unconditional theory content: direction-selection is fully determined on the generic (lattice) substrate, and determined down to a residual 2-dimensional degeneracy — plus one forced, exactly massless mode — on the orbit substrate. See AZ.13 below for the further consequence this unlocks for the tachyonic-sign question, and Part VI (AW.5/H.17 entry) for how this ties AZ.13's resolution to that entry's own standing lattice-vs-orbit question.

STATUS: PARTIAL, substantially advanced across three sessions. Discrete spectral precondition DERIVED (two independent constructions). All four requirements

selecting SU(3) among continuum structure groups have a derivation route: non-abelian, simple, dimension-8-as-consequence DERIVED unconditionally; representation type DERIVED CONDITIONALLY on Postulate 1's octonion candidate. Triad-sector stabilization term (item ii): DERIVED — a non-radial, octonion/triadic-motivated correction produces a genuine, dynamically-verified stable attractor (Addendum BD), though it is a fully SU(3)-broken configuration rather than one preserving structure useful to Postulate 8's open questions (Addendum BD.5) — now confirmed to rest on a genuine, newly-established site-dependent field theory (Addendum CP, Gate 1) rather than an unverified single-site construction. Soliton-based fermion route on this field theory tested and closed NEGATIVE (Addendum CP, Gate 2). A genuine dynamical candidate mechanism for AZ.13's direction-selection question found and composed with this postulate's own triad mechanism (Addendum CR) — see AZ.13 below for the precise, careful scoping of what this does and does not establish. NEW AT V37: CS.4's legitimacy gate is DECIDED (licensed, Addendum FP) — direction-selection is now DERIVED unconditionally on the generic (lattice) substrate, and PARTIAL (residual 2-dimensional degeneracy, plus a forced massless mode) on the orbit substrate. See Addendum FP and AZ.13 below.

2.8 Postulate 8 — Bipartite Clifford-Algebra-Valued Chiral Coupling

Introduces a bipartite, Clifford-algebra-valued chiral coupling structure to recover Dirac-type dispersion and supply the signature-selection mechanism used in Postulate 5.

Major revision, carried from v11–v13 (Addenda AZ.14, BB.1–BB.2). The originally-claimed necessity proof was found Route-B-conditional (nothing forces Route B over the simpler Route A, which achieves the same ∇^2 dispersion with an ordinary commuting generator, no Clifford structure needed). The fork was subsequently mapped precisely, not resolved: non-degeneracy (not isotropy) is Route B's actual forcing condition, established independent of H.17; Route A remains fully sufficient on its own terms. The fork reduces to whether the theory independently requires first-order (Dirac-type, chiral) dynamics — not derived from Postulates 1–7. This conditionality propagates into ALL of Postulate 5 and HALF of Postulate 6's floor (the chirality-parity leg specifically).

Major revision, carried from v15 (Addendum BG). Six candidate internal forcing arguments for the Route A/Route B fork were tested directly (companion-form doubling of Postulate 3's own dynamics, spin-statistics, gauge anomaly cancellation, bipartite spectral pairing, H.19's positive-definite-Hilbert-space premise, and a proposed connection to Postulate 7's SO(3)-breaking dynamics) — all six returned negative. An independent from-scratch reconstruction of g2 cross-confirmed Addenda Z.5/AJ's su(3)-stabilizer results by a different method, without revising them.

New at v22 (Addendum CO). Tested, for the first time, whether BB.1's non-degeneracy criterion — previously only ever achieved by postulating Clifford generators directly — could instead emerge generically from Route A's own machinery (Hermitian C_{ij}, Addendum M.3's graph-level-bipartite topology) at a band-touching point, the way it does in graphene-type condensed-matter systems. In 2D: YES, robustly — 186/300 generic complex hoppings and 194/300 octonion-derived hoppings produced a genuine band-touching point, 100% non-degenerate given existence in both cases (Cl(0,2) structure, confirmed via Sylvester's law). In 3D: NO, structurally, under strict bipartiteness — a 2-band model supplies only 2 real functions of k regardless of dimension, giving codimension-2 zero loci (points in 2D, lines in 3D); confirmed numerically (nodal line found, rank 2 of 3, one exact zero eigenvalue). Breaking bipartite purity (an intra-sublattice term) restores full rank 3 immediately — creating a direct tension with Addendum M.3's graph-level-bipartite reading.

New at v22 (Addendum CQ). This tension was resolved in favor of preserving Addendum M.3:

a genuine 4-component (Weyl-decomposed) representation — two components per sublattice, exactly what a real 3+1D Dirac fermion already requires physically — achieves full rank-3 non-degeneracy in 3D at near-universal rate (199/200 octonion-derived, 195/200 generic baseline, 100% non-degenerate given existence), with zero violation of bipartite purity anywhere in the construction. CO's 1-component (2-band) test had been structurally under-resourced relative to the actual physical target, not evidence against Route B's internal consistency.

Separately, Addendum CQ checked whether this 2-component structure could be derived canonically from Postulate 1's octonion algebra itself (rather than importing Pauli matrices as an external mathematical tool). Independently reconstructed g2 (cross-confirming Addendum BG.8 by an unrelated derivation, dimension 14, skew-symmetric to $\sim 10^{-15}$), located the octonion basis's 7 genuine associative triples (quaternion subalgebras, matching the classical Fano-plane count), and tested G2's action directly: 0 of 8 random G2 rotations preserved a fixed triple. This is Addendum AZ.13's own obstruction (no G2-invariant selection among equivalent directions), recurring one level up (no G2-invariant selection among equivalent quaternion subalgebras) — not a new, independent gap.

New at v26 (Addendum DK.1–DK.9, DK.18) — applying the g2-classification technique, for the first time, to the specific-realization question BG's search-based attempts never addressed. Given Route B and requiring the generators to be G2-equivariant (the natural non-patching request), the spinor module is forced to be $S \cong V \oplus \mathbb{R}$ (the only 8-dimensional real G2-module admitting any Clifford structure at all), and the octonion-multiplication realization of the generators on that module is forced, unique up to a trivial rescaling of the two summands and one unavoidable discrete chirality choice — shown, exactly, to be octonion conjugation (an orientation reversal of $\mathrm{Im}(\mathbb{O})$). This closes AZ.12/AZ.13's "nothing forces the octonion-derived realization over an independently-postulated internal tensor factor" gap, conditional on Route B and G2-equivariance. It does not touch BG's own six-for-six negative result on whether Route B is required at all, and does not resolve AZ.13's separate "which 4 of 7" question (see Postulate 7, Addendum DK.10–DK.12). A flagged resonance between the found chirality sign and Postulate 6's "chirality-parity" naming was checked directly against Addenda AZ.2/BB.5 and closed by dissolution (DK.18): no standalone parity operator exists in the record to identify it with.

STATUS: UPDATED AT V13 (Addendum BB.1–BB.2) — necessity of Clifford structure remains CONDITIONAL, precisely characterized rather than open. Does not force Route B over Route A. Propagates into ALL of Postulate 5 and HALF of Postulate 6. Specific multiplication rule remains POSTULATED as a matter of record, though UPDATED AT V26 (Addendum DK): its AZ.12/AZ.13 realization gap specifically is now closed — conditional on Route B and G2-equivariance — rather than merely well-motivated. Six internal forcing candidates tested at v15, all negative (Addendum BG) — the Route A/Route B necessity question remains exactly as open as before, now substantially more exhausted (search and forcing-argument methods, plus a classification-technique test of the realization sub-question at v26). Unaffected by v17–v19's own work, which addressed Postulate 1/H.17 and Postulate 3 specifically. UPDATED AT V22: Route B's internal machinery, once/if adopted, is now shown to work robustly and consistently in the physically relevant 3D case using only structure already established elsewhere in the project (Addenda CO, CQ), with no violation of Addendum M.3. The soliton-based sidestep of the fork was tested and closed negatively (Addendum CP). None of this bears on BG's six-for-six negative result on whether Route B is actually required — that question is untouched. CONDITIONAL, 75% / POSTULATED, 25% by the architect's own working estimate — unchanged overall.

2.9 AZ.13 — Updated Entry (Referenced from Postulates 7 and 8)

Original finding (v11, unchanged and not reopened): G2 acts transitively on S^6 — proven, not merely unattempted — so no G2-invariant *selection rule* can prefer one imaginary octonion direction over another. Addenda BC–BD (v14) then classified the deeper "3 of 6" reduction and found the only subgroup pattern admitting a real 3-dimensional invariant subspace is an unbroken residual SO(3) symmetry — then proved, via Schur's lemma applied to SO(3)'s vector representation, that no nonzero triad-sector configuration of any kind can achieve this (SO(3) acts transitively on every nonzero-radius sphere in its own representation, so any nonzero configuration necessarily breaks it further). This is a permanent, structural closure of the strong (symmetry-invariant) version of the question, reconfirmed via primary-source retrieval at v15 (Addendum BG.11–BG.12) rather than re-derived.

New at v22 (Addendum CR) — precisely scoped, not a reopening. CR tested a different question: not whether an invariant *selection rule* or an unbroken *residual symmetry* can prefer a direction (both already proven impossible), but whether a G2-invariant *potential* can have a symmetry-broken minimum that dynamically selects one — exactly the distinction Addendum BD already exploited once, one level down, for SU(3) itself. Two findings:

- A G2-level triad (3 vectors, using the associative 3-form φ as the direct analogue of BD's det term) does spontaneously break G2 — confirmed via two independent methods (Hessian zero-mode count and direct g2-generator rank), both giving stabilizer dimension 3 exactly — but to a *different* pattern (a preferred 3-plane with a 4-dimensional complement) than AZ.13's original 1-then-3 target, and not a re-derivation of BD's own SU(3)-level result.
- A 2-vector "collapse" mechanism — only a pairwise term and radial terms are available with 2 vectors, no cubic invariant — under a standard tachyonic-mass potential, dynamically breaks G2 to exactly SU(3) (stabilizer dimension 8, confirmed two independent ways; confirmed complex-linear on the orthogonal complement via the octonion-induced structure J=L_n̂, genuinely "SU(3) on $\mathbb{C}^3$" in the standard sense, not merely a dimension-matched subgroup). This is the first genuine dynamical mechanism found for AZ.13's *first* stage — direction selection — which every prior treatment (BC, BD, BG, and Postulate 7 itself) simply took as given.

A combined 5-vector system (the direction-selecting pair plus BD's own triad, with an explicit orthogonality-enforcing coupling) was checked and found to compose without interference — the combined minimum energy is exactly the sum of the two isolated minima, the triad's overlap with the selected direction is exactly zero, and the full system's final stabilizer is exactly 0 (complete breaking, G2→SU(3)→{1}). The "3 of 6" identification this produces is via vacuum selection (BD's triad lands on 3 specific directions because that is where the dynamics puts it) rather than unbroken symmetry (which BD.5–BD.6 already proved impossible) — a different, weaker, but genuine and newly-available sense of "preferred," not a contradiction of the prior closure.

Honest limitation, stated in the addendum and repeated here: the mechanism requires an unforced input — a tachyonic mass sign (γ<0) in the direction-selecting potential, without which the trivial (unbroken) configuration is the true minimum and no selection occurs at all. Nothing in Postulates 1–3 currently forces this sign. This is a well-precedented kind of assumption in physics (the same status the Standard Model's own Higgs potential sign has) but it is an assumption, not a derivation, and should not be reported as more than that.

New at v26 (Addendum DK.10–DK.12). The "3 of 6" leftover-direction question is sharpened from an empirical search finding to a proof. Building su(3)=Stab_g2(n) directly and computing its commutant on the leftover 6-dimensional space (exactly 2-dimensional, ≅ℂ, admitting only trivial idempotents), DK shows no nontrivial invariant subspace of the leftover space can exist — of any dimension, not only a real 3-dimensional one. This does not change this

entry's PARTIAL status or discharge CR's own unforced tachyonic-sign contingency, but it permanently forecloses any future purely algebraic/symmetry-based selection attempt: CR's dynamical mechanism is now provably the only remaining route, not merely the one that happened to be tried. *New at v26 (Addendum DK.14–DK.17).* Addendum CS.3's tachyonic-sign coupling proposal (a coupling of CR's direction-selecting vectors to M, the substrate's second-moment tensor) is reopened after correcting a scope slip in CS.5 (the coupling needs no site index on v, contrary to CS.5's stated reasoning) and shown to occupy the natural (equal-weight) point in an exactly 2-dimensional space of G2-invariant v–M couplings — not a cherry-picked functional form. CS.4's deeper legitimacy gate (whether substrate data may enter CR's potential at all) is unaffected and remains the live blocker; AZ.13's $\gamma<0$ contingency stands.

New at v36 (Addenda FI–FN). EE.6's own recommended precedent audit was carried out in full (Addendum FI): no prior cross-sector potential coupling exists anywhere in the addenda history, and the theory's one genuine attempt at a cross-postulate connection (Addendum EF) found the opposite of precedent. CS.3's coupling was then tested directly, substrate-coupling CR.3's pair to both real substrate constructions: the mechanism survives fully intact at every tested strength, BD's triad sector was directly confirmed decoupled, and the coupling resolves rather than damages the previously-flat direction — fully on the lattice substrate, down to one forced, exactly massless mode (traced to an exact closed-form spectral theorem and G2's own compactness) on the orbit substrate. CS.4's legitimacy gate itself remains untouched and open; this characterizes the coupling's consequences, not its licensing.

New at v37 (Addendum FP). CS.4's legitimacy gate is DECIDED — the architect's explicit call, licensing the v-M coupling (Addendum EE.2's reading (b)). This is the larger unlock, not merely for CS.3 itself: prior to this decision, CR.3's mechanism still required an unforced bare tachyonic mass $\gamma<0$, with the coupling only riding on top of that separate assumption. With the coupling licensed, Addendum EZ.2's reduction $\gamma\_\mathrm{eff}(n)=\gamma-\kappa\cdot n^{T}Mn$ becomes available in full: $\gamma\_\mathrm{eff}<0$ is achievable from the coupling term alone, at $\gamma\geq 0$, a genuine substrate-anisotropy-driven effective mass rather than a bare postulated sign. Whether this discharge is generic or fine-tuned is substrate-dependent, per Addenda EZ.4–EZ.8's own figures: close to forced on the orbit substrate (anisotropy 1.0 exactly, 8/8 seeds), still closer to fine-tuned on the lattice substrate (near-isotropic, 5/5 seeds). AZ.13's resolution is therefore now tied directly to Part VI's own standing AW.5/H.17 lattice-vs-orbit question, more tightly than Addendum EZ.9 first flagged.

New at v42 (Addendum GL). The discharge condition was tested on a third substrate type for the first time — the Heisenberg-group k-NN construction (Addenda GG–GJ) — where M has rank exactly 3/7 with anisotropy exactly 1.0 (though this is identified as a near-tautological consequence of the construction's single-3-plane embedding, not an independent geometric result). Composed successfully with genuine Postulate-3 dynamics, the cleanest direction-determination outcome of the three substrate types tested to date — joining the lattice and orbit substrates Addenda FJ–FN already tested. Does not resolve the substrate-dependence question itself; adds a third data point to it.

New at v44 (Addendum GR), not previously entered here — backfilled at v45. The discharge condition was tested on a fourth substrate for the first time: $H_3(\mathbb{Z})$, the rank-3 Heisenberg group, dimension exactly 7, embedding natively into all seven octonion imaginary directions with no coordinate compression — avoiding the shared root cause (GJ/GK/GL's single-3-plane embedding) behind the near-tautological readings of prior positive results. M is genuinely full rank (7/7), with a non-degenerate anisotropy figure (0.386, at the radius-6 substrate GR used), and composes cleanly with the actual 35-parameter joint dynamics — the first AZ.13 discharge result on record not flagged as tautological. Direction-selection under the full dynamics is unconditional on any nonzero coupling tested (no threshold gating). Adds a fourth data point to the substrate-dependence question; does not resolve it.

STATUS: PARTIAL. Strong (symmetry-invariant) version of the 3-of-6 question: CLOSED, NEGATIVE, permanent (Addenda BD.5–BD.6, unchanged). Direction-selection (first stage): PARTIAL — genuine dynamical candidate found (Addendum CR), contingent on an unforced parameter choice. Weaker (vacuum-selection) version of the 3-of-6 question, composed with the first stage: PARTIAL — genuine candidate mechanism found (Addendum CR), same contingency. NEW AT V26: the strong version's negative closure (BD.5–BD.6) is now reinforced by an independent, general proof for the leftover-space case (Addendum DK.10–DK.12); CS.3's coupling route is reopened and its functional form shown tightly constrained (Addendum DK.14–DK.17), though CS.4's legitimacy gate remains open. NEW AT V36: the precedent-audit question is closed (Addendum FI, no precedent found) and the coupling's consequences tested directly on both real substrates (Addenda FJ–FN, mechanism survives, previously-flat direction resolved), though CS.4's own legitimacy gate is unaffected and remains open. NEW AT V37: CS.4's legitimacy gate is DECIDED (licensed, Addendum FP), converting the $\gamma<0$ contingency from simply unforced to substrate-dependent — close to forced on the orbit substrate, still fine-tuned on the lattice substrate — and tying this entry's resolution directly to Part VI's AW.5/H.17 question. See Addenda DK, FI-FN, FP. NEW AT V44 (Addendum GU): EZ.4–EZ.5's "orbit substrate" figures cited above are found to describe a different, misnamed construction — a literal orbit under $\exp(\tau D)$, not by1 as used consistently throughout Addenda BY–CB, CD, DO–DP, and GK. Recomputing M on the verified by1 substrate (8 seeds × 3 weighting schemes) gives full rank and anisotropy 0.10–0.16, not the claimed rank-5/anisotropy-1.0 forced result. BU/BV's lattice figures were independently re-verified for parity and hold up (anisotropy 0.085–0.150). AZ.13's discharge mechanism is therefore weak on every substrate the project actually and consistently uses, not substrate-dependent as previously stated. FP.0's licensing decision itself rests on independent textual grounds (Postulate 7's own wording, EE.2's reading (b)) and is unaffected by this correction. See Addendum GU in full.

New at v45 (Addendum GW.8–GW.9), correcting in place rather than reversing the entry immediately above. Reconciling an independent reproduction of GR's own $H_3(\mathbb{Z})$ figure (0.386) surfaces a scope gap not previously flagged anywhere in this record: that figure uses GR's own raw-covariance anisotropy convention (distinct from EZ.5's unit-sphere-normalized one, used elsewhere in this entry) at a single radius, and is not scale-invariant under that convention — tracked across radii 3 through 7 it climbs steadily (0.125→0.090→0.256→0.386→0.487) with no sign of convergence, consistent with $H_3(\mathbb{Z})$'s c-coordinate accumulating quadratically in its own group law while a, b accumulate linearly. GR's qualitative framing (non-degenerate, not tautological, in contrast to GL's Heisenberg result) is not overturned — GR.10's own finding that direction-selection is unconditional on any nonzero coupling makes it plausible the qualitative conclusion does not depend on which radius's anisotropy value is used — but the specific 0.386 figure is not itself a well-defined, radius-independent property of the substrate the way this entry's other anisotropy figures are. Not independently verified at low-anisotropy radii; flagged as open rather than resolved.

New at v29 (Addendum EE.1–EE.5). CS.4's own internal pointer ("resolved in CS.5") is checked directly and found inaccurate — CS.5 only established spatial uniformity, a fact about the coupling's form, not an answer to CS.4's actual question of what v1, v2 physically represent; CS.4's gate remains open for a now precisely identified reason (EE.1). A previously unflagged drift is found between Postulate 7's own text ("algebraic triads of substrate elements") and BD/CR's realized construction, which builds v1, v2 as free internal fields rather than quantities built from the substrate's own realized configuration — two readings left open, not adjudicated (EE.2). The project's own postulate-patching criterion (Addendum CT.1–CT.3) is extracted and applied mechanically to CS.3's coupling: it passes on a literal reading (EE.3–EE.4), but the criterion was calibrated only on same-sector extensions (the triangle term reaching deeper into Postulate 2/3's own already-integrated data), and CS.3 is the first cross-

sector case (Postulate 7's v to Postulate 1/2's M) it has ever been asked to judge — a gap in the standard itself, not merely in this application of it (EE.5). CS.4 remains open; AZ.13's $\gamma<0$ contingency is unchanged.

New at v29 (Addendum EF.1–EF.8). Tests directly whether CR/BD's already-verified dynamical selection mechanism could also supply Postulate 8's separate "which 4 of 7" generator question (Addendum DK.10–12's no-go forecloses only symmetry-invariant selection, not vacuum selection). A general algebraic fact is verified from scratch: DK.7–8's chirality-operator construction ($\omega^2=I$, full anticommutation) works identically for any orthonormal 4-frame in Im(𝕆), not only the canonical placeholder DK.8 checked (EF.1). An initial counting objection (CR/BD's frame plus Postulate 8's own four would require 8 of only 7 available directions) is raised and then corrected in place: DK.8's own text already treats a clean 4+3 partition as the natural reading, under which $\hat{n}$ is one of Postulate 8's four and no obstruction exists (EF.2–EF.4). What remains is narrower than "which 4 of 7": if $\hat{n}$ is one of Postulate 8's four, the other three must lie in a specific 3-dimensional subspace L, the orthogonal complement of $\hat{n}\oplus$BD's triad (EF.5). g2 is independently re-derived and CR.6's "complete breaking, G2→SU(3)→{1}" result is reproduced generically, not only for whatever triad the original session happened to find (EF.6). This shows L is a genuine flat direction — no residual G2 symmetry relates different bases of L (so it is not gauge), but CR+BD's own potential never references L at all (so nothing dynamical constrains it either); "which 4 of 7" is sharpened into "which orthonormal basis of L," a smaller, precisely bounded open question, not a closure (EF.7–EF.8).

New at v34 (Addendum FB.1–FB.2). $\hat{n}$ — computed via eigendecomposition of D=[L_a,L_b]+[L_a,R_b]+[R_a,
throughout Addenda Z–EZ — has a closed-form expression, not previously stated anywhere in the record: $\hat{n} = \pm\mathrm{Im}(a\cdot b)/|\mathrm{Im}(a\cdot b)|$, verified to machine precision across ten independent trials (alignment 1.000000), orthogonal to both a and b (maximal distance to span(a,b)). No eigendecomposition is required to compute $\hat{n}$ going forward. Since Addendum AW.1/BU.1 already established $\mathbb{Z}^7$'s closure under octonion multiplication, this identity further shows that lattice-valued generator pairs (a,b) give exactly lattice-valued (pre-normalization) fixed axes — a previously unnoted exact bridge between the orbit- and lattice-picture readings of this postulate's own generator-selection construction (see Part VI, AW.5/H.17 entry).

NEW at v43 (Addendum GQ.1–GQ.15) — supersedes EF.7, above. Directly resolves EF.7's flat direction (quoted above: "no residual G2 symmetry relates different bases of L… nothing dynamical constrains it either") rather than replacing that finding silently — the original text is left untouched above and explicitly superseded here, per the project's own no-silent-corrections standard. L's basis is exactly and canonically supplied by the octonion associator. For BD's dynamically-selected triad {t1,t2,t3} (an associative triple — a genuine quaternion subalgebra Q, since it maximizes φ — confirmed exactly rather than assumed) and $\hat{n}$ (orthogonal to Q, by CR.4's zero-interference composition), the three vectors [t1,t2,$\hat{n}$], [t1,t3,$\hat{n}$], [t2,t3,$\hat{n}$] form an exact orthogonal basis of L: Gram matrix = $4\cdot I_3$, not approximately but exactly, proven directly via Cayley-Dickson doubling of O = Q ⊕ Q$\hat{n}$ rather than left as a numerically-verified pattern — $(t1\,t2)\cdot\hat{n} = t3\hat{n}$ while $t1\cdot(t2\hat{n}) = -t3\hat{n}$, giving $[t1,t2,\hat{n}] = 2t3\hat{n}$ exactly, cyclically. The identity is fully general (depends only on the triad's associativity and $\hat{n}$'s orthogonality to it, not on $\hat{n}$'s specific construction) and was replicated across four independent full dynamical realizations, each re-deriving both $\hat{n}$ and the triad from the actual CR.3/BD potentials from scratch — not a random stand-in satisfying only the group-theoretic constraint, which an earlier pass within the same addendum used first and found gave no information (φ vanishes identically on an associative triple's own orthogonal complement, since the triad's cross products stay entirely within its own span; a general, previously unstated identity, $F4\cdot F4^{T}=c\cdot I_3$ for any orthonormal triad, with the physically-real associative triad sitting at its exact zero-c degenerate limit, is what forced the pivot from φ to the associator). Two further corrections logged in place within the same addendum: Addendum GA.5's proposed link between this question and Addendum BI.7's seed-point arbitrariness does not hold, for two independent reasons — this resolution never references the seed point e_0 at all, and BI.7

itself was already closed by Addendum BJ (traced to an orbit-parametrization artifact, not a genuine free constant) well before GA.5 was written. Recommended: Postulate 6/8's "which 4 of 7" entry upgraded from CONJECTURE/OPEN toward DERIVED at this consolidation. Note, logged rather than silently reconciled: this document's own v40 historical consolidation summary (elsewhere in this Part) records Addenda GA–GD as already having closed EF.7 via a separate argument about the seed point e_0 of an unrelated D-eigenvector construction; that claim rested on an unverified premise (GA.5) never independently re-checked by GB, GC, or GD, and does not survive contact with this addendum's direct construction, which references neither e_0 nor that construction at any point — see Addendum GQ.16 for the full account.

New at v29 (Addendum EG, cross-referenced from Postulate 4 above). A separate thread tested Postulate 4's companion Cheeger-type dichotomy (Addendum B.6) against the v28 H.17 thread's contextual-gluing construction, finding a genuine third regime; a candidate explanation connecting it to Addendum EB.2's own open puzzle was tested in a controlled model and found to be a transient crossover rather than a stable effect; two attempts to extend the real construction to much larger N (via source recovered through conversation_search) produced one honestly-caveated large-N result and one correctly-halted failed calibration. Does not bear on AZ.13 directly; logged here only for cross-reference completeness. See Postulate 4 and Addendum EG for the full account.

3 Part III — Session Addenda Moved at v15, unchanged at v16. Full text of all session addenda (Addendum A through the current letter, chronological, append-only) lives in the companion file SFT_Addenda_Archive.docx. This master document's Parts II, V, and VI continue to cite addendum letters directly (e.g. "see BM.4", "Addendum BT"); readers following a citation into full derivation detail should open the archive file at the cited letter.

DOCUMENTATION BACKFILL AT V46: the Addenda AA–AH thread (session date 2026-07-17) is restored to the record. Investigating the AA–AH gap flagged by this project's own H.17 investigation audit (2026-08-26) found that these eight addenda — a single continuous investigation into the octonion/G2 orbit construction's locality-emergence mechanism, running from Addendum Z's unfinished priority queue through a full closed-form resolution — were never folded into any master-document consolidation and, contrary to the general rule stated above that addenda before Addendum BG live in Master_Document_v14.docx, are not actually present there either. This is significant beyond a bookkeeping gap: this document and the companion archive already cite AA–AH's results by letter in three places (Addendum FB.1–FB.2's $\hat{n}$ construction; Addendum FM's orbit-substrate M-spectrum theorem, which explicitly invokes "Addendum AF's own… Dirichlet(1/2,2,1) distribution"; and Addendum GE.2's calibration convention, which cites "Addendum AA.0" directly) — so a reader following any of those citations before this backfill would have found no such addendum anywhere in the record. AA–AH's own arc, briefly: AA reconstructs Addendum Z's pipeline from its documented method (the exact build seed and site-spacing were never recorded) and confirms Z.41–Z.45's amplitude-weighted-decay finding generalizes across all five of Z.28's original seeds; AB scans 60 independent instances and finds an 88.3%/11.7% split governing whether the locality-emergence mechanism activates; AC discovers this split is a circulant-matrix artifact of one specific discretization convention and retracts AB's generality claim; AD system-

atically tests the alternative, non-resonant convention and finds multi-block clustering is the typical outcome there, apparently confirming AC's retraction; AE then resolves the entire dispute by reformulating the problem in continuous (dt-independent) terms, showing AC's "arbitrary convention" framing was itself mistaken and AB's original split was the correct, exact answer all along — AD's multi-block results are reclassified as aliasing artifacts of the wrong discretization, not a competing physical regime; AF derives the model's three governing harmonics in closed form and identifies the exact probability distribution governing which dominates, Dirichlet(1/2, 2, 1); AG sanity-checks the resulting double-arc prediction against direct simulation (finding it holds only above a spectral-gap threshold) and derives the underlying eigenvalue pattern via trace invariants; AH completes the derivation with a full symbolic (computer-algebra) proof of two of the three needed trace identities, with the third verified via exact rational arithmetic at every tested instance but not yet proven symbolically for all inputs — a stated computational, not conceptual, gap. Full text of all eight addenda is now in the companion archive, inserted immediately before Addendum BG per its own chronological position; see SFT_Addenda_Archive.docx for the complete record, including this backfill's own provenance note repeated at the head of each of the eight entries. No formal status-tag changes anywhere in the theory as a result of this backfill — every finding AA–AH contains was already, in substance, part of the theory's working record via the **DOCUMENTATION BACKFILL AT V47: the Part III §G snapshot and the Addenda AO–AX thread (session dates 2026-07-18/19, nine addenda) are restored to the record.** This is the second half of the same recovery effort that closed the AA–AH gap at v46, continuing directly from the same H.17 investigation audit (2026-08-26). Two distinct gaps are closed here.

G was never a separate lettered addendum at all — it is Part III §G of an early Master Document snapshot, predating the v1–45 versioning scheme and even Addendum H itself (Addendum H's own header confirms this: "Continues Part III of the Consolidated Master Document as §H"). It addresses fermion doubling: the Wilson-term route is scoped and retired on fine-tuning grounds (severe, $\sim 10^{20}$, with no internal suppression mechanism found), overlap/Ginsparg–Wilson fermions are adopted in its place, and the resulting locality cost and Cheeger gap-versus-growth tradeoff are worked through in full — ending on one genuinely open item (G.6, the overlap-admissibility spectral-gap estimate) later closed by the unlettered k0-damping/adjacency-lemma session and Addendum H's own §H.1–H.15. G is cited by name downstream (Addenda EX.9/EY.8, "the original Wilson-term-vs-overlap-fermion decision") and was confirmed genuinely absent from SFT_Master_Document_v14.docx despite the general rule that material before Addendum BG lives there — the same kind of exception AA–AH required at v46.

AO through AX (AR excepted, since it is confirmed present in Master_Document_v14) form one continuous investigation, directly following on from Addendum AN's closure of the static/symmetric branch of H.17's fork: does Postulate 3's own existing selection dynamics, or some substrate-derived construction, supply the missing piece K–N left open? The arc, briefly: **AO** tests Postulate 3's own dynamics on a large discrete-subgroup orbit substrate and finds real, positive, four-for-four evidence of dimension reduction under selection — but flags its own decisive limitation: the substrate has a hard compactness ceiling (AN's own theorem), so the finite-size-scaling test needed to confirm genuine growth-class escape cannot be run on it at all. **AP** finds two scalable substrates to run that blocked test: a continuous G2-orbit (clean $d \approx 1.0$ growth, but possibly already presupposing lattice structure) and a genuinely random bounded-degree graph (a real but noisy signal that selection changes growth-saturation behavior). **AQ** averages the noisy signal over 8 realizations, resolving the methodological ambiguity but not the physics question — both log and power-law fits remain comparably good. **AS** independently reconstructs the entire pipeline from scratch (seeds preserved this time) and extends the test 8× further in N (to 102,400): the survivor/raw diameter ratio is flat across a full 512× range, resolving AQ's ambiguity pessimistically — selection changes the log-diameter law's multiplicative constant, not its growth class. **AT** corrects a stale line in AS's own priority queue (caught by checking the master document

directly) and fully diagnoses AS's anomalous non-fixed-point branch as ordinary RK4 numerical stiffness, confirming AS's headline finding is dt-robust. **AU** constructs Postulate 2's antisymmetric coupling Γ directly from the octonion algebra for the first time and proves — not just observes — that the resulting exact chiral rotating-wave solution always has at least one unstable direction, extending Addendum O.5's saddle-point mechanism to a genuinely new solution class. **AV** finds the first candidate discretization anywhere in this thread to achieve confirmed polynomial growth (the integral octonions / E8 lattice, imported from Coxeter 1946) but judges H.17's circularity to resurface in algebraic rather than geometric form. **AW** removes the import: the project's own naive $\mathbb{Z}^8$ lattice is trivially closed under its own structure constants, independently re-verifies polynomial growth by a second method, and surfaces an unplanned convergence with the theory's already-established Z7×Z3 symmetry (Addendum Z) — while sharpening AV's tension into the **lattice-vs-orbit tension (AW.5)**, one of the most heavily-cited results in the whole H.17 record downstream. **AX** tests whether the full 6-dimensional orbit space (not one curve through it) can achieve polynomial growth on its own: it can (AX.4, connectivity-verified across a 2000× range in N) — but connecting back to the real octonion coupling kernel shows geometric locality mechanically forces uniformly- positive coupling, so growth and competitive selection cannot coexist on this construction (**AX.10**, the other most heavily-cited result in this stretch, and the direct predecessor of AW.5/AX.10's eventual resolution at v42/GL and v44/GR — see Part VI's H.17 register). AU's own working notes continue directly into a further general result (labeled AY in the original consolidation, sharing AU's session date) proving no coupling matrix whatsoever can stabilize the bare cubic master equation; that derivation was already on file verbatim in Master_Document_v14 under a combined "Addenda AU, AY" header prior to this recovery effort and is not duplicated here — only AU's own primary text, which was itself missing, is restored.

Full text of G and AO–AX is now in the companion archive: G inserted at the very start of the body, before Addendum AA, since it predates the entire addendum-letter convention; AO–AX inserted immediately after Addendum AH and before Addendum BG, per their own chronological position. Ten new Contents entries added (G, AO–AX). Separately, this consolidation also closes the companion archive's own long-standing FV–GU Contents-index gap (flagged and left open, unaddressed, at every consolidation from v34 through v46) — all 26 missing entries are added in their correct position; see SFT_Addenda_Archive.docx for the complete record. No formal status-tag changes anywhere in the theory as a result of either backfill — G's findings were already superseded in substance by Addendum H's own later, more rigorous closure of the same question, and every AO–AX finding was already, in substance, part of the theory's working record via Part V §2's and Part VI's own narrative summaries of this era; this closes two citation-traceability and navigation gaps, not a physics gap.

downstream addenda that cited it; this closes a citation-traceability gap, not a physics gap.

4 Part IV — Macroscopic Emergence, Effective Action, and Falsifiable Predictions

4.1 1. Random Matrix Theory Sparsity Bound

May's stability criterion, combined with random matrix theory, forces sparsity and bounded degree on the coupling structure (Postulate 4): dense, unstructured coupling produces chaotic dynamics rather than a clean survivor manifold.

4.2 2. Lorentz-Invariance Violation — corrected to quadratic

The original linear-dispersion LIV claim does not survive: the symmetric coupling structure (Postulate 2's real part) forces quadratic, not linear, dispersion. This was checked against real gamma-ray-burst timing data (GRB 090510, GRB 221009A) and is compatible with current observational bounds at the quadratic order, whereas the original linear claim would already be excluded.

STATUS: DERIVED (quadratic form) — original linear-LIV claim dissolved and replaced

4.3 3. Speed of Light and the Continuum Limit

c is redefined as the maximum propagation velocity of informational phase shifts across substrate linkages; when discrete linkages are coarse-grained into the continuous spatial Laplacian, the coupling terms directly set the scaling coefficient c^2. This derivation now explicitly assumes polynomial substrate volume growth (Part III, B.6) — see Postulate 1 and H.17 above for the current, substantially narrowed but still open state of that assumption. **Resolved (negatively) at v33:** the dependency flagged at v32 was investigated directly rather than left as a cross-reference. Addendum EX found the picture is more layered than the original flag stated — B.5's Lieb–Robinson route likely was not even the right tool (a more direct theorem exists requiring only bounded degree and a spectral gap, no volume-growth condition) — but H.14's own multiscale induction carries an independent, previously unflagged polynomial-growth dependency of its own, so the more direct route's required input remains unestablished on the honest substrate either way. Addendum EY then measured SFT's actual disorder strength directly, using quantities already established elsewhere in the theory (Postulate 2's real coupling, Addendum H.13's own disorder definition), against a connectivity-matched Anderson/Bethe-lattice threshold: the threshold is not cleared at any reliably-measured substrate configuration, sparse or dense, with the shortfall worsening at higher connectivity. This is an evidential negative result, not a forcing theorem — see Addenda EX, EY and Part VI for full honest scope. **Overlap admissibility (H_W spectral gap) is downgraded to CONJECTURE/OPEN and restored to the Part VI register at v33.**

4.4 4. Wave-Attenuation and Radiative-Naturalness Bound

Two branches were tested for damped vs. lossless wave propagation. Branch A (damped propagation) catastrophically fails observational attenuation bounds, requiring suppression around 10^{-31} — excluded. Branch B (exact decoupling) requires exact decoupling that would contradict wave propagation itself unless resolved by a radiative-naturalness argument, giving a concrete bound $\Gamma_0 \lesssim 6\times10^{12}\ s^{-1}$.

STATUS: DERIVED — Branch A excluded observationally; Branch B survives with a concrete falsifiable bound on Γ_0

4.5 5. Effective Action (Grand-Unification Claim Retired)

$$S = \int\sqrt{-g}\, d^4x \left[(R-2\Lambda)/16\pi G - \tfrac{1}{4} \Sigma_k F^k_{\mu\nu} F^{\{\mu\nu\ k\}} + \Psi(i\gamma^\mu D_\mu - m)\Psi \right]$$

Gravity and the cosmological constant $(R-2\Lambda)/16\pi G$: recovered, with Λ driven by the baseline decay rate γ_i — the residual, non-divergent vacuum energy addresses the dark-energy fine-tuning problem without new input.

Gauge fields $F^k_{\mu\nu}$: retained in general, but the SU(3) sector is explicitly conditional on Postulate 7 (Triadic Substrate Postulate). Fermion sector $\Psi(i\gamma^\mu D_\mu - m)\Psi$: built on overlap/Ginsparg–Wilson fermions, avoiding the Nielsen–Ninomiya doubling obstruction

while preserving the protected vector Goldstone. *The overlap kernel's spectral-gap admissibility (whether H_W genuinely has the gap this construction needs) is CONJECTURE/OPEN as of v33 — see Part VI and Addenda EX–EY; the doubling-avoidance and Goldstone-protection results above do not depend on admissibility and are unaffected.*

4.6 6. Falsifiable Predictions (current state)

- Quadratic LIV dispersion in high-energy photon timing, constrained against GRB 090510 and GRB 221009A data.
- A concrete wave-attenuation bound $\Gamma_0 \lesssim 6\times10^{12}$ s^{-1} (Branch B radiative-naturalness result).
- Quantum Decoherence Floor: a non-thermal 'dissipative floor' from substrate selection pressure — still an open-ended qualitative prediction.
- CMB polarization anomalies from primordial-transition network fractures — still qualitative.

5 Part V — External Critique Tracking

An external analysis of five ChatGPT-generated documents (SFT1–SFT5) found that only SFT1 engaged with the actual theory. SFT1's two substantive critique points are tracked here permanently. Further external review episodes are tracked at §2b–2c below.

5.1 1. Postulate-patching pattern

Critique: Postulate 8 (Clifford-valued coupling) and the Dirac-fermion derivation risk being a patch introduced specifically to reproduce known fermion phenomenology, rather than a structurally forced consequence of the substrate postulates.

Independent-motivation verdict: mixed. The bipartite topology itself is independently motivated. The Clifford multiplication rule specifically remains postulated to match fermion phenomenology.

STATUS: CORRECTED AT V11 — the necessity proof is Route-B-conditional; the one non-circular rescue depends on H.17. The honest status is PARTIALLY ANSWERED, CONDITIONAL. Specific multiplication rule remains POSTULATED, with a candidate (not yet derived) closing route proposed at AZ.11–AZ.13. Unaffected by this consolidation's work.

5.2 2. Adjacency-before-geometry circularity

Critique: deriving spacetime adjacency/geometry from the substrate presupposes an adjacency structure (via W_ij / C_ij) that is itself defined without reference to geometry — a structural circularity, not merely a fine-tuning gap. Explicitly cross-referenced to causal set theory's unresolved primitive-order-relation problem, which faces the same class of difficulty.

Formalized as a permanent open problem, distinguishing precisely what SFT actually derives (the survival pattern over a given W) from what it presupposes (W's adjacency structure itself).

Addendum (Part III, H.9 and H.17). H.17 sharpens this into a concrete either/or, grounded in classical random-graph diameter theory: staying pregeometric with only bounded degree

(Postulate 4) forces the generic realization to be a random sparse ensemble, which has logarithmic diameter almost surely, not polynomial growth — while obtaining polynomial growth requires a lattice/Cayley-like realization that presupposes geometric structure by hand.

Addendum (Part III, K–N, AN–AQ, AS–AY). A candidate dynamical-rewiring route was tested and substantially closed (K–N). The octonion-substrate discrete-subgroup route was closed (AN). A direct test of Postulate 3's own selection dynamics gave a real, statistically robust, noise-eliminated quantitative effect (AO–AQ) whose qualitative resolution (AS) leans pessimistic — the survivor/raw diameter ratio is flat across a 512× range in N, favoring a larger logarithmic constant over genuine growth-class escape. Separately, two new, genuinely non-subgroup discretization candidates were found and verified (AV–AW: discrete subrings/lattices of the octonion algebra), revealing a new tension (AW.5, lattice picture vs. the bounded-orbit picture used throughout the theory's dynamical work) and a new trade-off (AX.10, the orbit side can achieve genuine growth but only by forcing uniformly-positive coupling, precluding selection).

Addendum (Part III, BT–BW). Both of AW.5's and AX.10's demonstrated obstructions now have at least one working counterexample in this project's own record, using the triangle-term mechanism (Addenda BE–BQ) that did not exist when either obstruction was first identified. Addendum BT found the triangle term, unlike the pairwise coupling AX tested, produces genuine spatially coherent partial selection directly on AX's own orbit-based substrate, at zero pairwise coupling — with a mechanical explanation tracing to Addendum BM.1's exact vanishing theorem for the associator near coincident points. Addenda BU–BV built a genuine, independently-verified octonion lattice and found growth and triangle-term selection coexist there too, first via normalized point-directions then confirmed via the full unnormalized octonion product directly — closing an honestly-flagged gap in the first construction. Addendum BW extended both findings toward Addendum AX's own N=100,000 scale, under properly density-controlled conditions (a real sampling-density confound was found and fixed along the way): growth-class discrimination remains genuinely ambiguous at the tested measurement depth, while the selectivity phenomenon itself persists robustly across a 30× increase in N, but its spatial coherence appeared to substantially degrade over that same range (BW.6) — the sharpest new open question that consolidation left behind.

Addendum (Part III, BX–CB), at v17. BW.6's coherence-degradation question is now understood, not merely flagged. Addendum BX first established, independent of H.17's substrate entirely, that the triangle term's locality is driven by overlap concentration — triangles sharing nodes — rather than raw triangle count, degree-controlled and replicated on toy substrates. Addendum BY confirmed this transfers to H.17's actual substrate and, critically, found that distant overlap-anchored partial-selection clusters coexist independently and superpose exactly rather than interacting — which supplied a direct, testable mechanism for BW.6: if large substrates ignite many such independent clusters simultaneously, a global contiguity statistic would report fragmentation even where every individual cluster remains internally pristine. Addendum BZ executed this prediction, registered before testing, at Addendum BW's own N≈90,778 scale: a mechanism-constructed 10-cluster state reproduces BW's reported 14% global contiguity almost exactly (0.171) while per-component direction-coherence (0.893 mean) exceeds every reference level in the project's record, including BT's original N=3,000 result. BW.6's apparent degradation is, on this evidence, an artifact of measuring a multi-cluster state with a single-cluster statistic — not locality failing at scale. Full adjudication against BW's own literal survivor state (rather than a mechanism-constructed one at matching N) remains open, pending recovery of BT/BU/BV/BW's exact coupling-reduction parametrization — the same open item Addenda CA and CB independently ran into from a different angle (see below) and which v17 treated as one standing item rather than three separate flags. Separately, Addendum CA tested the leading candidate explanation for that parametrization discrepancy directly — whether a signed rather than magnitude associator reduction would let generic initial conditions reach stationary partial states, as the archived

record reports and this project's reconstructions had not — using the natural G2-invariant 3-form, and found it does not, on either substrate, ELIMINATING that specific explanation while logging a new IC-dependent ignition-threshold disorder signature. Addendum CB then explained that disorder signature: ignition proceeds through a sharp activity bottleneck, and the surviving nucleation sites are significantly overlap-enriched (degree-controlled), directly connecting CA's finding back to the BX–BZ mechanism and reframing this entry's own standing question — "where do the structured seeds required for selection come from" — into the sharper, smaller question of what determines how much initial amplitude lands on high-overlap sites, and whether that is itself derivable from Postulates 1–2.

Addendum (Part III, CC–CE), at v18. Addendum CC introduced R_i (pairwise triangle-overlap) as a stronger predictor of individual-seed ignition threshold than the register's existing T_i, and found the octonion associator's specific magnitude contributes little beyond topology, given fixed topology. Addendum CD demonstrated that a prior session's exact generating code can be recovered via conversation search rather than requiring a surviving cached file or the architect resupplying material, and confirmed every CC finding on the real by1 substrate. Addendum CE applied the same technique to BU/BV directly, recovering their exact lattice substrate and coupling-scale convention to a verified match, and eliminated three further candidate explanations for the BT/BU/BV discrepancy, one by a general mathematical proof.

Addendum (Part III, CF), this consolidation: the BT/BU/BV generic-IC discrepancy is RESOLVED. Following CE.6's recommendation, conversation search located BT.4's and BV.3's own literal dynamics code. The cause: generic ICs in the original code are drawn from Uniform(0.3,0.8), not the Uniform(−0.3,0.3) every session since BY.1 had used, including CC–CE earlier in this same session. Reproducing BT.4's and BV.3's exact configurations with the recovered code and correct IC gives genuine partial selection from generic ICs — qualitatively on a reconstructed orbit substrate and near-exactly on the verified-recovered lattice substrate, matching BV.3's own reported figures at two of five tested points. A clean isolation test found the IC convention alone necessary and sufficient. This closes the parametrization-recovery item as it pertains to the generic-IC discrepancy specifically; CA–CE's own findings stand, now understood as characterizing a distinct, real IC regime rather than a failed reconstruction of BT/BU/BV's own claim.

Addendum (Part III, DS–ED), at v28: the most extensive single-thread investigation of this entry to date — twelve addenda testing whether purely algebraic (non-metric) adjacency, static or genuinely sequential, can escape the log-diameter default without presupposing an ambient metric. Two independently-designed constructions (Addendum DS: a static associator-filtered random 2-complex, and a sequential boundary-driven gluing process) both land in the log-diameter universality class — the sequential version even more decisively (R^2=0.9936) — closing off the specific hope that a purely local, non-metric, algebra-native adjacency rule could supply escape for free. A third, purely deterministic variant (no injected randomness after the seed) closed into an exactly finite structure (N=33, machine-precision-exact, 8/8 seeds) — an independent, dynamically-derived confirmation of Addendum AN's compactness obstruction. Addendum DT executed a previously-flagged, unpursued lead (Bianconi–Rahmede Network Geometry with Flavor, noted at Addendum K) by generalizing the construction to a dimension-parametrized family; all tested dimensions (d=2,3,4,6) remain log-diameter, ruling out gluing dimension as an escape route on its own. Addendum DU then sharpened the entry's stakes directly: Cheeger's inequality plus Friedman's theorem on random regular graphs establish, and this addendum verifies numerically rather than only asserts, that H.17's generic/default branch is not merely unable to run Part IV's specific derivation — it is structurally incapable of supporting any real-space coarse-graining scheme at all. But not every log-diameter construction shares this fate: the sequential-gluing construction, despite its own log-diameter volume growth, was found to have a genuine, shrinking conductance bottleneck unlike a plain expander, confirmed via the

Fiedler vector (DU.2–DU.4) — diameter and conductance are logically independent properties, and this construction has the favorable one. Addendum DV showed this extends to a full two-diagnostic match against a genuine hyperbolic random-geometric-graph reference (bounded Gromov δ-hyperbolicity and shrinking spectral gap together, matching neither a flat nor a plain-expander reference), though specific to low gluing dimension and not yet understood mechanistically. A methodological error — mixing two different implementations when comparing across gluing dimension — was found, logged, and corrected in place (Addendum DW), which confirmed the diameter-scaling conclusion (DT) survives the correction while retracting a specific conductance-dimension claim (DV.4) that did not. Addenda DX–DZ and EA properly isolated a real, seed-robust, degree-independent driver: a construction-selectivity ratio (how many fresh candidates versus existing-vertex checks inform each growth step), collapsing cleanly onto a single saturating curve (conductance exponent $\approx -0.29 \cdot \ln(\text{ratio})$, $R^2=0.953$) that reaches a precisely-measured ceiling (exponent ≈ -1, $R^2=1.0000$ at its cleanest) once dimension and degree are properly disentangled from it — the strongest, cleanest signal in the whole thread. Addenda EB–EC then found this ceiling's geometric identity remains genuinely open: it numerically matches 2D flat-lattice spectral-gap scaling on one diagnostic while directly contradicting flat geometry's own δ-hyperbolicity signature, and a numerically-corrected, curvature-parametrized sweep of the hyperbolic reference family (Addendum EC) found a genuine trade-off rather than a match — lower curvature moves the conductance exponent toward the ceiling but loses the bounded-δ signature that would make the match meaningful, and no curvature tested reproduces both. Finally, Addendum ED tested and answered a standing methodological question directly: whether this whole construction family's interesting behavior depends on continuous background access during growth. It does not — the double signature survives intact under the strictest possible finite realization of Postulate 1 (a fixed pool of exactly N pre-given configurations, no further continuous sampling at any point) — but this narrows rather than resolves the deeper, still-fully-open question the whole thread sits inside: Addendum AJ.6's own standing discretization question, identical to H.17's own root difficulty per Addendum AN's original finding, and untouched by any of DS–ED's work.

Addendum (Part III, EM), at v31. The v30 EH–EL thread's near-critical decoupling finding (conductance weakens, δ-hyperbolicity strengthens, as glue_bias approaches DS.7/DS.8's own collapse boundary) was tested for generality against two structurally distinct constructions. A dimension-parametrized simplex-gluing family (d=3, verified to reduce exactly to the d=2 case before use) shows the same conductance-weakening direction approaching its own, much higher, collapse boundary (glue_bias between 1.2 and 1.3, versus d=2's ≈0.70), but the δ-hyperbolicity trend, while moving toward positive, never completes a sign crossing within the safely-testable range — the construction's own near-flat diameter (4–5 hops across the full tested N range, unlike d=2's steadily growing diameter) may be suppressing the signal rather than it being genuinely absent. DS.2's static associator-filtered 2-complex — structurally unlike either sequential construction, tuned toward its own percolation threshold rather than a growth-collapse boundary — shows the identical pattern: conductance weakens approaching criticality, δ-hyperbolicity trend moves toward positive without crossing, and the construction is markedly noisier throughout ($R^2\approx 0.4$–0.6 vs. d=2's >0.9), consistent with its established characterization as closer to a plain expander. It does show one unambiguous signature the other two did not: mean diameter itself grows sharply near the percolation threshold (6.8→18.0 at matched N), a textbook correlation-length divergence. Taken together: the conductance-weakening direction is now confirmed general across three structurally distinct constructions; the δ-hyperbolicity sign-crossing that made the original d=2 finding notable appears specific to (or at minimum, much harder to detect outside of) low-dimensional sequential edge-gluing — genuinely underdetermined which, since neither counter-example was pushed fully to its own collapse edge. No formal status-tag change; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY.

Addendum (Part III, EU–EV), at v32. Two independent diagnostics, neither previously

applied to any H.17 construction, were run directly on the DS.7/DS.8-family hybrid substrate. EU computed spectral dimension via random-walk return probability (heat-kernel flow), calibrated first against three reference cases with known answers (2D torus lattice locks exactly to d_s=2 and stays flat; 3D torus lattice and a plain random-regular expander both show expected finite-size behavior) before being trusted on the real construction. Pushed to N=50,000 via a sparse Hutchinson stochastic trace estimator (verified against exact dense diagonalization before use, and made tractable only after a 20× build-speed optimization — batching the associator-norm evaluation that profiling identified as the actual bottleneck, verified exact to machine precision against the original scalar construction): the finite-N spectral-dimension peak grows with N rather than converging (2.66→3.01→≥3.20→4.02 across N=6,000→50,000), passing visibly through d_s=4 without leveling off there — a second, independent line of evidence for Addendum DU's Cheeger-based conclusion that this construction class supports no genuine macroscopic dimension at any tested scale. EV then asked a different question of the same substrate: whether physical dynamics riding on top of an expander-like graph can be confined even though the bare graph itself is not dimension-locked. Using Postulate 2's own already-derived complex coupling (Addendum Z.7/AU's C_ij = W_ij + iΓ_ij, restricted to the substrate's actual edges, no new structure introduced) as a tight-binding Hamiltonian: bare topology is delocalized (inverse participation ratio → 0, consistent with EU), but switching on Postulate 2's real coupling alone flips this to genuine Anderson-style localization (IPR plateaus at a nonzero, N-independent value, confirmed against seed-to-seed noise before being trusted, and independently confirmed via real-space localization length, R^2=0.78–0.95 once fit correctly on distance-shell averages rather than raw per-site values). The genuine octonion phase term Γ_ij specifically weakens this confinement rather than strengthening it (plateau IPR roughly a third of the real-only case) — the discrete-graph analogue of weak antilocalization, reported precisely rather than smoothed toward the motivating hypothesis that disorder and phase would both add to confinement. No formal status-tag change; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY.

Addendum (Part III, FE–FH), at v35: the session's own opening hunch, chased to exhaustion — no escape found on any axis, but the project's actual diameter measure is finally mechanistically explained. Before testing anything, the hunch itself (a non-compact algebraic substrate might escape the compact/finite wall FA.7/FB.5 had just converged on) was checked against Addendum AN.5–7's own general argument and found already closed on both branches, independent of any specific construction: compact discrete Aut(𝕆) subgroups are finite by elementary topology; non-compact split-octonion routes are closed twice over (Addendum Z.16–17's locality-violating growth, and independently, Milnor–Švarc/Gromov's exponential-growth guarantee for lattices in simple Lie groups). Addendum FE located the one genuinely untested cell instead: {lattice-ordered vertex set} × {non-metric algebraic adjacency} — every DS-family construction to date uses Haar-random points on S^6, and Addendum BU's own lattice substrate uses k-NN (ambient-distance, "secretly metric") adjacency, not DS.2's associator-threshold rule. Built from scratch (independent octonion algebra, machine-precision validated) and tested at N=98→88,272 (900×, exceeding the project's own 512× convention): still log-diameter, 3/3 seeds, decisively (R^2=0.994–0.997 vs. 0.977–0.983 for power-law); direction-degeneracy, a structural difference from Haar-random points, checked and found negligible (≈1% duplication). Addendum FF tested FE's own follow-on hunch — that DS.2's independent-triple-sampling architecture, not vertex-set compactness, is what forces log-diameter — directly on the lattice substrate: a full-boundary sequential/correlated triple-selection rule gives statistically indistinguishable results (log R^2=0.979–0.997, 3 seeds), and this axis turns out to have already been settled by the project's own prior record — grow_hybrid's genuine sequential construction (Addendum DS) is the *more* decisively log-diameter of the two (R^2=0.9936), a check that should have been run before the hunch was proposed and is logged as such. A sharper, fixed-window frontier-restriction attempt produced a dramatic near-linear diameter blowup; checked before being reported

as a positive result and correctly diagnosed as a pure construction artifact (a global fixed window forces the entire build into a serial chain once N exceeds it, giving diam≈N/window mechanically, confirmed 2/2 seeds against the exact predicted ratio). Addendum FG built the corrected, properly N-scaled (proportional-fraction) version of that frontier and found a real effect the fixed-window version was not: an order-of-magnitude increase in log-diameter's own slope, cross-validated against a simple geometric-decay formula at three independent fraction values (predicted-vs-observed R^2=0.995) — still within the log-diameter class, but the first result in this whole arc distinguished by more than a marginal fit-quality difference. Addendum FH then tested whether that slope formula was algebraic before attempting to derive it from the octonion structure, as originally asked: a null model with the identical proportional-window architecture but zero octonion content (fixed-probability Bernoulli acceptance replacing the associator threshold) reproduces FG's slopes within ordinary seed noise at every fraction — DERIVED (negative), confirmed by direct substitution rather than argued: the effect is a fact about sequential random-graph construction, not the octonion algebra, the associator, or G2 symmetry. The correct first-principles derivation, supplied in its place, is graph-theoretic: a worst-case greedy backward-jump argument reproduces FG's formula exactly and was rigorously bracketed — a proven, algebra-independent lower bound, and a Monte Carlo-verified upper bound (2,000 trials, eight orders of magnitude in N, R^2=1.00000). A follow-up test of that bracket's practical relevance then caught a genuine bug in its own setup (vertex *labels*, assigned by a random permutation, had been mistaken for true temporal build order) — corrected, the finding survives and strengthens (the true newest-to-oldest pairwise distance is disconnected in 36/36 trials, not 4/36) — and, pursued further, replaced a vague diagnosis ("a scattered low-degree population") with a precise, directly-verified mechanism: an asymmetric recency boundary layer. The most recently introduced ≈5–10% of vertices are systematically starved of connectivity (mean degree falling from ≈7.5–8 to 3.80 in the final 5%, 43.5% at degree ≤2) because the construction simply ends before they can accumulate received edges — confirmed not just against degree but directly against measured eccentricity, which tracks a source's normalized temporal position in a close-to-monotonic relationship. An exact closed form for this project's own multi-source-measured slope, now precisely targeted rather than vaguely gestured at, remains open. No formal status-tag change; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY.

Addendum (Part III, FO), at v37. The one remaining untested cell in the FE–FH arc — sequential/correlated edge-selection (Addendum DS.7/DS.8's own hybrid gluing rule specifically) re-run on a lattice-ordered vertex set rather than its native Haar-random points — is closed. A first pass at DS.7/DS.8's own 256× range convention was genuinely ambiguous (log R^2=0.9655 vs. power-law R^2=0.9679), unlike every other construction in this arc; extended to 1024× range (N=200–204,800), the log fit pulls decisively ahead (R^2=0.9764 vs. 0.9656), corroborated by two independent mechanistic signals — a shrinking free power-law exponent (0.1205→0.1074, the same degenerating-exponent pattern DS.8 itself showed) and decelerating diameter increments per doubling. Negative: the fifth structurally independent construction axis tested across this project's history on H.17's log-diameter branch, and the fifth to return log-diameter. No formal status-tag change.

Addendum (Part III, FQ–FR), at v37: the fork itself is revisited directly, and restated. Every session from DS through FO attacked H.17's first horn (pregeometric, bounded-degree, presumed log-diameter); none had returned to the second horn's circularity charge — that polynomial growth requires an imposed lattice — since it was first raised. Addendum EO's minimal-vector-adjacency construction (v31, Postulate 4 dependency) already bore on this directly without being framed that way: bounded degree and the correct growth degree from an order in the octonions, a purely algebraic object with no metric input at any step. FQ sharpens and tests this. Building Coxeter's canonical maximal order (the octavian integers) from scratch — independently reproducing, and correcting via a coordinate swap, the historical Kirmse's-mistake closure failure, verified against controls — FQ shows growth degree is exactly 8 by theorem regardless of which order is chosen (closing a gap Addendum

EO.5 itself flagged), and that Postulate 3's own non-patching dynamics forces the canonical choice: the non-canonical integer lattice's minimal-vector graph is exactly triangle-free (bipartite by coordinate-sum parity), so the entire triangle-term mechanism vanishes identically on it, while the canonical maximal order has over a million triangles in a comparably sized ball. On a genuine bulk region of this substrate (constant-degree shells, ruling out the ball-boundary artifact an initial pass mistook for a real result), polynomial growth and genuine algebraic selection are shown to coexist for the first time in this project's history — but the selected sets prove to be exactly centrally symmetric against 300 random controls, no spatial localization at all. H.17's second horn is restated: not circularity, but symmetry — a substrate symmetric enough to force genuine polynomial growth is symmetric enough that any admissible coupling becomes constant on group orbits, and orbit-restricted selection is necessarily globally spread. Addendum FR then tests this restatement to its limit, within the same single addendum, correcting its own working hypothesis twice along the way (an informal argument that un-normalizing the coupling shouldn't matter is tested and found false; a prediction that a signed, un-magnituded coupling would be genuinely chiral is tested and found false) before landing on the actual mechanism: exact invariance under global negation, $x \to -x$, forced by -1 being central in the octonions and fixed by every automorphism — not orbit structure under some specific finite symmetry group, which the first (false) hypothesis had proposed and the second test directly refuted. This is verified, not merely argued, for every literal-coefficient candidate the theory currently admits as non-patching (associator-norm, $|\varphi|$, φ^2, independently built and tested, bit-identical results), and confirmed to survive two structurally distinct chirality-escape attempts by two independent mechanisms. A final test of whether abandoning additive-lattice substrates entirely could escape the obstruction finds it coherent in principle — the project's own orbit substrate is a real, verified example of a forced construction that is not negation-closed — but not rescuing: cataloguing all four construction paradigms this project has built (compact orbit, non-compact orbit/split-octonion, closure-based lattice, non-metric algebraic-adjacency), each fails to deliver both genuine polynomial growth and freedom from this symmetry, each for its own independently-established reason (compactness; Milnor–Švarc/Gromov; the parity obstruction itself; log-diameter). No formal status-tag change; a fifth, unconstructed paradigm — unbounded, non-closure-based, still forced rather than patched — is not ruled out, but every construction method actually built has now been tried and independently closed.

Addendum (Part III, FS–FU), at v38: the fork itself set aside, a genuinely new diagnostic lens applied for the first time — holographic tensor-network structure — and pursued to a decisive, if still negative, conclusion. Rather than continuing to characterize which construction paradigms can or cannot achieve growth and locality together (Addenda DS through FR), FS asked a different, sharper question of the theory's own established substrate: does it show the specific minimum-cut scaling signature (logarithmic in boundary-region size) that genuine holographic/HaPPY-code structure requires, as opposed to the linear scaling generic expanders show regardless of any local curvature signature they may carry? Validated against a regular tree (positive control) and a random 4-regular graph (negative control) before being trusted, the diagnostic returns a clean negative on the substrate as previously constructed: linear-dominant, despite the real local hyperbolic signatures (bounded Gromov δ-hyperbolicity, a shrinking conductance exponent) already on record from Addenda DU–DV and EM. A constructive response — enforcing a hard $Q=7$ vertex-valence cap, the minimal genuinely hyperbolic $\{3,Q\}$ triangulation ($\kappa = -1/6$) — succeeds for the first time in this project's history at producing a real, connected, negatively-curved, proportionally-growing consistent core, but does not move the full-graph RT result: the core is swamped by a defect-laden majority regardless of which of three independent suppression levers (search width, soft-cap tolerance, load-balancing) is applied, with search width the only one that helps at all, and only by converging toward the identical finite-closure phenomenon Addenda AN and DS.6 already established by unrelated methods — a third structurally independent route to the same trap. FT tests and closes the one hypothesis this result structure left open (that the

core itself is holographic but full-graph noise conceals it: refuted directly, the isolated core's own RT scan is if anything more linear-dominant than the full graph), then implements the structurally correct constraint FS's own result had diagnosed as missing: genuine edge-level {3,Q}-manifold consistency (every edge capped at exactly 2 incident faces), not merely vertex-valence capping. True consistency under this corrected, stronger definition is two to three orders of magnitude rarer than FS's vertex-only core at fixed search width, but a proportional-search regime sustains a genuinely stable, non-vanishing density across a full 32× range in N — the first stable consistency density this entire H.17 arc has produced. This is immediately and substantially qualified: the resulting true-consistent vertices form no connected structure whatsoever (42 isolated singletons at N=128,000) and show zero spatial clustering under a rigorous, well-powered statistical test against a random-placement baseline, confirmed across three separate density regimes spanning a 25× range — local combinatorial consistency and spatial organization are, on this evidence, fully decoupled properties of the construction, not merely rare-but-correlated ones. FU takes the one remaining candidate explanation FT's own closing discussion named but did not test — a complete absence of any feedback channel between spatially separated closure events in the underlying growth rule — and tests it directly by building an explicit feedback-biased variant and sweeping its strength: the result is a sharp collapse threshold rather than a rescue, with feedback strong enough to plausibly matter destroying growth entirely via a runaway preferential-attachment dynamic, and feedback weak enough to preserve normal growth showing no clustering benefit at all. This closes FT's own open mechanistic question in the negative and constitutes a seventh independent, structurally distinct manifestation of the growth-vs-global-consistency tension now documented across this project's full history (AN, DS.6, FS's three levers, FT's proportional-search thread, and now FU's engineered feedback). No formal status-tag change; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY.

STATUS: CONJECTURE/OPEN, PERMANENT ENTRY — genuine open structural problem, not resolved by any session to date. Sharpened from a diffuse pattern into a concrete, citably-grounded either/or (H.17); several candidate routes tested and closed (K–N, AN); a direct dynamical test gives a real but qualitatively inconclusive positive signal that leans pessimistic on further testing (AO–AQ, AS); two new lattice discretization candidates found, revealing a new lattice-vs-orbit tension (AV–AW, AW.5) and a new growth-vs-selection trade-off (AX, AX.10); both trade-offs given working counterexamples using the triangle-term mechanism, with growth-class scaling and cross-instance generality left open (BT–BW). At v17 (BX–CB): the triangle term's locality mechanistically understood as an overlap-concentration effect, degree-controlled and cross-validated on toy and real substrates; BW.6's coherence-degradation question reframed and its proposed resolution SUPPORTED for mechanism-constructed states, literal-state adjudication pending; the underlying structured-seed requirement for selection dynamics — visible independently in BN.6, AO/AC/AL, and now CA/CB — mechanistically explained for the first time as overlap-nucleated ignition. At v18 (CC–CE): a stronger per-seed-variance predictor found (R_i); BU/BV's substrate and coupling-scale recovered and verified. This consolidation (CF): the generic-IC discrepancy itself — RESOLVED, cause identified as an IC-convention mismatch, not any property of the coupling, substrate, algebra, or dynamics order. Growth-class scaling toward AX's N=100,000 benchmark (BW), AW.5's lattice-vs-orbit tension itself, and full adjudication against BW's own literal survivor state remain untouched and open. At v23 (Addendum CV): the negative case on the honest pregeometric branch is extended from the bare pairwise mechanism (AS) to the fully-characterized triangle-term coupling class (Addendum CU). Triangle count on the honest substrate (genuinely random regular graphs, no presupposed adjacency) is O(1) as N→∞ — a classical fact about configuration-model random regular graphs, verified across a 64× range in N — so the triangle-term mechanism's reach vanishes as a fraction of the substrate and cannot affect the

asymptotic diameter-scaling class, independent of dynamics calibration. Combined with Addenda BT–BW's confirmation that the same mechanism succeeds only on presupposed-structure (lattice/full-orbit) substrates, every dynamical selection mechanism tested to date — bare pairwise (AS, AX) and triangle-term (BT–BW, CV) — shows the identical pattern: works only where geometric structure is already assumed, fails or cannot engage on the honest branch. The negative case is now comprehensive rather than partial. At v28 (Addenda DS–ED): the log-diameter default branch's coarse-graining incapacity is now independently VERIFIED (not just argued) via Cheeger/Friedman; a specific sequential-gluing construction, despite itself being log-diameter in the ordinary volume-growth sense DS/DT test for, is found to have a genuine, hyperbolic-reference-matching double signature (bounded δ-hyperbolicity + shrinking conductance) on two independent diagnostics, driven by a construction-selectivity ratio, degree- and dimension-independent, reaching a precisely-measured ceiling (conductance exponent ≈ −1) whose geometric identity remains open after testing against flat, expander, and curvature-swept hyperbolic references; confirmed independent of continuous background access under the strictest finite realization of Postulate 1, which in turn clarifies rather than resolves this entry's shared root with Addendum AJ.6's own standing, still fully open discretization question. The entry's formal status is unchanged. At v30 (Addenda EH–EL): a return to Postulate 1's own wording reconnected this thread to a previously-closed Trugenberger/Kelly mapping and found a genuine, verified, N-independent bimodal nucleation transition in the growth process, confirmed on two independently-coded constructions; tested and found this substrate still does not escape Addendum AX.10's growth-vs-selection trade-off despite clearing Addendum CV's specific negative mechanism; identified the driving mechanism via direct basin-counting (a converged 2-attractor landscape versus a genuinely rugged reference) and traced roughly half of it to a real, partial, statistically-confirmed causal channel; formally closed (without resolving) the standing DZ.3/grow_contextual contradiction as a permanent open discrepancy; and found the same bimodal signature recurring a third time, in landscape type itself, with its severity shown to peak at an internal critical boundary — sharpening the entry's synthesis into a critical-fluctuation-bimodality claim confirmed at three independent levels. At v32 (Addenda EU–EV): two new diagnostics, spectral dimension and Anderson-style localization, both run for the first time on this construction family. Spectral dimension shows no dimension-lock at any tested scale up to N=50,000 (a second, independent confirmation of Addendum DU's conclusion); physical dynamics on top of the substrate, using Postulate 2's own real coupling, is nonetheless genuinely localized — the bare graph's lack of a fixed dimension does not by itself mean physical fields propagating on it are unconfined. At v35 (Addenda FE–FH): the session's opening hunch that some untested construction axis might let a non-compact substrate escape the log-diameter default was checked exhaustively and found not to — compactness is closed on general argument (re-confirming AN.5–7), and independent-vs-correlated edge-sampling architecture is closed on direct test, corroborated by the project's own prior grow_hybrid record (R^2=0.9936). A third axis, frontier breadth, yields a real, mechanistically-explained, cross-validated (R^2=0.995 across three fraction values) effect on the log-diameter slope — but proven, via direct null-model substitution, to be algebra-independent: a fact about sequential random-graph construction, not about the octonion algebra or G2 symmetry. Chasing the resulting graph-theoretic bracket caught a genuine indexing bug in its own extreme-pair test (corrected, strengthening the finding) and surfaced this entry's first precise account of what actually drives the project's own diameter measure: an asymmetric recency boundary layer (the most recently introduced ~5–10% of vertices, systematically under-connected because the construction ends before they can mature), confirmed directly against measured eccentricity. No formal status-tag

change; the entry's formal status remains unchanged. At v37 (Addenda FO, FQ–FR): the last flagged gap in the log-diameter branch is closed (FO — sequential/correlated edge-selection on a lattice-ordered vertex set, still log-diameter, the fifth independent negative). The fork's second horn is revisited directly for the first time since it was raised and RESTATED: not circularity but symmetry — genuine polynomial growth and genuine algebraic selection are shown to coexist for the first time in this project's history (FQ, on Coxeter's canonical maximal order in the octonions, itself forced by Postulate 3's own dynamics since the non-canonical alternative is exactly triangle-free), but selection is confirmed unable to localize, verified across every literal-coefficient candidate the theory admits and surviving two independent chirality-escape attempts, traced to exact invariance under global negation forced by −1's centrality (FR). Every construction paradigm this project has built across its full history — compact orbit, non-compact orbit, closure-based lattice, non-metric algebraic-adjacency — now fails this entry's full requirement (growth and locality together), each for its own independently-established reason; a fifth, unconstructed paradigm is not ruled out. At v38 (Addenda FS–FU): a genuinely new diagnostic (Ryu–Takayanagi min-cut scaling, validated against tree and random-4-regular controls) finds the honest substrate is generic-expander-dominant, not holographic, despite real local hyperbolic signatures already on record; a constructive fix (a hard Q=7 vertex-valence cap, FS) produces the project's first genuine connected negatively-curved consistent core, but does not rescue the full-graph result, and every tested defect-suppression lever either converges toward the same finite-closure trap found independently at AN/DS.6 or is counterproductive. Isolating that core directly does not help either (FT, decisively closed). The structurally correct, stronger constraint (genuine edge-level {3,Q}-manifold consistency, not merely vertex-valence capping) is two to three orders of magnitude rarer at fixed search width, but achieves a genuinely stable, non-vanishing density under proportional search across a 32× range in N — the first stable consistency density in this entry's entire history — though the resulting consistent vertices are found to be completely disconnected and spatially uncorrelated with one another, tested and confirmed across three separate density regimes (FT). A direct test of the one remaining candidate explanation — engineered local feedback between closure events — finds a sharp collapse threshold rather than a rescue: no feedback strength both preserves growth and produces clustering (FU), a seventh independent, structurally distinct manifestation of this entry's central growth-vs-consistency tension. Entry's formal status unchanged. At v39 (Addenda FV–FW): an eighth and ninth independent manifestation, testing whether Fano-plane-collinearity structure supplies genuine locality — boundary-confined selection found real but proven, via three exact theorems, incapable of carrying more than one discrete bit or escaping the bulk wall (PSL(2,7)-invariant, order 168). At v42 (Addenda GJ–GL): the Heisenberg nucleation mechanism confirmed a genuine bulk phenomenon across a 92× N-range (GJ); a generic-frame embedding escapes the φ≡ 0 degeneracy, enabling the first three-candidate literal-coefficient comparison (GK); AZ.13's discharge condition composes successfully on this substrate for the first time, five theory requirements cleared simultaneously (GL). At v44 (Addenda GR–GT): generalized to $H_3(\mathbb{Z})$ (dimension 7, no coordinate compression); a genuine substrate-type distinction confirmed after two rounds of self-correction — only $H_1(\mathbb{Z})$ plateaus among six constructions tested, cause not yet identified (GR); four external hunches tested against this entry, all negative (GS); the pattern behind four prior negative results proven as an actual theorem — no G2-invariant construction can select a preferred direction (GT). At v45 (Addenda GV–HC): a second, symmetry-based obstruction axis identified alongside GT's counting-based one (GV); nilpotency class and every tested ratio excluded as explanations for $H_1(\mathbb{Z})$'s uniqueness, seven constructions now tested (GW, GX); GJ's own plateau claim reconstructed at full rigor to radius 28 and CONFIRMED, but at

a materially lower value than originally reported (GY, HC); three further candidate escapes tested and closed across two parallel sessions (GZ, HA, HB). Entry's formal status unchanged. At v46 (Addenda HD–HF): the most thorough sweep GR.24 has yet seen finds eight constructions and six candidate discriminators now all negative except $H_1(\mathbb{Z})$'s own plateau, with a general mechanism identified for why M-tensor anisotropy specifically could never have discriminated (HD); separately, catching a real instance of this project's own working-memory failure mode, two of an external literature survey's four proposed H.17-resolution families are found already closed weeks before the survey ran (Family A/GV, Family B/K–N,L,EH), and the remaining two are then genuinely tested and closed — entanglement- derived adjacency by direct two-part correlation test (HE), on-shell symmetry by structural argument after reading its primary source in full (HF). All four externally-proposed families are now closed. Entry's formal status unchanged.

BACKFILLED AT V45 — not previously entered in this section (found stale since v38, logged per the no-silent-corrections standard, not silently patched). The following restores Addenda FV through GT's own contribution to this entry's narrative; all were already current in Part VI's own standing register throughout.

Addendum (Part III, FV–FW), at v39. An eighth and ninth independent manifestation, tested via a different route: whether Fano-plane-collinearity structure in the octonion algebra itself can supply genuine locality. FV builds the project's first positive-cone semigroup (N^8), genuinely escaping FR's negation-cancellation mechanism at the correct polynomial growth degree (8) — but three symmetric associator-based couplings show null localization, generalizing into a conjecture that any coupling invariant under the octonion's own structure constants inherits some discrete symmetry of them. A genuinely position-dependent coupling (Fano-plane collinearity of a point's active support) breaks this with real, mechanistically-forced, dynamically-robust selection (0/70 collinear-support survivors, 5 seeds × 4 regimes) — but confined to the cone's boundary, since the bulk (support ≥5) is provably flat by Fano design regularity. FW answers the resulting question directly with three exact theorems rather than further examples: the boundary signal can never carry more than a single discrete bit regardless of magnitude (trilinearity of the associator); the bulk wall is exact, generalizing to any Steiner triple system rather than being octonion-specific; and any coupling built purely from the collinearity data is invariant under exactly PSL(2,7) (order 168), the exact symmetry floor for this coupling class. A richer vector-valued coupling escapes the bulk degeneracy, but the escape is ~99.7% explained by generic vector-angle proximity rather than Fano-specific structure — an honest negative, not a route forward. No formal status-tag change; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY.

Addendum (Part III, GJ–GL), at v42. GJ runs a seven-scale study of the Heisenberg construction first cleared at v41 (Addenda GE–GI) — radius 9–28, N=2,845→261,815, a 92× range exceeding most of the project's working substrate sizes — and confirms the underlying nucleation mechanism is a genuine bulk phenomenon, not a finite-size artifact; separately finds the Postulate-3 literal-coefficient tie-back structurally inaccessible on this specific embedding, since φ is identically zero, confined to a single 3-dimensional subspace. GK escapes that degeneracy with a generic-frame embedding, enabling a three-candidate comparison for the first time: φ^2 widest in all four tested instances, associator-norm/$|\varphi|$ tied in all four — the tie proven structural, forced by any single-3-plane embedding. GL tests AZ.13's discharge condition on this same substrate for the first time (rank 3/7, anisotropy exactly 1.0, though later identified as a near-tautological consequence of the single-3-plane embedding rather than independent geometric content) and composes it successfully with genuine Postulate-3 dynamics — the cleanest of the substrate types tested to that point, meaning five separate theory requirements clear simultaneously on one construction for the first time in the project's history. No formal status-tag changes.

Addendum (Part III, GR–GT), at v44. GR builds $H_3(\mathbb{Z})$, dimension exactly 7, embedding na-

tively into all seven octonion imaginary directions with no coordinate compression — avoiding the shared root cause behind GJ/GK/GL's near-tautological findings — confirming the nucleation mechanism replicates in full and composing AZ.13's discharge condition cleanly via the full 35-parameter joint dynamics on a genuinely non-degenerate M-tensor. Extending to six constructions in total to isolate the variable behind a genuine decay-vs-plateau distinction, a tentative dimension-threshold hypothesis is proposed, directly falsified by a further construction, and retracted in place rather than removed (GR.19→GR.22) — the honest net position is that only $H_1(\mathbb{Z})$ plateaus among six independent constructions now tested, for a cause not yet identified. GS tests four hunches proposed by two different external LLM collaborators against this entry and AZ.13 — a correlation-structure locality diagnostic, growth measured on the dynamically-selected surviving sector, cone/M-tensor composition, and a non-associativity bounding claim — all four tested directly and found negative, each via an independent method with its own finite-size or artifact control. GT builds G2 = Der(O) directly from the project's own octonion structure constants and proves, rather than continues to observe empirically, that G2 has no nonzero fixed vector in Im(O) and acts transitively on S^6 — so no G2-invariant construction, of any kind, can select a preferred direction — and that $\Lambda^3(\text{Im O})$ has an exactly 1-dimensional G2-invariant subspace, the classical associative form φ recovered independently rather than assumed. This unifies Addenda FR, FV.7–FV.9, GL.4, and GS's own cone result as one fact viewed four times, and supplies a predictive filter for future proposed mechanisms: check the representation content before simulating. No formal status-tag changes; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. NOTE, found while backfilling this section at v45: GT had no dedicated paragraph in Part VI's own register either, prior to this consolidation's own update there — logged here rather than treated as unique to this section.

Addendum (Part III, GV–HC), at v45. GV checks whether the Bombelli–Henson–Sorkin no-go theorem for discrete Lorentz invariance bears on this entry's own fork; found not to apply directly (no presupposed continuum here), but its mechanism identifies a second, symmetry-based obstruction axis alongside GT's own counting-based one — currently untestable, since nothing has escaped the counting floor to reach a stage where it would bite. GW and GX exclude nilpotency class, center ratio, and derived-subalgebra ratio as explanations for why $H_1(\mathbb{Z})$ uniquely plateaus, extending the tested-construction count to seven; GW separately finds $H_3(\mathbb{Z})$'s own AZ.13 anisotropy figure (0.386, GR) is not scale-invariant under its own reported convention — a genuine scope gap, logged in the AZ.13 entry, not a reversal. GY and HC together complete Addendum GR.23's own long-flagged question: does GJ's original plateau claim survive reconstruction at full rigor out to its own largest tested scale? HC reaches radius 28 (N=261,815, exactly GJ's own largest documented point, after approximately 121 continuous hours — the longest single computation in this project's history) and finds the local decay exponent collapsing monotonically across the five largest radii, the signature of a genuine asymptote: $H_1(\mathbb{Z})$ does plateau, upholding GJ's qualitative claim, but at a materially lower value (est. 1.5–1.7%) than GJ's originally reported 2.05–2.39%. Separately, across two parallel sessions: GZ tests dynamically-gated growth admission and finds it asymptotically self-defeating via k-NN coupling saturation on a compact manifold; HA tests nonlocal Jordan–Wigner fermionization and, after two further self-corrections logged in place, finds $H_1(\mathbb{Z})$'s locality advantage is an ordinary growth-degree effect, not a special pipeline interaction; HB tests BFSS Matrix Theory and correctly retracts an initially exciting dimensional-collapse signal via calibration against synthetic data before trusting it. No formal status-tag changes; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY, on firmer empirical footing than before. See Addenda GV, GW, GX, GY, GZ, HA, HB, HC in full.

5.3 2b. SFT-Gemini Review (2026-07-19) — added at v12 (Addendum BA.2)

A Gemini-authored proposal was reviewed in-session on 2026-07-19. Verdict: factual recall was accurate, but the central synthesis overclaimed, reinventing an already-tested and explicitly-declined mechanism. Adoption was recommended against; nothing from the proposal enters the theory.

STATUS: TRACKED, CLOSED — reviewed, declined with citable reasons; no theory content adopted.

5.4 2c. Gemini Audit of Addenda BH–BK (2026-07-21) — added at v15 (Addendum BL)

An audit requested by this project, per its own standing policy that results of BH–BK's consequence require a skeptical pass from outside the generating session. The returned audit's headline claim (a "fatal algebraic cancellation" in the triangle-force law) was checked directly and found to misdescribe the actual implementation, resting on a non-reproducible verification script. A second claim was checked and found to be the opposite of what was observed. One point was confirmed legitimate and already-known: the natural-scale derivation chain (BJ) was, at the time, plausible but not fully first-principles.

STATUS: AUDIT COMPLETE (text/record level) – eleven total instances found and repaired across v11-v21. Numerical re-execution verification remains open as an optional distinct task. Fifth instance (v16): Addendum BN's own search methodology initially failed to reproduce a known result due to a parameter-provenance gap (guessed values substituted for unrecorded actual ones); caught, flagged PROVISIONAL rather than reported as final, and corrected once the architect supplied the source document – logged as a correction in Addendum BN.2-BN.5, not silently fixed. Sixth instance (v16): Addendum BS.3-BS.4 – an initially-favorable pooled statistic on cross-topology generalization was found, on further scrutiny, to be dominated by the derivation network's own in-sample effect rather than a genuine shared effect across fresh networks – caught and corrected within the same addendum before being reported as final. Seventh instance (v17): SFT_Addenda_Archive.docx was found to jump directly from Addendum BL to Addendum BX – Addenda BM through BW, all of v16's own consolidation, were never appended. REPAIRED AT v20 – see the ninth and tenth instances below.

5.5 2d. Gemini Audit of v21 (2026-07-23) — added at v22, no addendum letter

A Gemini-authored review of v21 was checked against primary sources before any response was given, per standing practice. Findings: two of its four "critical bottlenecks" were not new — the non-differentiability /G2-invariance objection (its S3.1) is a repeat of a claim already refuted at Addendum BL.4 (v15), now recurring a second time from an independent external-review session; the θ(N,k) decay point (S3.2) restates Addendum BM.5's own already-logged caveat, framed as a fresh discovery. One point (S3.4, spacetime fragmentation) was numerically accurate but ignored Addendum BZ's already-established reframing (per-component coherence 0.893 mean, exceeding all references; the cited 14% figure is a single-cluster statistic misreading a multi-cluster state). One point (S3.3, Route A/B fork) was accurate and added nothing beyond the existing record. All four of its "mandatory directives for v22" required no new action; the differentiability footnote it recommended was explicitly rejected as resting on a refuted premise.

STATUS: AUDIT COMPLETE. Recommended action: a standing pointer to Addendum BL.4 near Postulate 3's triangle-term discussion, since this specific objection has now recurred twice from independent external reviews. IMPLEMENTED at v25, following Addendum DA.2's discovery of a separate, genuine differentiability finding that needed to be clearly distinguished from BL.4's settled (and still correctly refuted) G2-invariance claim — see Part II, Postulate 3.

5.6 3. Documentation-Propagation Audit

STATUS: AUDIT COMPLETE (text/record level) — four total instances found across v11–v12, all repaired; a fifth and sixth found and logged at v16; a seventh found at v17; a twelfth confirmed and repaired at v24 (see below). Numerical re-execution verification remains open as an optional distinct task. Fifth instance (v16): Addendum BN**'s own search methodology initially failed to reproduce a known result due to a parameter-provenance gap (guessed values substituted for unrecorded actual ones); caught, flagged PROVISIONAL rather than reported as final, and corrected once the architect supplied the source document — logged as a correction in Addendum BN.2–BN.5, not silently fixed. Sixth instance (v16): Addendum BS.3–BS.4 — an initially-favorable pooled statistic on cross-topology generalization was found, on further scrutiny, to be dominated by the derivation network'**s own in-sample effect rather than a genuine shared effect across fresh networks — caught and corrected within the same addendum before being reported as final. Seventh instance (v17): SFT_Addenda_Archive.docx was found to jump directly from Addendum BL to Addendum BX — Addenda BM through BW, all of v16's own consolidation, were never appended. That gap remains unrepaired; full derivation text for BM–BW is still not recoverable from the archive.

New at v20: a ninth instance found and repaired – the seventh instance's own gap, closed. The Addenda BM-BW archive gap flagged at v17 (seventh instance) and still open at v19 was repaired this session: all thirteen addenda were recovered from the original source session via conversation_search and re-inserted into SFT_Addenda_Archive.docx in full, using the same recovery technique demonstrated at Addenda CD, CE, and CF. The recovery pass itself required a logged follow-up correction – an initial pass mislabeled part of Addendum BN (BN.6 omitted entirely, BN.8's content misnumbered as "BN.4"), caught and corrected via a second conversation_search pass against the same source, with the correction noted in place at BN's own header rather than silently overwritten. This is the first documentation-propagation instance in the project's history whose own repair itself needed a correction, logged accordingly.

New at v20: a tenth instance found and repaired – a recurring file-format bug, caught in two files. SFT_Addenda_Archive.docx was found to have been saved as plain text with a .docx extension rather than a genuine binary Word document – the same specific bug previously found and fixed once before, per the project's own v16 consolidation record, and evidently silently reintroduced at some point since. The same bug was also found, independently, in SFT_Master_Document_v19.docx itself while preparing this v20 consolidation. Both files were rebuilt as genuine binary docx and verified by rendering to PDF before being finalized.

New at v21: an eleventh instance found and repaired – the same file-format bug found at v20, reintroduced a third time. While preparing this consolidation, SFT_Master_Document_v20.docx itself was found to have been saved as plain text with a .docx extension – the identical bug fixed at v16 (first occurrence) and again at v20 (in both the archive and, separately, in what was then SFT_Master_Document_v19.docx). This is the third occurrence of the same specific bug, now in three different files across three consolidations. Rebuilt as genuine binary docx and verified by rendering to PDF before being finalized. No content was lost – the plain-text file was fully readable and its content carried forward intact – but the recurrence itself suggests whatever step is meant to produce

genuine binary output is not reliably being applied, and is worth a dedicated check at the next consolidation rather than treating each instance as independent.

New at v19: an eighth instance found and corrected. Four Part VI section headers — "Substantially resolved this consolidation" for Addenda BM–BQ, BR–BS, BT–BW, and BX–CB — had never been updated after the consolidation that actually resolved them (v16 for the first three, v17 for the fourth), and had propagated unchanged, and uncaught, through both the v17 and v18 consolidations. Two STATUS lines (Postulate 5, Postulate 8) similarly referenced "this consolidation" in a way that had gone stale across the same span. Caught during this consolidation's own header review, prompted directly by the architect's standing instruction to maintain consistency; corrected to name the actual resolving version (**"carried from v16"** / **"carried from v17"**) rather than silently left as found. No content changes resulted — only the version attribution.

New at v18: a second recovery route, distinct from the architect supplying original session material, was demonstrated and used successfully twice. Addendum CD located by1_substrate.npz's exact generating code (RNG seed, construction algorithm) via conversation_search over this project's own past sessions, re-executed it deterministically, and verified the result against that session's own contemporaneous sanity figures — not a recovery of literal cached bytes (no tool can retrieve those; session sandboxes do not persist), but a faithful regeneration from recovered source code, sufficient to upgrade several findings from "one reconstruction" to "the real substrate." Addendum CE applied the identical technique to Addenda BU/BV directly, recovering their exact lattice substrate and coupling-scale convention to a verified match — the first concrete progress on the BT/BU/BV parametrization-recovery item since it was opened at BY.1.

New at v19: the same technique, pushed one step further, closed the item entirely. Addendum CF located BT.4's and BV.3's own literal dynamics/integration code — not only their substrate — via the same route, and used it to resolve the standing generic-IC discrepancy outright (see Part VI). This is now a demonstrated, repeatable methodology for this project's archive-recovery problems generally, including the still-unrepaired BM–BW gap noted above, which remains open and is the natural next application.

New at v24: a twelfth instance of the plain-text-saved-as-.docx bug confirmed and repaired — this time decisively, not as a mount-artifact reading. v23's own text (above) had encountered the same plain-text symptom and, reasoning that project-knowledge search always extracts text regardless of the underlying file's real format, declined to treat it as a fresh corruption. This consolidation checked the actual stored bytes directly instead: both SFT_Master_Document_v23.docx and SFT_Addenda_Archive.docx, copied directly from /mnt/project/ and inspected with Python's zipfile module (a structural, format-agnostic test — not a text-extraction pathway of any kind), failed zipfile.is_zipfile() outright, and their first bytes on disk are literal readable text ("Table of Contents…", "SFT Addenda Archive…"), not the PK zip signature every genuine docx must have. This is a materially different and more direct check than the one v23 correctly declined to act on, and is not subject to the same ambiguity. Both files are confirmed genuinely corrupted as stored. Distinct from, and not to be conflated with, the still-unconfirmed v22 literal-asterisk corruption noted in the version history above, which remains a separate, separately-tracked item pending the architect's confirmation of its cause. No content was lost in either case — both files were fully readable as plain text and their content carried forward intact into this v24 build. Both rebuilt as genuine binary docx and verified, this time, by three independent checks before being finalized: the file command reporting "Microsoft Word 2007+", Python's zipfile.is_zipfile() returning True, and a full PDF render via LibreOffice with visual inspection of the output pages.

Process note, honestly scoped. This is the twelfth occurrence of the identical bug across many files and consolidations (v16, v20 ×2, v21, v23's near-miss, v24), consistent with the

standing observation (first made at v21) that whatever step is meant to guarantee binary output is not reliably applied. This consolidation's contribution to the process fix is narrow but concrete: the verification step now includes an explicit zipfile.is_zipfile() check, not only file and visual PDF inspection, since file's signature-sniffing has not reliably caught every instance in the record. What remains genuinely unaddressed is the actual injection point — whether corruption happens during generation, during delivery, or during the project's own storage/re-upload pathway downstream of delivery — which is outside what any verification performed before delivery can detect or guarantee against. The honest recommendation, not yet acted on: when a delivered master document is re-uploaded to the project's own knowledge store, a quick open-and-check on the architect's own end would catch corruption introduced after the point this project's own verification can see.

5.7 4. Document Structure Response

In response to the broader pattern of critique, the monolithic paper was restructured into a Normative Core and a set of Informative Notes, following ISO/RFC-style normative/informative conventions. This master document folds both back together for a single authoritative reference.

5.8 5. Construction Errors Caught and Corrected Before Being Trusted (Addenda CP, CR) — added at v22

Addendum CP.3: a hedgehog soliton's exact-center site was initially special-cased to U=I, incorrectly — the correct r→0 limit for any finite profile scale is U=−I (since sin(π)=0 kills the direction-dependent term exactly regardless of profile scale). The error silently zeroed out the soliton's core energy cost, producing a spuriously low energy at small profile scales that would have been reported as a false positive had it not been checked before use.

Addendum CR.5: a first attempt to verify the dynamically-selected SU(3) acts complex-linearly on the selected direction's orthogonal complement returned values far from zero (0.34, where ~0 was expected) — traced to comparing quantities computed from two different independently-randomized minimization runs (each landing on a different point on the same degenerate S^6 orbit of equivalent minima), not a physics discrepancy. Corrected by recomputing self-consistently within a single run; all quantities then matched to $\sim 10^{-8}$.

STATUS: Both recorded here per the project's no-silent-corrections standard, not because either affected the final reported results (both were caught before being reported), but because the project's own discipline requires logging the correction itself, not only the corrected outcome.

6 Part VI — Standing Open-Problems and Next-Steps Register

Consolidated across all sessions through Addendum GQ (updated at v43 from "through Addendum GN" following this consolidation's fold-in of Addenda GO–GQ; updated at v42 from "through Addendum GI" following this consolidation's fold-in of Addenda GJ–GN; updated at v41 from "through Addendum GD" following this consolidation's fold-in of Addenda GE–GI; corrected at v28 from the stale "through Addendum DK" marker, itself unnoticed since v26 — logged here rather than silently fixed; updated at v29, v33, v35, v36; corrected again post-v40, documentation-accuracy pass, no version bump, from the stale "through Addendum FN" marker left unaudited since v36 and 12 letters behind after Addenda FX–GD were folded in at v40 without updating this pointer). *Note: this section's own "**through Addendum X**" marker is now current as of v35; the front-matter "**compiled from… through Addendum***

***EV/EL/EG/ED**" markers elsewhere in Part III were not audited this consolidation — flagged for a future pass, not addressed here, per the scope of this consolidation's task*. Full addendum text lives in the companion file SFT_Addenda_Archive.docx, append-only and updated at each consolidation (not after each individual addendum, per the architect's 2026-07-22 correction to prior working practice) – the Addenda BM–BW gap logged at v17 (see Part V §3) was repaired at v20 and remains closed. This register continues to cite addendum letters directly. Ordered roughly by how directly each blocks a currently load-bearing result.

6.1 Highest priority — foundational, directly load-bearing NEW at v33: Overlap admissibility (H_W spectral gap, Addendum H §H.11–H.14) — REOPENED, downgraded from SUBSTANTIALLY RESOLVED to CONJECTURE/OPEN, restored to this register after removal at v3. Flagged at v32 as depending on B.5/B.6's Lieb-Robinson bound, itself dependent on polynomial substrate growth — a premise H.17's own work (Addenda DS–EU) has since shown does not hold. Addendum EX investigated directly: found B.5's own tool choice was likely wrong (the Hernández-Jansen-Lüscher theorem this project's overlap-fermion decision is actually based on needs only bounded degree and a spectral gap, no volume-growth condition), but found a second, independent, previously unflagged polynomial-growth dependency inside H.14's own multiscale induction (an explicit L_k^d volume factor) — so the more direct route's required spectral-gap input remains unestablished on the honest substrate regardless. A literature check confirmed Anderson localization on tree/expander-class graphs (DU/DV/EU's own characterization of the honest substrate) is threshold-gated, not automatic at any disorder strength, unlike finite-dimensional localization. Addendum EY then computed SFT's actual disorder-to-hopping ratio directly — using only quantities already established elsewhere in the theory (Postulate 2's real coupling, Addendum H.13's own disorder definition) on this project's own already-verified honest substrates (grow_hybrid, DS.2) — against a connectivity-matched Anderson/Bethe-lattice critical threshold, calibrated against an independent literature reference point before use. Result: H.11's deterministic Weyl certificate does not cover the honest substrate at any tested configuration (r_eff 5–9× past its own r_eff<1/2 threshold); the probabilistic route's own required disorder threshold is not cleared at any reliably-measured configuration or connectivity, sparse or dense, with the shortfall worsening as connectivity increases. A same-session correction (EY.9) found the scan's one apparent near-miss was a measurement artifact (disconnected fragments contaminating a whole-graph statistic, not restricted to the giant/connected component); the corrected sparse-end scan closes the item and finds the negative result more decisively supported, not less. Honestly scoped as evidential, not forcing: the disorder-to-hopping proxy imperfectly matches the literature's uniform-hopping convention (SFT's own hopping also carries disorder, direction of its effect on the threshold comparison not separately assessed); grow_hybrid and DS.2 are honest but not canonically-adopted substrates, since H.17 itself remains the open question of what the true adjacency rule is; the threshold formula is an accurate large-K approximation, not the exact Abou-Chacra–Thouless–Anderson solution; mean degree was used, not the full degree distribution. Consequence flagged, not yet acted on: per Addendum EX's own pre-stated fork, this is grounds to revisit the Wilson-term-vs-overlap decision (Addendum G), made assuming overlap admissibility was tractable. Does not affect Postulate 8's own CONDITIONAL status (the Route A/Route B fork) — admissibility is a downstream fermion-

late 7 sense) with a physical-completeness argument and converted the unforced sign into a threshold-coupling question ($\gamma\_eff(n)=\gamma-\kappa\cdot n^{T}Mn$), tested directly on both of the project's real substrate constructions: the mechanism is essentially free on orbit substrates (exact rank-5 anisotropy, forced by D's own eigenspace structure, 8/8 seeds) and weak on lattice substrates (near-isotropic, 5/5 seeds) — the first evidenced link between AZ.13 and AW.5. FA then tested AW.5 directly and found the tension is very likely not independently resolvable: any orbit-based substrate — single orbit or multi-orbit ensemble alike — lives on a compact space (S^6 at most, under G2), and shell growth saturates to ~100% of N within log(N) hops regardless of construction (FA.5), the same compactness Addendum Z.15's Poincaré-recurrence result already used, now shown to force H.17's own log-diameter fork-branch directly. FB, testing a third hunch, found a genuine new closed-form identity for D's fixed axis ($\hat{n}=\pm Im(a\cdot b)/|Im(a\cdot b)|$, FB.1 — now folded into Postulate 8's own text, Part II) and confirmed lattice-valued generators give exactly lattice-valued fixed axes (FB.2), but an initially promising growth signal from this route was caught and corrected — pure multiplicative closure from the seven basis vectors is exactly finite (14 points, FB.4), a fourth structurally independent confirmation of the same compact/finite obstruction (joining Addendum AN's discrete-subgroup finiteness, Addendum Z.15's Poincaré recurrence, and FA.5's own graph-distance saturation). Evidential, not a forcing theorem — does not rule out non-compact orbit modifications beyond Addenda Z.16–Z.20's split-octonion route. **Implemented at this pass: explicitly cross-referenced against the H.17/Postulate 1 discretization entry below, this section, rather than left for a future consolidation.** Full relocation into one contiguous entry was considered and set aside: this section's Highest-priority bullets are ordered chronologically by consolidation date across multiple topics (Postulate 3 stabilization and H.17/Postulate 1 discretization interleave throughout), not grouped by topic — moving this bullet would break that convention rather than genuinely merge content. Cross-reference applied instead, per the addendum's own stated fallback ("or at minimum cross-reference each from the other explicitly"). CS.4's own legitimacy gate (may substrate data enter Postulate 7's potential at all) remains untouched and open. No status-tag changes: H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. → See the H.17/Postulate 1 discretization entry, this section, below, for the full adjacency-before-geometry record this synthesis draws on. Next unused addendum letter: FC. See Addenda EZ, FA, FB in full.

NEW at v35 (Addenda FC–FD): the DZ/EA conductance-ceiling thread, standing open since v29 (Addendum EG), is closed as a decisive negative on continuum-limit relevance — while turning up a real, previously-untested independent axis in DS.2's own construction. FC completed all three of EM.14's own open items. d=3 (dimension-parametrized simplex family) was pushed to its own collapse boundary (gb=1.1, earlier than EM.10's reported 1.2–1.3): the δ-crossing still does not complete — flat negative (−0.025→−0.024) — strengthening rather than merely narrowing EM.13's conjecture that the sign-crossing is specific to d=2 sequential gluing. DS.2 was pushed to θ=0.50 (closer to its own percolation edge than EM.12's 0.55): the δ/diam slope does flip sign (+0.015→+0.021 across two seed batches) but noisily ($R^2\approx0.5$). d=4, tested for the first time, shows its own boundary (≈1.0–1.1) and a δ/diam slope moving further negative near it (−0.008→−0.018) — distinguishing d=4 from d=3's flat non-crossing rather than merely replicating it. A fourth, tentative finding (N-dependence of EM.8's own ratio=1 near-boundary noise, CV~N^0.123, R^2=0.567 across N=800–6,400) motivated FD's own follow-up. FD then delivered the thread's decisive result: grow_hybrid's diameter stays flat (9.78→10.33) across its entire stable range at N=3,200, against DS.2's sharp divergence (6.8→18.0) near its percolation threshold — confirming these are two genuinely different mechanisms (a percolation correlation-length effect specific to DS.2, versus a purely spectral/topological effect specific to d=2 sequential gluing) rather than one shared phenomenon, as EM's own framing had left open. A second, independent finding: DS.2's own M-ratio (candidate-triple count per vertex, fixed at 20N since the construction's introduction at Addendum DS itself) is a genuine independent axis, not redundant with θ — tested at matched giant-fraction (≈0.35), M-ratio=10/20/40 give

exponents −0.60/−0.56/−1.24 and δ/diam slopes +0.027/+0.016/+0.051, with the project's own original convention (M=20) producing the *noisiest* fits of the three (R^2=0.46–0.75, versus 0.91–0.96 at M=10 and M=40) — DS.2's entire prior record may have characterized its least-clean parameter choice by accident, logged here rather than silently adjusted. **Correction to FC.4, per the no-silent-corrections standard**: extending the CV/N check with one further point (N=12,800, CV=0.089, a sharp drop) reverses the fit entirely — CV~N^0.123 (R^2=0.567) becomes CV~N^−0.100 (R^2=0.028), essentially no trend across the full 16× range. FC.4's own tentative "suggestive of critical-point behavior" reading is retracted as a directional claim; this does not retract Addendum EM.8's own original near-boundary noise finding, only FC.4's specific extension of it into a claimed N-scaling trend. **Synthesis (FD.4)**: the conductance/δ-hyperbolicity double signature first found at Addendum DU–DV is real but fragile, mechanism-specific to d=2 sequential gluing rather than generic to associator-minimizing constructions, and does not currently constitute evidence for any particular continuum-limit geometry for H.17's substrate. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. See Addenda FC, FD in full.

NEW at v35 (Addenda FE–FH): a four-addendum, single-session arc directly chasing whether any untested construction axis lets H.17's substrate escape the log-diameter default — finding none, while resolving two genuinely new structural questions along the way and catching two of its own errors before they could stand. The session's opening hunch — that a non-compact, algebraically-generated substrate construction, beyond Addenda Z.16–Z.20's already-explored split-octonion route, might escape the compact/finite obstruction FA.7 and FB.5 had just converged on — was checked against Addendum AN.5–7's own general argument before being tested: closed on both branches already, by construction-independent reasoning (compact discrete subgroups are finite by elementary topology; non-compact split-G2 lattices are closed twice over, by Z.16–17's locality-violating cosh-growth and independently by Milnor–Švarc/Gromov's exponential-growth guarantee for simple Lie groups). Addendum FE located the actual untested cell instead: every H.17 adjacency construction to date (Addendum DS's static filter and its descendants) uses Haar-random points on S^6 as its vertex set; Addendum BU's integral-octonion-lattice substrate uses non-metric associator adjacency's own metric cousin (k-NN on ambient lattice distance) — the combination {lattice-ordered vertex set, non-metric algebraic adjacency} had never been tested. Built and run (from-scratch octonion algebra, validated to machine precision; $\mathbb{Z}^7$ lattice balls, N=98→88,272, a 900× range exceeding the project's own 512× convention; direction-degeneracy checked and found negligible, ≈1%): still log-diameter, decisively, 3/3 seeds (R^2=0.994–0.997 vs. 0.977–0.983 for power-law). Addendum FF then tested FE's own proposed mechanism (that DS.2's independent-triple-sampling architecture, not vertex-set compactness, forces log-diameter) directly on the lattice substrate — a full-boundary sequential/correlated triple-selection rule gives statistically indistinguishable results (log R^2=0.979–0.997 across 3 seeds) — and separately found this axis was already closed by the project's own prior record before being re-tested: grow_hybrid's genuine sequential construction (Addendum DS) is *more* decisively log-diameter (R^2=0.9936) than the independent static filter, not less. A sharper, fixed-window frontier-restriction attempt produced a dramatic-looking near-linear diameter blowup, caught before being reported as a positive escape and correctly diagnosed as a pure construction artifact (a global fixed window forces the entire build into a serial chain once N exceeds it, giving diam≈N/window regardless of any algebra — confirmed 2/2 seeds, ratio-to-prediction converging to 0.92–0.99). Addendum FG built the corrected, N-scaled (proportional-fraction) version of that frontier and found a genuine, replicated, mechanistically-explained effect that FF.4's artifact was not: an order-of-magnitude increase in the log-diameter slope, driven by how much of the accumulated structure a new element can see, cross-validated against a simple geometric-decay formula (diam≈ln(N)/ln(1/(1−fraction))) at three independent fraction values with predicted-vs-observed R^2=0.995 — a real result, still within the log-diameter class, not an escape. Ad-

dendum FH then tested whether this slope formula was algebraic at all, before attempting to derive it from the octonion structure as originally asked: a null model with the identical proportional-window architecture but zero octonion content (fixed-probability Bernoulli edge acceptance in place of the associator threshold) reproduces FG's slopes within ordinary seed noise at every fraction tested — DERIVED (negative): the phenomenon is a fact about sequential random-graph construction, not about the octonion algebra, G2 symmetry, or any Postulate 1–4 structure. The correct first-principles derivation, supplied in the null model's place, is graph-theoretic: a worst-case greedy backward-jump argument reproduces FG's formula exactly, and was rigorously bracketed (a proven, algebra-independent lower bound; a Monte Carlo-verified upper bound over eight orders of magnitude in N, R^2=1.00000) — but a follow-up test then found this bracket, while mathematically sound, bounds a quantity (newest-to-oldest pairwise distance) that is a poor proxy for the project's actual multi-source diameter measure (4–9× smaller, sometimes disconnected). Pursuing this further caught a genuine indexing bug in that very test (vertex *labels* were mistaken for temporal build order) — corrected, strengthening rather than reversing the finding (36/36 trials disconnected using the true extremes, versus 4/36 uncorrected) — and led to the arc's most durable positive result: the gap is driven by a precisely characterized, asymmetric *recency boundary layer*, not a scattered low-degree population as first suspected. The most recently introduced ~5–10% of vertices are systematically starved of connectivity (mean degree falls from ≈7.5–8 to 3.80 in the final 5%, with 43.5% at degree ≤2) because the construction simply ends before they can accumulate received edges — confirmed directly against measured eccentricity, not only degree, with a close-to-monotonic relationship between a source's temporal position and its eccentricity. **Overall**: two genuinely different construction axes opened by this session (compact/non-compact substrate structure; independent/correlated edge-sampling architecture) are both closed, on general argument and direct test alike; a third axis (frontier breadth) yields a real but non-escaping effect, now proven algebra-independent; and the mechanism actually responsible for the project's own diameter measure is identified for the first time, with an exact closed form for it left as a sharply-defined, still-open target. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. Next unused addendum letter: FI. See Addenda FE, FF, FG, FH in full.

NEW at v36 (Addenda FI–FN): CS.4's legitimacy gate is directly confronted for the first time from two angles — a precedent audit, and a direct build-and-test of its physical consequences. FI executes Addendum EE.6's own recommended follow-up: a systematic sweep of the addenda history for any other cross-sector potential coupling. None found — the closest candidate (BU.2/EP.4's reuse of Postulate 2/3's own machinery on lattice points) is same-sector under EE.5's own test, and the one genuine cross-postulate attempt (Addendum EF) found the theory's mechanisms actively resist referencing other-sector data in their potentials. CS.3 stands confirmed as the theory's first proposed cross-sector coupling, with nothing to judge it against. FJ–FN then tested the coupling directly rather than continuing to argue about it: CR.3's exact G2→SU(3) pairwise-collapse mechanism, coupled to each real substrate's own second-moment tensor M, survives fully intact at every tested coupling strength on both substrates (FJ), with BD's triad sector confirmed decoupled under genuine joint optimization (FK.1). The previously-flat S^6 direction CR.3 left undetermined is resolved rather than damaged — full determination on the generic lattice substrate; on the orbit substrate, an exactly massless, permanently protected mode at every tested κ, verified via direct Hessian computation (FK.3–FK.4, with one apparent-threshold artifact caught and corrected in place, FK.2). Addendum FL checked and corrected an existing claim along the way: EZ.6's $\hat{n}$/M-eigenvector alignment is real but was mislabeled — $\hat{n}$ sits at M's isolated, non-degenerate middle eigenvalue, not its top one, ruling out a hoped-for connection to the massless mode. Addendum FM then derived the orbit-substrate result from first principles: an exact closed-form theorem (Haar-averaging within any bounded rotation plane of a compact-group generator) predicts M's entire spectrum from D's own eigenspaces alone, verified to 7+ decimal places — identifying the massless mode as a fourth manifestation of the identical G2-compactness

obstruction already documented in Addenda AN, Z.15, and FA.5, with its exact manifestation rate (91.73%) recovered for free from Addendum AF's own decade-old Dirichlet(1/2,2,1) distribution, and shown identical whether the k=1 or k=2 harmonic dominates. Addendum FN closed the remaining methodological question — whether an AG-style spectral-gap conditionality applies — negative, via a direct stress test at 30× tighter neighboring-eigenvalue gap than previously tested: the claim is an exact symmetry of the coupled potential, not a linear approximation of nonlinear dynamics, and so is gap-independent by construction, unlike AG's own conditionality. CS.4's legitimacy gate itself — whether substrate data may enter Postulate 7's potential at all — remains untouched and open; none of this licenses the coupling, only characterizes its consequences if licensed. No formal status-tag changes anywhere in the theory; AZ.13's $\gamma<0$ contingency is unchanged. See Addenda FI, FJ, FK, FL, FM, FN in full.

NEW at v37 (Addenda FO, FP, FQ–FR): CS.4's legitimacy gate is DECIDED, and H.17's own circularity-charge horn is directly revisited and restated for the first time since it was raised. FP records the architect's explicit decision on CS.4 — the v-M coupling tested unconditionally in Addenda FJ–FN is LICENSED, per Addendum EE.2's reading that Postulate 7's own text ("algebraic triads of substrate elements") was always the operative intent, with BD/CR's free-field construction an unreconciled stand-in — the first use of the project's DECIDED status tag. This promotes FJ–FN's results to unconditional theory content and opens a substrate-dependent route to discharging AZ.13's previously-unforced tachyonic sign via Addendum EZ.2's effective-mass mechanism (close to forced on the orbit substrate, still fine-tuned on the lattice substrate), tying AZ.13's resolution directly to this entry's own standing AW.5/H.17 lattice-vs-orbit question — now the single most direct route to sharpening AZ.13 further. FO closes the FE–FH arc's one remaining flagged gap (sequential/correlated edge-selection on a lattice-ordered vertex set): still log-diameter, at 1024× range, the fifth independent construction axis and fifth negative result on this entry's own log-diameter branch. FQ–FR then revisit the entry's OTHER horn directly — the circularity charge against polynomial growth requiring an imposed lattice — untouched since first raised. Building Coxeter's canonical maximal order in the octonions from scratch (reproducing and correcting the historical Kirmse's-mistake closure error), FQ shows growth degree is exactly 8 by theorem for any order, that Postulate 3's own dynamics forces the canonical choice (the non-canonical alternative is exactly triangle-free), and — for the first time in this project's history — that genuine polynomial growth and genuine algebraic selection coexist on a real bulk substrate region. But selection proves exactly centrally symmetric against controls: no spatial localization. FR traces this, after correcting two of its own working hypotheses along the way, to exact invariance under global negation forced by -1's centrality in the octonions — verified for every literal-coefficient candidate the theory admits, surviving two independent chirality-escape attempts, and shown to be a structural feature of every construction paradigm this project has built (compact orbit, non-compact orbit, closure-based lattice, non-metric algebraic-adjacency), each independently failing this entry's full requirement for its own previously-established reason. H.17's circularity horn is RESTATED as a symmetry horn. No formal status-tag change to H.17 itself; see H.17 entry above (Part V) for full detail. See Addenda FO, FP, FQ, FR in full.

NEW AT V44 (Addendum GT): the symmetry-horn restatement above (Addenda FQ–FR) is derived rather than merely observed to recur. Building $G_2 = \mathrm{Der}(\mathbb{O})$ directly from this project's own octonion structure constants, GT proves G_2 has no nonzero fixed vector in $\mathrm{Im}(\mathbb{O})$ and acts transitively on S^6 — so no G_2-invariant construction, of any kind, can select a preferred direction — and that $\Lambda^3(\mathrm{Im}\,\mathbb{O})$ has an exactly 1-dimensional G_2-invariant subspace, supported on the Fano lines (φ), fully exhausting the invariant content available to any coupling built from up-to-trilinear octonion data. This unifies Addendum FR, FV.7–FV.9, GL.4's near-tautological finding, and Addendum GS's cone/M-tensor result as one fact viewed four times, and supplies a predictive filter: any future proposed mechanism built as a G_2-invariant or fully basis-symmetric function of structure-constant data can now be ruled out without

simulation. No formal status-tag changes. See Addendum GT in full.

NEW at v38 (Addenda FS–FU): a genuinely new diagnostic lens — holographic tensor-network structure — applied to this entry for the first time, and pursued to a decisive negative, then twice qualified further. FS ran a Ryu–Takayanagi min-cut scaling test (validated against tree and random-4-regular controls) on the project's established H.17 substrate: linear-dominant, not the logarithmic signature genuine holographic structure requires, despite real local hyperbolic signatures already on record (Addenda DU–DV, EM). A constructive response enforcing a hard Q=7 vertex-valence cap ($\kappa=-1/6$, the minimal hyperbolic {3,Q} triangulation) produces this project's first genuine connected, negatively-curved, proportionally-growing consistent core — but the full-graph RT result is unmoved, and every one of three independent defect-suppression levers tested (search width, soft-cap tolerance, load-balancing) either converges toward the identical finite-closure phenomenon Addenda AN and DS.6 already established by unrelated methods (a third structurally independent route to the same trap) or is actively counterproductive. FT closes the one hypothesis this left open — that the core itself is holographic but full-graph noise conceals it (refuted directly: the isolated core is, if anything, more linear-dominant than the full graph) — then implements the structurally correct constraint FS's own result had identified as missing (genuine edge-level {3,Q}-manifold consistency, not merely vertex-valence capping): true consistency under this corrected definition is two to three orders of magnitude rarer than FS's vertex-only core at fixed search width, but sustains a genuinely stable, non-vanishing density under proportional search across a full 32× range in N — the first stable consistency density this entry has ever produced. This is immediately and substantially qualified: the resulting true-consistent vertices form no connected structure at all and show zero spatial clustering under a rigorous statistical test against a random-placement baseline, confirmed across three separate density regimes spanning a 25× range — local consistency and spatial organization are, on this evidence, fully decoupled properties of the construction. FU tests, rather than merely proposes, the one remaining candidate explanation FT's own closing discussion named but did not test — a complete absence of any feedback channel between spatially separated closure events — by building an explicit feedback-biased variant and sweeping its strength directly: the result is a sharp collapse threshold, not a rescue, with no feedback strength found that both preserves growth and produces clustering. This constitutes a seventh independent, structurally distinct manifestation of this entry's central growth-vs-global-consistency tension, now documented across AN, DS.6, FS's three suppression levers, FT's proportional-search thread, and FU's engineered feedback test. NEW AT V39 (Addenda FV–FW): an eighth and ninth independent manifestation, tested via a different route — whether Fano-plane-collinearity structure in the octonion algebra itself can supply genuine locality. FV builds the project's first positive-cone semigroup (N^8): genuinely escapes FR's negation-cancellation mechanism at the correct polynomial growth degree (8), but three symmetric associator-based couplings show null localization, generalizing into a conjecture that any coupling invariant under the octonion's structure constants inherits some discrete symmetry of them. A genuinely position-dependent coupling — Fano-plane collinearity of a point's active support — breaks this: real, mechanistically-forced (associativity of the generated quaternion subalgebra), dynamically-robust selection (0/70 collinear-support survivors, 5 seeds × 4 regimes), but confined to the cone's boundary (support ≤4), since the bulk (support ≥5) is provably flat by Fano design regularity — every point lies on exactly 3 of 7 lines. FW answers the resulting open question — is boundary-localized selection genuine H.17 locality, or a cone-boundary artifact? — with three exact theorems rather than further examples: the boundary signal can never carry more than a single discrete bit regardless of magnitude, by trilinearity of the associator; the bulk wall is exact, provable via inclusion–exclusion on the Fano plane's own 2-design axioms (3 lines/point, 1 line/pair), and generalizes to any Steiner triple system, not specifically to octonions; and any coupling built purely from the collinearity data is invariant under exactly PSL(2,7) (order 168) and cannot be pushed to a smaller symmetry group. A richer, coherent vector-valued coupling does escape the exact bulk degeneracy, but regression shows that

escape is ~99.7% explained by generic pairwise vector-angle proximity — a property of any alternative algebra — rather than Fano-specific structure, an honest negative rather than a route forward. For the entire class of couplings built as closed-form functions of the structure constants, no route to genuine bulk locality is found, now for well-characterized reasons in each case; what remains untested is real Postulate-3 dynamics on one concrete random realization, which carries extra non-invariant data no closed-form coupling has, and is plausibly why FS–FU's actual simulations found a genuine if swamped hyperbolic core. No formal status-tag change; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. See Addenda FV, FW in full. NEW at v40 (Addenda FX–FY): FX tests whether real Postulate-3 dynamics, run on one concrete random realization rather than a closed-form invariant coupling, can escape the exact degeneracy FV/FW proved for the support≥5 bulk — it does (a genuine, converged, IC-independent PARTIAL attractor, replicated across 10 independent substrate realizations), but tracing the mechanism (local triangle-associator mass predicts survival, r=+0.142, 10/10 realizations) finds it reduces to the same generic vector-angle geometry FW.9 already identified for static couplings (R^2=0.726 from cosine alone), extending that honest negative to the dynamical case. FY directly tests the positive-cone semigroup itself (Addendum FV) as a candidate answer to Addendum AJ.6's own discretization question: a genuine bounded-degree, magnitude-corrected adjacency was built on it for the first time (two earlier attempts — raw magnitude coupling, and a properly degree-capped version — failed via hub concentration and fragmentation respectively, both traced to a magnitude bias in the raw octonion product), and its diameter growth remains log-diameter class with an elevated constant, paralleling Addenda AQ→AS→AT's own earlier resolved case (steeper slope, not growth-class escape) at far smaller N-range (32× vs. AS/AT's 512×). This is an eighth structurally independent construction family converging on the same wall. AJ.6 and H.17 are formally stated as one problem, not two, per AN.3–4's original argument: any discretization of Postulate 1's substrate must support a well-defined large-N growth class, which is exactly H.17's fork. No formal status-tag changes. **NEW at v41 (Addenda GE–GI): AZ.13's link to H.17/AW.5 tested directly at the construction level for the first time, and a genuinely new construction family (non-abelian nilpotent groups) found and characterized.** GE built a quaternionic (Fano-line) sublattice, achieving genuine polynomial growth together with AZ.13's discharge mechanism at its exact maximal, mechanistically-forced value (7/7 Fano lines) — a real decoupling of AZ.13 from H.17/AW.5's compactness wall — but found this specific construction fails Postulate 3 on three independent grounds (identically-zero triangle term, since a Fano line is a closed associative subalgebra; a sharp extinct/blowup transition under pairwise dynamics with no broad partial regime; and the same forced-centrosymmetric-selection obstruction as any negation-closed additive lattice). GF converted two pieces of the eight-family pattern from case-by-case testing into general theorems: a universal bounded-degree diameter floor (any bounded-degree graph has diameter at least log(N)/log(k−1), explaining why log-diameter is the default rather than a recurring mystery), and a negation-closure theorem showing any additive-lattice discretization, under any admissible magnitude-type coupling, is forced into centrosymmetric Postulate-3 survivor sets — generalizing Addendum FR's single-lattice finding to the entire family, verified independently on GE's own construction. This left one door open (Gromov's theorem characterizes growth as virtually nilpotent, not merely abelian; the theorem's argument does not obviously extend to non-abelian nilpotent groups). Testing it directly, GF built the discrete Heisenberg group from scratch (verified nilpotency class 2, growth converging to the exact Bass–Guivarc'h prediction of 4 at three independent scales) and found the first construction in this project's history to simultaneously clear every obstruction tested to that point — genuine growth, escape from the negation-forcing theorem, and genuinely nonzero associators — then confirmed this is a general property of non-abelian nilpotent constructions, not a Heisenberg-specific accident, via a second, structurally independent class-3 construction (UT(4,Z), built by literal matrix multiplication, growth converging toward its own theoretical prediction of 10, negation-closure failure even stronger than Heisenberg's). This construction nonetheless failed the theory's actual purpose on a third, previously uncatalogued ground: its Postulate-3 dynamics proved almost entirely topology-

determined rather than genuinely competitive at the parameter range first tested (99.86% agreement with a static zero-dynamics predictor) — a new failure mode distinct from compactness, the parity obstruction, or associativity collapse. GG found this was an artifact of the tested parameter range rather than a structural property: a fine h-scan (matching Addendum ES's own precedent on an unrelated substrate) located a genuine, spatially coherent, initial-condition-sensitive nucleation domain in a narrower window, the first real competitive selection found on any growth-achieving construction in this project's history. GH proposed a mechanism (negation-orphan status driving the nucleation domain) that GI, testing at a larger radius matching Addendum GF's own growth-verification scale, found did not survive proper scale control and retracted explicitly; GI identified the actual mechanism instead — ambient scale driving triangle-associator mass via the associator's own degree-3 homogeneity, producing a single-parameter critical transition (monotonic core/dead fractions, unimodal swing fraction) — and confirmed this mechanism replicates on the independent class-3 construction to three significant figures (r=0.7173 vs 0.7169). Net position: the wall is not proven closed (a construction exists clearing every obstruction tested prior to GF) and not proven open (that construction's own Postulate-3 dynamics, now properly characterized rather than assumed, needed two further sessions to locate genuine competition at all, and its mechanism — while general across two independent constructions — is not yet shown to scale to the substrate sizes the theory's other results depend on). No formal status-tag changes to H.17 or AJ.6; AZ.13's own reading is updated with a caveat rather than resolved (decoupling from H.17/AW.5 is confirmed real in principle, but no substrate found to date achieves it while remaining otherwise viable for the full postulate stack). See Addenda GE, GF, GG, GH (superseded in part by GI), GI in full. UPDATE AT V42, not previously cross-referenced here: GI's scale caveat is substantially addressed.** Addendum GJ (folded in at v42, filed under the Postulate-3 tie-back thread and never linked back to this entry) ran a seven-scale study of the same Heisenberg construction — radius 9–28, N=2,845→261,815, a 92× range exceeding most of the project's working substrate sizes — and confirmed the GG/GI nucleation mechanism is a genuine bulk phenomenon (GJ.1–7), not a finite-size artifact, with an initially-apparent radius-16 outlier resolved as a scan-depth artifact once corrected. GI's "not yet shown to scale" caveat is therefore no longer accurate as stated; what remains open is narrower: (a) whether the mechanism generalizes beyond the two tested nilpotent constructions (Heisenberg, UT(4,Z)), and (b) whether it composes with AZ.13's discharge condition at this same scale — GL composed the two successfully, but only at the smaller GK/GL-scale substrate, not against GJ's own 92×-range study. Net position revised accordingly: the wall is not proven closed and not proven open, but the specific gap blocking closure is now scale-generality, not scale itself. See Addendum GJ.1–7. Next unused addendum letter: GR [corrected at v43 from a stale "GJ", itself left unaudited when Addenda GJ–GQ were folded in at v42–v43 without updating this pointer — logged here rather than silently fixed, per the project's own standard].

NEW AT V44 (Addenda GR, GS, folding in the provisional update above): both items raised previously are now answered. (a) Generalization beyond Heisenberg/UT(4,Z): DERIVED — Addendum GR built $H_3(Z)$, dimension exactly 7, embedding natively into all seven octonion imaginary directions with no coordinate compression; the GG/GI nucleation mechanism replicates in full. A genuine substrate-type distinction was found and, after two rounds of self-correction logged in place (GR.19, GR.22), confirmed rather than explained: across six independent constructions now tested ($H_1(Z)$, UT(4,Z), $H_2(Z)$, $H_3(Z)$, Filiform-4, $\mathbb{Z}^3$), only $H_1(Z)$ plateaus — every other construction decays, with no variable yet isolated that accounts for the difference. (b) AZ.13's discharge condition composes cleanly on $H_3(Z)$ at its own scale: DERIVED, via the full 35-parameter joint dynamics, with a genuinely non-degenerate M-tensor (full rank 7/7, anisotropy 0.386) — see the AZ.13 entry above (Addendum GU) for how this third data point now fits the corrected picture. Addendum GS independently tested four external hunches proposing mechanisms to escape this entry's core obstruction (correlation-structure locality, dynamically-selected sector growth, cone/M-tensor composition, and a non-associativity

bounding claim) — all four negative, each via an independent method with explicit finite-size or artifact controls, extending the honest-branch negative case rather than opening a new route. No formal status-tag changes to H.17, AJ.6, or AZ.13. See Addenda GR, GS in full. Two open items remain: why $H_1(Z)$ specifically plateaus, and whether that plateau survives at N beyond GJ's tested range (radius 20 unreconciled per GR.23, radius 24+ untouched). Next unused addendum letter: GV.

NEW AT V45 (Addenda GV, GW, GX, GY, GZ, HA, HB, HC): both open items from the entry above are now addressed, one substantially, one completely. GV checked whether the Bombelli–Henson–Sorkin no-go theorem for discrete Lorentz invariance bears on this entry's own second horn; it does not apply directly (no presupposed continuum here), but identifies a second, symmetry-based obstruction axis alongside GT's own counting-based one — currently untestable, since nothing has escaped GF.2's counting floor to reach a stage where it would bite. A redirection toward order-relation primitives this suggested was tested directly and found not currently actionable (the algebraic route is already closed twice in this project's own record; causal set theory's own manifoldlikeness problem shows the redirection would not evade the tension regardless). GW and GX together close off ratio-based and class-based explanations of GR.24's own open question decisively: center-dimension and derived-subalgebra ratios both fail as discriminators (GW), and a purpose-built construction matching $H_1(\mathbb{Z})$'s nilpotency class exactly while exceeding its own best ratio still decays (GX) — extending the tested-construction count to seven, $H_1(\mathbb{Z})$ still the sole plateau, with dimension, class, and every tested ratio now excluded as candidate explanations. Why $H_1(\mathbb{Z})$ specifically plateaus (item (a) above) remains genuinely open. GY and HC together resolve item (b): whether the plateau survives at N beyond GJ's tested range. GY reaches radius 20 and 24 and finds the trend still declining, below GJ's claimed floor, leaving the question open with the evidence leaning against GJ's specific claim. HC completes the reconstruction at radius 28 (N=261,815, exactly GJ's own largest documented scale, after — at approximately 121 continuous hours — the longest single computation in this project's history) and settles it: the local decay exponent between successive large radii collapses monotonically ($0.462 \to 0.428 \to 0.324 \to 0.174 \to 0.055$), the actual mathematical signature of an approaching asymptote rather than continued decay. $H_1(\mathbb{Z})$ genuinely plateaus — GJ's qualitative claim is UPHELD — but at a materially lower value than GJ originally reported (an estimated 1.5–1.7%, against GJ's claimed 2.05–2.39%), which explains rather than merely disputes GR.23's own partially-reconciled radius-20 discrepancy: GJ's own original value was itself somewhat too high, not an artifact of seed count or search methodology alone. GY's own interim speculative reading (that $H_1(\mathbb{Z})$ might be merely the slowest-decaying construction, not a qualitative exception) does not survive HC's completed result and is retracted here — the exponent-collapse signature is categorically different from ordinary slow decay. Net position: this entry's central question is now sharper than at any point in its history, not resolved. The plateau is empirically settled; the seven-construction survey (GW, GX) has excluded every structural-invariant hypothesis tested so far; no mechanism explaining $H_1(\mathbb{Z})$'s uniqueness has been found. Separately, three further constructions were tested against this entry across two parallel sessions, all negative, each independently reinforcing the entry's central pattern: GZ's dynamically-gated growth admission is asymptotically self-defeating via k-NN coupling saturation on a compact manifold (a further manifestation of AX.10's own growth-vs-selection trade-off); HA's nonlocal Jordan–Wigner fermionization test, after two further self-corrections logged in place, finds $H_1(\mathbb{Z})$'s apparent locality advantage is an ordinary growth-degree effect rather than a special construction-pipeline interaction; HB's BFSS Matrix Theory test (motivated by the division-algebra classification putting the octonion case among only four viable dimensions) finds no genuine dimensional collapse on the tested construction, an initially exciting signal correctly retracted via calibration against synthetic data of known dimension before being trusted. None of these three bears on $H_1(\mathbb{Z})$'s own plateau mechanism directly; each closes one further candidate escape route from this entry's central obstruction. No formal status-tag changes; this entry remains CONJECTURE/OPEN, PERMANENT ENTRY. See

Addenda GV, GW, GX, GY, GZ, HA, HB, HC in full. Next unused addendum letter: HD.

NEW AT V46 (Addenda HD, HE, HF): item (a) above — why $H_1(\mathbb{Z})$ specifically plateaus — is now the most thoroughly negative-mapped open question in this project's history, and, separately, all four of an external literature survey's proposed H.17-resolution families are closed. HD extends the structural sweep to eight constructions and six candidate discriminators. It completes the natural 2×2 design GX began (class-fixed/rank-varied vs. rank-fixed/class-varied) with a new construction, F(2,3), which decays cleanly across a 72× N-range on the architect's own external hardware (radii 9–10 attempted but genuinely impractical within available time, reported as an honest hardware ceiling rather than extrapolated). It redesigns the M-tensor-anisotropy diagnostic (GW.7) to remove an ambient zero-padding artifact that had forced every sub-7-dimensional construction to read identically (anisotropy≡1.0) — the corrected version still finds no discrimination, and a follow-up test (HD.3) identifies the mechanism precisely: raw anisotropy is forced toward 1 with radius for any non-abelian construction as a consequence of mixed-degree coordinate scaling, so the metric was only ever tracking abelian-vs-non-abelian status, already known not to split plateau from decay — voiding it as a candidate more decisively than a flat negative would. Two dynamics-level hypotheses, tested for the first time rather than only structural ones (h_ijk associator-coupling distribution width; h_scale coverage), are both falsified — the second only after an initial sandbox pass was confounded by mismatched seed counts and correctly left unreported until resolved by a seed-matched external rerun, which finds $H_1(\mathbb{Z})$ has the *lowest* coverage of the three constructions compared, the opposite of the tested hypothesis. Item (a) remains genuinely open — no formal status-tag change — but the negative space around it is now unusually complete: dimension, class, rank, center ratio, derived-subalgebra ratio, and M-tensor degeneracy (in both its original and corrected forms) are all excluded, with a general mechanism (not just an empirical failure) identified for why M-tensor anisotropy specifically could never have discriminated. Separately, HE and HF close a bounded, independently-motivated program: an external literature-survey session (2026-08-21) had proposed four candidate H.17-resolution families as though all four were fresh open leads. HE's own direct project-knowledge search finds this framing substantially wrong before testing anything — Family A (causal sets) was already closed by Addendum GV and Family B (Trugenberger/Kelly-style symmetry breaking) by Addenda K–N/L, both weeks before the external survey ran, unknown to it since it worked from web search alone with no access to this project's own record — a concrete, caught instance of this project's own recurring working-memory failure mode, logged rather than repeated. Family D (entanglement-derived adjacency) is then read from primary sources (Van Raamsdonk arXiv:1005.3035 in full; the Swingle/MERA line; the ER=EPR counter-example paper) rather than the survey's secondhand summary, and tested in two parts directly on the real octonion substrate: a clean exponential correlation-vs-distance relationship is found when a k-NN graph is already presupposed (HE.2), but this is caught before being reported as evidence, since it is close to tautological on an already-adjacent graph; with the presupposed graph genuinely removed and every point-triple made a valid coupling triangle, correlation vs. true S^6 geodesic distance comes back indistinguishable from noise at every tested time point (HE.3) — all-to-all coupling saturates the system into one global collective mode before any geometric signal can differentiate it. Family D does not escape H.17's circularity: real content survives only where geometric structure is presupposed by hand. HF closes the remaining family, C (on-shell symmetry emerging from dynamics), by structural argument rather than simulation — the same precedent Addendum AN set for the discrete-automorphism route — after reading its actual primary source in full (Jiménez & Tsujikawa, arXiv:2608.02458) rather than relying on the survey's summary. Two independent grounds: a category-level type mismatch (the mechanism operates entirely inside an already-existing continuum FLRW spacetime, presupposing exactly the kind of prior geometric existence H.17 asks about, not a mechanism applicable to it), and, even charitably reframed onto graph-level observables, no content beyond what SFT has already tested (Addendum AX.10's presupposed-structure

substrates; Addenda CR/BC/FJ's already-tested spontaneous-symmetry-breaking mechanism class, for the distinct but structurally related AZ.13/Postulate-7 direction-selection question). With HE and HF, all four of the external survey's proposed H.17-resolution families are now closed — two (A, B) were already closed internally before the survey ran; the other two (C, D) have now been tested or scoped as thoroughly as their own source literature permits. This closes a specific, bounded program, not H.17 itself or this project's own much larger internal investigative history (Addenda K through at least FK and beyond), most of which remained outside this session's working memory in full going in — all three addenda independently recommend a dedicated audit pass to build one authoritative, complete map of that history, a recommendation executed directly following this consolidation (see SFT_H17_Investigation_Audit_v1.md, 2026-08-26). No formal status-tag changes; this entry remains CONJECTURE/OPEN, PERMANENT ENTRY. See Addenda HD, HE, HF in full. Next unused addendum letter: HG.

NEW AT V48 (Addenda HG, HJ): the project's first genuine attempt at FR's own flagged "fifth, unconstructed paradigm" is built, tested, and closed negative — sharpening rather than resolving this entry's central obstruction, and confirming the dynamical-rewiring route (Addenda K–N) remains correctly closed. HG uses the associator generatively for the first time — as a literal displacement direction for a growing walk — rather than as a static filter, motivated by Addendum DN's uniqueness result for the associator among G2-equivariant vector-valued trilinear maps. The natural linearizations of this idea are proven exactly skew-symmetric, a fifth independent confirmation of this entry's own compactness/rotation obstruction (joining AN, Z.15, FA.5, FB.5); a normalized-step variant escapes that specific trap and is checked, not assumed, to be genuinely negation-odd per realization — the first checked (not hand-waved) instance of ensemble-symmetric/realization-asymmetric vertex-set structure, distinct from Addendum FR's own closed mechanism, which concerned coupling formulas on already-symmetric vertex sets. A dynamical collapse onto a lower-dimensional subspace is caught and fixed before use. Once fixed, the resulting vertex set still gives log-diameter under both independent-triple-sampling and sequential/correlated adjacency — an eighth structurally independent construction landing in this class — sharpening Addendum FE.10a's own hypothesis from "possibly the sampling architecture" to "most likely the associator-norm-threshold adjacency rule itself," which appears to force quasi-random/expander-like connectivity on essentially any sufficiently generic point set regardless of provenance. HJ then checks HG's own two constructions against a diagnostic never applied to them: Addendum DU's finding that conductance and diameter are independent properties, with DS.7/DS.8's sequential gluing showing a genuine shrinking bottleneck despite an identical log-diameter verdict to DS.2's flat-expander case. Before running this test, HJ confirmed directly that the Trugenberger/Kelly dynamical-rewiring route (Addenda K–N) remains closed for good reason — no forcing mechanism for graph-selection dynamics exists in Postulates 1–4 without promoting Postulate 2's coupling matrix to a ninth postulate — and was correctly not re-opened. The conductance test itself (pipeline validated against DU's own published control figures before use) finds HG's independent-sampling construction reproduces DS.2's flat-expander signature almost exactly on a completely different vertex-generation mechanism, while its sequential construction shows a real but much weaker decay than DS.7/DS.8's own clean result — sharpening DU.2's own finding rather than merely failing to reproduce it: the earlier bottleneck-preserving property looks specific to Addendum AW's lattice-ordered vertex set, not a general property of sequential-vs-independent adjacency architecture. This is now a ninth structurally independent construction sharing both this entry's log-diameter default and its flat-conductance signature. No formal status-tag changes; this entry remains CONJECTURE/OPEN, PERMANENT ENTRY. The live implication for future attempts, sharpened by this consolidation: a genuinely different adjacency rule — not another vertex-generation scheme, and not another sampling architecture layered on the same associator-norm-threshold rule — is what the next real attempt on this entry needs. See Addenda HG, HJ in full. Next unused addendum

letter: HK.

Postulate 3 stabilization: specific form underdetermined (Addenda O–U, tracing back to J.9–J.12). The triangle-term candidate remains the best-characterized of the four live options (R/S/T plus the triangle term), and is now the most fully derived: existence and locality (BH), coupling strength in pattern and magnitude including its natural angular scale from first principles (BI–BJ, BM), robustness (BK), external audit (BL), confirmation that R/S/T do not share its locality property (BN), and a generator-selection rule validated out-of-sample for the network it was derived on (BO–BQ) but shown NOT to generalize across topology instances (BR–BS). The mechanism BEHIND the triangle term's locality is understood (BX–CB) — overlap concentration among shared triangle-nodes, degree-controlled and cross-validated on toy substrates and on H.17's own real substrate. The triangle term's own generic-IC behavior, previously an unreproduced and doubtful claim, is fully reproduced and mechanistically explained (Addendum CF).

New at v21 (Addenda CG–CN): the evidential asymmetry between the triangle term and R/S/T is substantially sharpened, on four independent lines, though the four-candidate underdetermination itself remains formally open. (1) The single-site-invariant derivation route for R/S/T's free parameters (μ, s, A_c) is closed by a general G_2 orbit-transitivity theorem (CG), strictly generalizing Addendum AR's earlier associator-squared-specific finding to every possible construction of that type. (2) R/S/T's non-locality (Addendum BN) is confirmed directly on the theory's own real H.17 substrate, not only toy networks, with long-time convergence verified rather than assumed after an initially-promising false lead was caught and corrected in-session (CH). (3) The reason no derivation route — single-site or global-aggregate — has ever worked for R/S/T is now understood as one structural fact rather than three unexplained negative results: unlike the triangle term's own defining equation, which already contains a genuine Postulate-1 geometric object (the associator) for a constant to normalize against, R/S/T's defining equations contain no Postulate-1 content anywhere in their functional form (CI). (4) Separately, triangle-closure itself was tested against the weaker open "wedge" motif and found NOT required for escaping Addendum AY's no-go theorem — only the degree-matching relationship between driving and restoring terms matters for bare stabilization (CJ) — but triangle-closure DOES produce a substantially larger, more robust, more easily-discoverable local-attractor basin than the open-wedge alternative, a ~5–35× discoverability gap confirmed across three independently-seeded Watts-Strogatz graphs and replicated in direction (though smaller in magnitude, ≈5×) on the theory's own real octonion-orbit substrate (CK–CN, correcting an intermediate overclaim at CK.5 along the way). None of CG–CN changes Postulate 3's formal status or proves the theory requires the triangle term over R/S/T; the four-candidate underdetermination remains this entry's core open question, unchanged in kind, now with a substantially more detailed and more honestly-scoped evidential record behind the asymmetry.

New this consolidation (Addenda CW–CZ): a stable-spiral dynamical signature is characterized for the first time and found to be generic, not distinguishing — a negative result for the hope it might sharpen the four-candidate question further, but a genuine new piece of understanding about the dynamics itself. Following up Addendum O's corrected signature for genuine decaying-oscillation behavior (real $\kappa < -\eta^2/4$ at an otherwise-stable fixed point), Addendum CW found it present in ≈99% of a large stable-fixed-point sample on the triangle term's original toy case; Addendum CX–CY found complete (25/25) coverage on two independent real-network attractors and on the theory's actual N=3,000 octonion-orbit substrate (3000/3000) under the associator-norm; Addendum CZ found the identical complete coverage under φ on the same real substrate and on R's own genuinely-stable quintic branch (25/25) — six fixed points, three mechanisms, toy and real substrates, complete or near-complete coverage throughout. CZ.1 logs a correction to CY.8's own framing along the way (AY's theorem does not foreclose R/S/T's stability, as CY.8

had claimed; R/S/T's disqualification is on postulate-patching grounds, not instability). Honest reading (CZ.6): this looks like a property of η=0.8 second-order dynamics generally, not a lever for adjudicating among Postulate 3's candidates. Not yet tested: S and T directly (though likely by analogy to R); the signature's behavior away from η=0.8; whether any stable fixed point in the theory's history fails it (the genuinely informative case, not yet found).

New this consolidation (Addenda DA–DI): the literal-coefficient choice between φ and the associator-norm was pushed on five independent angles, narrowing the viable class and closing (for now) the differentiability angle as a discriminator between them, though not as a physical phenomenon in its own right. φ is now provably the unique G2-invariant alternating trilinear form on Im(𝕆) (DB.3–DB.5, a classification theorem, not a search result). Signed φ fails the theory's own generic-IC protocol outright while |φ| succeeds once calibrated (DD) — the viable literal-coefficient class is now {associator-norm, |φ|}, not {associator-norm, φ}. DB.5's theorem does not reach the associator-norm (different representation spaces, Sym^3 vs Λ^3 — DC), and the "economy of structure" argument for φ's reuse elsewhere in the theory mostly restates DB.5 rather than adding independent evidence (DE) — both downgraded on their own follow-up scrutiny, not silently dropped. Both candidates' confirmed non-differentiability was shown to have genuine dynamical consequences via exact zero-crossing tests (DH–DI, signal-to-background ratios up to ~59,000x), reopening a question Addendum DA had first found inert — but extensive multi-substrate testing (DG: 14 paired proximity comparisons, p=0.013 in the opposite direction from an initial single-substrate finding; DI: 22-substrate severity comparison, p=0.351) found no discriminating advantage for either candidate. The four-candidate underdetermination remains formally open; the honest state after DA–DI is a more precisely characterized two-candidate residual choice with no forcing argument yet for either.

New at v27 (Addenda DL–DR): the two-candidate residual is extended to three, with the new candidate substantially advantaged and the advantage now analytically explained rather than merely observed. DL–DM first closed out v25's own flagged follow-ups (kink propagation confirmed on the real N=3,000 substrate; the residual "natural near-degeneracy" puzzle in |φ|'s survivor set fully traced to a mundane triangle-count/degree confound, not new physics). DN proved the associator is independently, uniquely forced as the G2-equivariant vector-valued trilinear map on Im(𝕆) (a companion to DB.5's scalar-case theorem), and closed a proposed patch-argument against |φ| as a false lead — both |φ| and the associator-norm are the same canonical norm operation, applied to different-rank objects. DO–DP found and stress-tested φ^2: smooth everywhere (sidestepping DA–DI's differentiability question by construction), non-negative by construction, and wider in discoverability window than both established candidates on four independent substrates spanning a 10× range in N, with the advantage shown to exceed seed-to-seed noise by roughly an order of magnitude — the most heavily cross-checked discoverability claim in the project's history for this class of question. DQ–DR supplied the analytic mechanism: a single-triangle self-consistency threshold, first calibrated to two empirical percentiles and validated out-of-sample (DQ), then superseded by a zero-fitted-constant cascade theory matching the true nonlinear dynamics to within a few percentage points and, at the node level, to 3–4 significant figures in amplitude (DR) — tracing φ^2's advantage to its higher statistical skewness, itself a consequence of squaring an approximately sign-symmetric quantity. None of DL–DR forces a selection or changes Postulate 3's formal status; the four-candidate-turned-three-candidate underdetermination remains formally open, now with the most thoroughly characterized evidential asymmetry of any pairing in the theory's history favoring one candidate (φ^2) without proving it.

[Cross-referenced at v34 from the AW.5↔H.17 synthesis entry above, this section (Addenda EZ–FB): AW.5's lattice-vs-orbit tension is now evidentially linked to this entry via AZ.13 — see above for the full synthesis.] Postulate 1's discretization question / H.17 adjacency-before-geometry, merged with the Postulate 3 triangle-term mechanism thread at v17 (Addenda Z–AJ, AN, AV–AX, BT–BW, BX–CB), continued at

v18 (CC–CE), and substantially advanced this consolidation (CF). With the free-algebra reading eliminated (AJ) and the discrete-subgroup route closed (AN), two genuine lattice discretizations exist (AV, AW) in a demonstrated tension with the bounded-orbit picture used throughout the theory's dynamics (AW.5), and the orbit side's own growth trades off against selection when connected to the real coupling kernel (AX.10). BX–BZ established and confirmed the overlap-concentration mechanism, degree-controlled, on toy and real substrates, and found BW.6's coherence "degradation" is very likely a global-statistic artifact. CC–CE advanced the per-seed nucleation-variance question (a new predictor, R_i, tested and confirmed on the real substrate) and made verified progress on substrate/coupling-scale recovery for BT/BU/BV.

New at v23 (Addendum CV): the negative case on the honest pregeometric branch is extended from the bare pairwise mechanism to the fully-characterized triangle-term coupling class — a backfill correction, this item was never entered into this register's own narrative at the time. Following Addendum CU's full characterization of the non-patching literal-coefficient class, CV tested whether the triangle-term mechanism — BT–BW's own success case, immediately above — can reach H.17's honest pregeometric substrate: genuinely random regular graphs with no presupposed adjacency. Triangle count there is O(1) as N→∞, a classical configuration-model fact verified across a 64× range in N, so the mechanism's reach vanishes as a fraction of the substrate and cannot affect asymptotic diameter-scaling, independent of dynamics calibration. Combined with BT–BW's own confirmation that the mechanism succeeds only on presupposed-structure (lattice/full-orbit) substrates, every dynamical selection mechanism tested to date — bare pairwise (AS, AX) and triangle-term (BT–BW, CV) — shows the identical pattern: works only where geometric structure is already assumed, fails or cannot engage on the honest branch. No formal status-tag change; the negative case is now comprehensive rather than partial.

This consolidation (Addendum CF) closes the standing BT/BU/BV generic-IC discrepancy outright — the single item this entry has carried as its highest-leverage open task since Addendum BY.1. Following CE.6's recommendation, conversation search located BT.4's and BV.3's own literal dynamics code, not merely their prose description. The cause was a genuine initial-condition convention mismatch: every reconstruction since BY.1, including CC–CE in this same session, used generic ICs drawn from Uniform(−0.3, 0.3) by inherited convention; BT/BU/BV's own code draws from Uniform(0.3, 0.8) — substantial, uniformly positive amplitude. Reproducing BT.4's and BV.3's exact configurations with the recovered code and the correct IC gives genuine partial selection from generic ICs: qualitatively on a reconstructed orbit substrate (4/4 seeds, matching BT.4's outcome class), and near-exactly on the verified-recovered lattice substrate (5/5 seeds, matching BV.3's own reported node counts almost precisely at two of five tested h_scale points). A clean two-factor isolation test (CF.4) found the IC convention alone is necessary and sufficient — swapping only the IC on otherwise-identical first-order dynamics reproduces the full second-order result to the exact node; swapping only to second-order dynamics with the old IC still gives complete extinction. **RESOLVED, not merely narrowed**: not a property of the coupling, substrate, octonion algebra, or dynamics order — all correctly cleared of blame across five addenda (CA–CE) of careful, honestly-reported elimination.

At v28 (Addenda DS–ED, twelve addenda — the most extensive single-thread investigation of this entry to date). Now fully backfilled into this register's own narrative, superseding the flag-only placeholder first entered at v28's initial consolidation. Two independently-designed algebraic-adjacency constructions — a static associator-filtered random 2-complex and a sequential boundary-driven gluing process (DS) — both land in the log-diameter universality class, the sequential version even more decisively (R^2=0.9936), closing off the specific hope that a purely local, non-metric, algebra-native adjacency rule could supply escape for free; a third, purely deterministic variant (no injected randomness after the seed) closed into an exactly finite structure (N=33, machine-precision-exact, 8/8 seeds) —

an independent dynamical confirmation of Addendum AN's compactness obstruction. Addendum DT executed a previously-flagged, unpursued lead (Bianconi–Rahmede Network Geometry with Flavor, noted at Addendum K) by generalizing to a dimension-parametrized family: no tested dimension (d=2,3,4,6) escapes log-diameter, ruling out gluing dimension alone as an escape route. Addendum DU sharpened the stakes directly — verifying numerically, via Cheeger's inequality and Friedman's theorem on random regular graphs, that the log-diameter default branch is structurally incapable of supporting any real-space coarse-graining scheme at all — but found the sequential-gluing construction, despite its own log-diameter volume growth, has a genuine shrinking conductance bottleneck unlike a plain expander, confirmed via the Fiedler vector. Addendum DV extended this to a full two-diagnostic match (bounded Gromov δ-hyperbolicity + shrinking spectral gap) against a genuine hyperbolic random-geometric-graph reference, specific to low gluing dimension and not yet understood mechanistically. A cross-implementation methodological error — mixing implementations when comparing across gluing dimension — was found, logged, and corrected in place (DW): DT's diameter-scaling conclusion survives; a specific conductance-dimension claim (DV.4) is retracted. Addenda DX–DZ and EA isolated the true driver: a construction-selectivity ratio, degree- and dimension-independent, collapsing onto a single saturating curve (conductance exponent $\approx -0.29 \cdot \ln(\text{ratio})$, $R^2=0.953$) reaching a precisely-measured ceiling (exponent ≈ -1, $R^2=1.0000$ at its cleanest) — the strongest, cleanest signal in the whole thread. Addenda EB–EC tested that ceiling's geometric identity against flat, expander, and curvature-swept hyperbolic references and found a genuine trade-off rather than a match: lower curvature moves the exponent toward the ceiling but loses the bounded-δ signature that would make the match meaningful. Addendum ED confirmed the whole construction family's behavior is independent of continuous background access, surviving intact under the strictest possible finite realization of Postulate 1 — narrowing but not resolving this entry's shared root with Addendum AJ.6's own standing discretization question. No formal status-tag changes result from any of DS–ED; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY.

New at v29 (Addendum EG). DZ/EA's ~ -1 conductance ceiling, reported above, is now flagged as unconfirmed at the N range so far tested, not merely unexplained. Checked against Addendum B.6's early Cheeger-type dichotomy for the first time (EG.1–EG.3): the contextual-gluing construction satisfies bounded degree and non-polynomial growth while showing a shrinking, non-O(1) gap — a genuine third regime outside B.6's two stated branches, logged as an extension to B.6 rather than a violation of it. A candidate explanation for this — that the ceiling is a stable echo of B.6's branch-1 decay rate sitting inside branch-2's growth regime, offered as a possible account of Addendum EB.2's own previously-unresolved geometric-identity puzzle — was tested directly in a controlled Watts-Strogatz interpolation (EG.5–EG.6): the echo is real and reproducible at DS–EA's own N scale, but pushed to N=51,200 it proves to be a transient crossover, drifting from exponent $\approx -1.7/-1.8$ at small N toward $\approx -0.3/-0.5$ at large N, not a stable third regime. Two attempts to test this directly on the real construction (source code recovered via conversation_search, not reconstructed from memory) reached different outcomes: a reconstruction of Addendum DS.7/DS.8's earlier grow_hybrid mechanism was pushed to N=32,000 and stayed shallow and only mildly drifting (exponent -0.14 to -0.17), but at parameters that don't match either DU.2's or DZ/EA's own reported numbers for this construction family, so this result is not read as confirming or refuting the ~ -1 ceiling; a reconstruction of the later, more developed grow_contextual mechanism was calibrated directly against Addendum DZ.3's own reported exponents before any extension was attempted, and the calibration failed (context_size showed no measurable effect, versus DZ.3's large, monotonic, well-fit effect) — correctly halted rather than extended past an unvalidated build. No formal status-tag changes; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. Recommended next step: recover the actual cached contextual_gluing.py construction and parameters, or rebuild context_size's mechanism correctly from scratch and calibrate before any further large-N claim about the ~ -1 ceiling.

New at v30 (Addenda EH–EL): a single continuous thread, returned to Postulate 1's

own wording after the grow_contextual reconstruction above stalled, that reframed this entry's whole recent arc and found the same qualitative signature recurring at three independent levels. EH reconnected the DS–EG construction family to Addenda K–N's previously-closed Trugenberger/Kelly mapping — closed at the time because "a genuine selection dynamics over the coupling C_ij is structurally absent from Postulates 1–4" (Addendum L.5) — on the observation that DS–EG's own constructions already are such a dynamics, generated rather than rewired. Directly testing grow_hybrid (verified literal recovered source, not reconstruction) found a genuine bimodal survival transition — every run lands at either full sustained growth or collapse to N<50, zero intermediate outcomes across 80+ runs — confirmed N-independent (identical per-seed fate at N=300/600/1200) and independently reproduced on a second, differently-coded pre-fix construction along a different control axis, which also explained a previously undocumented finding from Addendum EG.8. EI then tested whether this substrate escapes Addendum AX.10's own growth-vs-selection trade-off: it clears Addendum CV's negative mechanism (triangle count scales linearly with N, not O(1) as on CV's honest random-regular substrates) but reproduces Addendum BT.2's negative extinct/universal dynamical outcome rather than BT.3's positive partial selection under Postulate 3's own recovered triangle-term dynamics, after three controlled isolation tests (topology, marginal coupling-value distributions, two-point spatial correlation) each eliminated the obvious explanation and one (topology) validated the dynamics implementation against BT.3's own native construction style. EJ identified the actual mechanism via direct basin-counting on the dynamics landscape: grow_hybrid's landscape converges to exactly 2 attractors, confirmed saturated across a 50–500 seed range, against a genuinely rugged, still-growing basin count (4 major basins, converged, plus slowly-accumulating finer sub-structure) for an honest i.i.d.-uniform substrate — and traced roughly half this gap to a real, statistically-supported (Spearman $p<0.05$, strengthening rather than washing out under noise-averaging) but partial causal channel, global S^6-coverage discrepancy, via a properly calibrated 24-realization test. EK returns to and formally closes, without resolving, this session's own opening contradiction: a direct proof that grow_contextual's recovered description — a fixed FIFO boundary queue in which every step unconditionally succeeds — makes the resulting topology, and hence conductance, provably independent of context_size, which cannot be reconciled with Addendum DZ.3's own reported context_size-dependent exponents; four candidate reconciling mechanisms were tested and none both matches the recovered description and reproduces DZ.3's numbers. Logged as a permanent open discrepancy, sharpening Addendum EG's own caution: DZ/EA's ~ -1 ceiling is now "unreproduced at any N, by four independent mechanisms," not merely "unconfirmed at larger N." EL found the same bimodal signature recurring a third time, now in landscape type itself rather than in growth or dynamics outcomes: the naive pre-fix construction does not have a single landscape ruggedness but two, splitting roughly evenly and categorically by construction seed (4 smooth, 3 rugged of 7 checked to convergence), and this split's severity was shown to peak specifically near the construction's own internal glue-onset critical boundary and damp on both sides — the textbook signature of critical fluctuations. Taken together, EH–EL sharpen this entry's standing synthesis from "associator-minimization is a landscape-smoothing/winner-take-all operator" into the more specific claim that such constructions exhibit critical-fluctuation-like bimodality — variance maximized near internal phase boundaries, damped away from them — independently confirmed at three levels of examination (growth success, Postulate-3 selection dynamics, landscape structure) and on two independently-coded constructions. Stated as CONJECTURE, well-motivated by convergent evidence, not yet a proof; no forcing argument exists. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. Recommended next step: test whether the same critical-fluctuation pattern replicates on grow_hybrid itself via a fine glue_bias sweep near its own 0.36–0.38 boundary, the most direct possible cross-construction confirmation.

New at v31 (Addenda EM–EP): four addenda, spanning H.17's own construction-generalization question, AW.5's lattice-vs-orbit tension, Postulate 4's dependency

structure, and a direct test of Postulate 2's founding wording. EM found the near-critical conductance-weakening signature from v30's EH–EL thread generalizes across three structurally distinct constructions, while the accompanying δ-hyperbolicity sign-crossing does not — see H.17 entry above for full detail. EN proved AW.5's tension has a hard mathematical core (full G_2-covariance and additive translation invariance cannot both hold for a nonzero antisymmetric single-difference coupling, a consequence of G_2's sphere-transitivity), built a constructive lattice-native alternative satisfying the weaker discrete symmetry, and showed — more generally than Addendum AU originally demonstrated — that this does not reopen Postulate 3's stabilization question, since the chiral-state instability was never sensitive to Γ's construction. EO found Postulate 4's bounded-degree requirement follows directly and with zero free parameters from Postulate 1's own lattice structure (minimal-vector adjacency), independent of Postulate 3's dynamics or May's criterion — a genuine, if partial, structural decoupling, verified against Addendum AW's own reported kissing number and growth degree exactly. EP tested whether Postulate 1's own lattice clarity (EO) motivates rewriting Postulate 2 around a triadic coupling: found well-motivated (the associator's antisymmetry is forced by the algebra alone) but not actionable without further work — it inherits EN's translation-invariance obstruction, explaining Addendum BF's historical failure and Addendum BU–BV's normalization convention mechanistically for the first time, but a concern this raised about origin-privileging in the coupling field was tested directly against BV's own real substrate and dynamics and found not to threaten BT–BW's core results. No formal status-tag changes anywhere in the theory; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 4 remains DERIVED, now on two independent grounds for its bounded-degree component.

Addenda CA through CE are not retracted by this resolution. They characterized real, reproducible phenomena of this dynamics under the Uniform$(-0.3,0.3)$ IC regime — a reasonable reading of "generic initial conditions" adopted transparently at BY.1 in the absence of recovered code, and used consistently and honestly by every session since, including this one's own CC–CE prior to CF's search succeeding. That regime's findings (the overlap-nucleation mechanism, R_i as a stronger-than-T_i predictor, topology-dominance over algebraic magnitude, the Gram-matrix associator-invariance proof) stand as their own line of results, now understood to be distinct from BT/BU/BV's own generic-IC claim rather than a series of failed attempts at reconstructing it.

What remains open in this entry, honestly: BT.4's own exact reported node counts (34–109) were not numerically reproduced, because BT's own original orbit-substrate seed was never recovered — only a statistically faithful reconstruction (by1) was available; this is a substrate-instance gap, not a methodological one. AX.10's growth-vs-selection trade-off is unaffected by this consolidation and remains as previously characterized. **AW.5's lattice-vs-orbit tension is no longer independent of this entry**: New at v34 (Addenda EZ–FB), above, found it evidentially linked to H.17 itself via AZ.13 — the same compactness obstruction (Addendum AN's discrete-subgroup finiteness, Addendum Z.15's Poincaré recurrence, Addendum FA.5's graph-distance saturation, Addendum FB.4's multiplicative closure) now confirmed to force H.17's own log-diameter fork-branch directly. *[Correction to this entry's own earlier framing, logged per the no-silent-corrections standard rather than silently updated: the "**remain as previously characterized**" language above predates that finding and is superseded for AW.5 specifically — unaffected for AX.10.]* → See the New at v34 (Addenda EZ–FB) entry, this section, above, for the full evidential chain. The R/S/T-vs-triangle-term underdetermination (Postulate 3's top-level status) is likewise unaffected — CF's resolution concerns the triangle term's own generic-IC behavior specifically, not its uniqueness among the four candidates. Recommended next steps, in priority order, now that the BM–BW archive recovery (formerly item 1 here) is complete as of v20: (1) an AS-scale growth-class scaling effort on the lattice substrate (BW, still untouched); (2) the still-partially-open per-seed nucleation-variance residual (CC–CD's R_i explains ~55–68%, not all, of it); (3) AW.5's lattice-vs-orbit tension itself, still untouched since it was first identified; (4) Postulate 3 stabilization proper

— the four-candidate underdetermination remains this project's single highest-priority open item and was not addressed by the BM–BW recovery.

Flagged at v25, not yet started: generalizing the DH–DI kink-propagation finding. Confirmed on real substrate triangles via exact zero-crossings, but tested on freshly-built substrates matching the documented recipe, not the theory's own canonical by1 N=3,000 instance used throughout the CW–CZ/CU record. Recommended: (1) repeat the exact-crossing-plus-full-dynamics test directly on the canonical by1 substrate; (2) extend DI's 22-substrate severity comparison with a larger, higher-success-rate batch (current pipeline's crossing-search succeeds only ~41% of the time per candidate, consistent with the already-logged narrow-wedge parameter sensitivity, BL.5/BH.6); (3) determine whether the propagated-discontinuity phenomenon has any implication for the theory's dynamics beyond the two literal-coefficient candidates it was tested on — e.g., whether it affects trajectories that pass near, rather than exactly through, a kink locus during ordinary time evolution, which DA–DI's tests (deliberately constructed zero-crossings) do not address.

New at v48 (Addendum HI): the invariant-theory landscape behind the three literal-coefficient candidates is now completely mapped, closing off rather than opening a fourth-candidate search — the three-way underdetermination itself remains exactly as open as before. Checking the specific gap Addendum U's own recommended next step names (does Postulate 1–2's algebraic substrate origin constrain the choice among the candidates), HI found Addenda AL/AM/AR — cited as having executed this — actually address a different, superseded framework (the pre-Addendum-CT R/S/T stabilization mechanism) and never touch associator-norm, $|\varphi|$, or φ^2 at all; nobody had actually asked the question for the current candidate set. Building the G2-Casimir decomposition of $\Lambda^3(\mathrm{Im}\,\mathbb{O})$ from scratch (independently reconstructing $g2=\mathrm{Der}(\mathbb{O})$, confirming the classical $1\oplus7\oplus27$ split structurally rather than citing it), HI found and verified an exact new identity, $\varphi(x,y,z)^2+\tfrac{1}{4}|[x,y,z]|^2=\det(\mathrm{Gram}(x,y,z))$, holding to machine precision across thousands of trials and multiple seeds — not previously in the project's record. Two errors were caught and corrected in place while interpreting this identity (an incorrect first hypothesis about what the associator-norm term captures; a numerically-found bound later explained rather than merely confirmed), converging on a complete resolution via Schur's lemma: since the octonion associator is alternating (a classical consequence of alternativity) and G2-equivariant, it defines a G2-module homomorphism $\Lambda^3V\to V$ that Schur's lemma forces to vanish on the 1-piece and the 27-piece entirely — the associator was never measuring anything beyond the 7-dimensional piece, confirmed directly (a fixed ratio of exactly 16.000000, not merely close, across independent trials). Combined with φ^2 already being exactly the 1-piece, this shows the theory's two existing candidates jointly and exactly determine all three pieces of a triple's G2-decomposition — there is no fourth independent invariant anywhere in this structure for a new literal-coefficient candidate to draw on. φ^2 remains the best-evidenced candidate on record, now supported by two structurally different lines of evidence (Addenda DQ–DR's dynamical/skewness account, and this addendum's clean algebraic role as the sole term appearing directly and additively in an exact conserved decomposition) without either constituting a forcing proof. No formal status-tag change: Postulate 3 remains POSTULATED, UNDERDETERMINED among {associator-norm, $|\varphi|$, φ^2}. See Addendum HI in full.

6.2 Medium priority — scoped, pending dedicated work

W_ij pairwise non-atomicity (H.13); Clifford multiplication rule (Postulate 8); quantum decoherence floor and CMB polarization predictions; Postulate 4 two-sided degree bound formal adoption; heterogeneous k_min verification; near-threshold admissibility regime; exponential-tail circularity interaction — all unaffected by this consolidation.

Triadic Substrate Postulate (Postulate 7) — item (i) (which 3 of 6 octonion generators, AZ.13): the strong (symmetry-invariant) version remains CLOSED, NEGATIVE, permanent

(Addenda BC, BD). UPGRADED at v22 (Addendum CR): a genuine dynamical (vacuum-selection) candidate mechanism was found for direction-selection and composed with Postulate 7's own triad mechanism, moving this from "no mechanism attempted" to PARTIAL — contingent on an unforced tachyonic-mass-sign input Postulates 1–3 do not currently force. INVESTIGATED at v23 (Addendum CS): the natural next step — does anything in Postulates 1–3 force that sign? — was taken up directly. Two candidate routes were closed (DERIVED negative): a literal Coleman-Weinberg / RG-running derivation is foreclosed because $\tau = -\ln \Lambda$ carries no genuine renormalization-group content (CS.2, also permanently closing Addendum B.3's Fork B), and a CP-style nearest-neighbor coupling is structurally blind to the uniform mode the sign question concerns (CS.6). One route remains genuinely open but unforced: a direct coupling of the direction-selecting sector to the substrate's own inhomogeneous background, which would require justifying a new term (CS.3–CS.4). The contingency therefore stands, now better mapped — two plausible-looking derivations ruled out with mechanisms, one honest option left standing. NEW AT V26 (Addendum DK.14–DK.17): the remaining open route was revisited after correcting a scope slip in CS.5 (CS.3's proposed coupling needs no site index on v — M is already summed over sites; logged as a correction, not silently worked around). The full space of G2-invariant v–M couplings was then classified directly (commutant of g2 on $\mathrm{Sym}^2(V)$, exactly 2-dimensional) and CS.3's proposed form confirmed to sit at the natural equal-weight point within it — tightening the functional form without discharging CS.4's actual gate, whether coupling CR's spatially-uniform order parameter to substrate configuration data is licensed at all. That legitimacy question remains the live blocker; AZ.13's $\gamma < 0$ contingency is unchanged. Also new at v26 (Addendum DK.10–DK.12): the strong (symmetry-invariant) closure of the "3 of 6" question (BD.5–BD.6, below) is reinforced by an independent, general proof that the leftover 6-dimensional space admits no nontrivial su(3)-invariant subspace of any dimension — permanently foreclosing any future algebraic/symmetry-based selection attempt, not only the specific patterns BD.5–BD.6 checked. Item (ii) (triad-sector stabilization) is DERIVED (Addenda BC, BD), now confirmed at v22 (Addendum CP, Gate 1) to rest on a genuine, newly-established site-dependent field theory rather than an unverified single-site construction.

[Documentation defect found and fixed: misfiled priority tier.] This four-addendum v31 update (Addenda EM–EP — H.17's construction-generalization question, AW.5's lattice-vs-orbit tension, Postulate 4's dependency structure, and a Postulate 2 triadic-coupling test) was logged in this Medium-priority section instead of the Highest-priority H.17 chain, where the bullet's own content belongs by the register's own stated criterion (it states H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY — explicitly a Highest-priority, directly-load-bearing item). No v31 entry existed anywhere in the H.17 chain prior to this correction — the chain jumped from v30 directly to a v25-dated flag, silently skipping this consolidation's findings (including EN's no-go theorem for Postulate 2's Γ). Relocated to the Highest-priority H.17 chain above, inserted chronologically after the v30 (EH–EL) bullet. Logged here per the no-silent-corrections standard rather than deleted outright.

Route A/Route B fork (Postulate 8) — substantially exercised at v22 (Addenda CO, CP, CQ) without resolution of BG's core "is it required" question, which remains untouched. The soliton-based sidestep is closed (CP, Gate 2, NEGATIVE). Route B's internal 3D consistency is now well-established (CO, CQ), using only structure already established elsewhere in the project, with no violation of Addendum M.3's graph-level-bipartite reading. NEW AT V26 (Addendum DK.1–DK.9, DK.18): a new angle was found and run — the g2-classification technique (built at Addendum DB, reused at DJ) applied to the specific-realization sub-question, not the necessity question itself. Given Route B and G2-equivariance, the octonion-multiplication generator realization is forced, unique up to gauge and one discrete chirality choice, shown exactly to be octonion conjugation — closing AZ.12/AZ.13's realization gap. BG's own six-for-six negative result on whether Route B is required at all is untouched. A flagged resonance between the found chirality sign and Postulate 6's "chirality-parity" naming was checked and closed by dissolution (DK.18) — no standalone parity operator exists in the record. H.17's

resolution remains the more load-bearing dependency going forward, per Addendum AZ.14's own tracing, unaffected by this update.

Standing pointer to Addendum BL.4 in Part II's Postulate 3 text — IMPLEMENTED at v25. Flagged at v22's Gemini-review audit (Part V S2d), re-flagged at v23 (Addendum CS.1), and finally written into Part II's Postulate 3 text at v25 once Addendum DA.2 gave it something more precise to say than a repeat refutation: BL.4's G2-invariance claim stays correctly refuted; DA.2/DH–DI's differentiability finding is separate, later, and narrower, and future reviews should be checked against that record instead.

Which Postulate 3 stabilization mechanism to adopt, if any: four originally-characterized candidates (R, S, T, triangle/associator-norm). R, S, and T are permanently disqualified from non-patching status by a standing theorem (Addenda CG, CI): their defining equations cannot, for any construction of their type, ever carry Postulate-1 geometric content, and by the project's own postulate-patching standard they are patches, not forced mechanisms (Addendum CT). The non-patching class for the triangle-term position is now fully characterized (Addendum CU): exactly one new G2-invariant 3-point scalar exists beyond the pairwise Gram data already used via W_ij — the associative 3-form $\varphi(a,b,c)=\langle a,bc\rangle$ — with the associator-norm an exact derived combination of φ and that pairwise data ($|[a,b,c]|^2=4(\text{Gram\_det}-\varphi^2)$, verified to machine precision). At v25 (Addendum DB.3–DB.5), φ's status was strengthened from CU's search-based finding to a proven classification result: φ spans the unique 1-dimensional G2-invariant subspace of $\Lambda^3(\mathbb{R}^7)$ (35-dimensional total, matching the classical $1\oplus7\oplus27$ decomposition), re-derived directly from this project's own octonion implementation rather than cited. This does not extend to the associator-norm, which is fully symmetric (Sym^3) rather than alternating (Λ^3) — a different representation space DB.5's classification does not cover (Addendum DC). Both the associator-norm and |φ| (not signed φ — Addendum DD found signed φ fails the theory's own generic-IC protocol outright, narrowing the viable literal-coefficient class to exactly these two) produce converged, spatially coherent local attractors on independently-replicated real substrates, both have a natural non-fine-tuned coupling scale (BM; CU.10), and both carry a confirmed, dynamically consequential non-differentiability at their own zero locus with no established advantage for either on proximity (Addendum DG) or propagated severity (Addendum DI). STATUS: POSTULATED; the choice between |φ| and the associator-norm as the literal coefficient remains open, but the class itself is closed and now provably exhausted on the alternating side — not an open-ended or merely-two-member family by coincidence. Addendum CT.3's earlier "uniquely privileged" framing is superseded by CU.9, logged as a correction, not silently overwritten; CU.4–CU.5's "|φ| shows greater discoverability" framing should now be read alongside DD's finding that this was already implicitly comparing |φ| against associator-norm, never signed φ against either.

Whether the saturating-coupling and amplitude-dependent-λ stabilization mechanisms have the full two-branch robustness confirmation the quintic case received (Addendum T: 151/151 for saturating coupling; Addendum U: only a handful of points for amplitude-dependent λ).

Addendum BW.6–BW.7's coherence-degradation question — SUPERSEDED at v17, see the merged Highest-priority entry above (Addenda BX–CB) and "Substantially resolved," below. No longer listed as a standalone medium-priority item.

Carried from v17, partially advanced at v18, unaffected by this consolidation: does the overlap-nucleation mechanism (Addendum CB) generalize beyond the one substrate and one signed coupling it was found on? CB found ignition nucleates preferentially at overlap-concentrated sites under the G2 3-form signed coupling on the N=3,000 orbit substrate. CC–CE collectively tested the (distinct but related) magnitude-coupling seeding mechanism across three substrates — a toy reconstruction, the real by1 orbit substrate, and the real BU/BV lattice substrate — finding structured/overlap-informed seeding is required in all three, under the Uniform(−0.3,0.3) IC regime. Still untested: whether CB's own signed-coupling result specifically generalizes to the lattice substrate or to other h; this remains open

in its original form.

Carried from v17, substantially advanced at v18, unaffected by this consolidation: the per-seed nucleation variance (Addenda BY.2–BY.4, BZ.3, CB.4). Note for clarity, given Addendum CF's resolution of the separate generic-IC discrepancy this consolidation: this entry concerns *seeded/localized* ICs under the Uniform(−0.3,0.3) regime specifically, confirmed by CF.6 to be a real and legitimate (if distinct from BT/BU/BV's own) line of results — not superseded or called into question by CF. A candidate explanation has been tested (Addendum CC's R_i metric, $R^2 \approx 0.55$–0.68 across two samples); T_i is largely redundant with it once R_i is known. A residual of roughly 32–45% remains genuinely unexplained by either metric.

CA.4's four candidate explanations for the generic-IC parametrization discrepancy — SUPERSEDED this consolidation, see the Highest-priority entry above and "Substantially resolved," below (Addendum CF). Final disposition: (a) alternative signed/vector-valued associator reductions — DISCONFIRMED (CC), correctly cleared. (b) a different initial-condition convention — CONFIRMED AS THE ANSWER (CF): generic ICs drawn from Uniform(0.3,0.8), not Uniform(−0.3,0.3). (c) a different convergence/classification criterion — DISCONFIRMED (CC), correctly cleared. (d) re-examination of the archived claims against primary artifacts — this is, in effect, what CF did once the literal code was recovered, and the claims held up exactly as originally reported. No longer listed as a standalone medium-priority item.

6.3 Documentation / writeup tasks (no new physics required)

I.2(b): the originally-planned (3,1) vs. (2,2) linearized-stability comparison — SUPERSEDED BY H.19, per I.2's own recommendation.

Documentation-propagation bug — process fix. Still outstanding as a standing process gap, but with a new data point at v23. When the v22 master was read from project knowledge at the start of this consolidation, it again surfaced as plain-text UTF-8 rather than binary docx — the same signature as the eleven recorded instances. This session did not treat that as a fresh twelfth corruption: consistent with the mount-artifact reading logged this session (the plain-text appearance is how project-knowledge surfaces the file's content, not necessarily proof the stored file is corrupt), the v23 master was built and verified as genuine binary docx (PDF-render confirmed) regardless. The underlying process gap — no automated build-verification step guaranteeing binary output — remains unaddressed and is still the right dedicated fix. The separate twelfth, distinct-mechanism instance (literal markdown-style asterisks left as text, found across two of three uploaded v21 copies) remains noted in the front matter, not yet formally numbered, pending the architect's confirmation of its cause.

6.4 Substantially resolved at v16 (Addenda BM–BQ) — carried from v16

- BJ.6/BK.4's flagged natural-scale first-principles gap — CLOSED (Addendum BM): derived from Postulate 1's fixed substrate dimension (6) and an exact algebraic vanishing theorem, matching prior empirical calibration to ~4% independently, with zero circularity.
- Whether R/S/T share the triangle term's locality property — RESOLVED, no (Addendum BN), on a fully validated network and search methodology, after an explicitly-logged mid-session correction when the validating baseline itself initially failed to reproduce.
- BK.4's principled, non-random generator-selection question — RESOLVED for the network it was derived on (Addenda BO–BQ): an independent replication (BO), a real selec-

tion criterion found and corrected (BP, after an earlier confounded signal), and out-of-sample validation at 90% accuracy with zero false positives (BQ).

6.5 Substantially resolved at v16 (Addenda BR–BS) — carried from v16

- Whether BQ's validated selection criterion generalizes across topology instances, in absolute form — RESOLVED, no (Addendum BR): refuted by a clear, non-boundary counterexample on a fresh network, and by real successes on a third network falling below the threshold entirely.
- Whether a network-relative (percentile-rank) version of the criterion generalizes instead — RESOLVED, no (Addendum BS), after explicitly catching and correcting an initially-misleading pooled statistic dominated by the derivation network's own in-sample effect. BK.4's original question is now understood to be closed for one network, not in general — reframed as a distinct, larger open problem.

6.6 Substantially resolved at v16 (Addenda BT–BW) — carried from v16

- Whether Addendum AX's growth-vs-selection trade-off is fundamental to the octonion/G2 substrate, or specific to the pairwise mechanism AX tested — RESOLVED, specific to that mechanism (Addendum BT): the triangle-term mechanism, unavailable when AX was written, produces genuine spatially coherent selection directly on AX's own reconstructed substrate, at zero pairwise coupling, with a mechanical explanation traced to Addendum BM.1's own vanishing theorem.
- Whether Addendum AW.5's lattice-vs-orbit tension admits any working reconciliation — YES, on one instance (Addenda BU–BV): a genuine, independently-verified octonion lattice supports both classical/structural polynomial growth and genuine triangle-term selection, first shown via normalized point-directions (BU) then confirmed via the full unnormalized octonion product directly, closing an honestly-flagged gap in the first construction (BV).
- Whether these results scale toward Addendum AX's own N=100,000 benchmark, and generalize across substrate instances — PARTIALLY TESTED, mixed (Addendum BW): a real sampling-density confound was found and fixed; growth-class discrimination remains genuinely ambiguous even controlled for it, mirroring Addendum AQ's own pre-AS ambiguity; cross-instance generalization at matched scale is positive (a narrower test than BR/BS's cross-topology failure); the selectivity phenomenon itself persists robustly across a 30× increase in N, but spatial coherence substantially degrades over that range — the sharpest open question v16 left behind.

6.7 Substantially resolved at v17 (Addenda BX–CB) — carried from v17

- Whether Addendum L's clustering bias and the triangle term's own locality (BH/BM) are genuinely connected or only superficially similar — RESOLVED, genuinely connected (Addendum BX): overlap concentration, isolated from raw triangle count and from degree density via two rounds of increasingly strict controls (an initial confound caught and corrected mid-thread), measurably increases locality propensity, replicated at two independent toy configurations.
- Whether this overlap mechanism transfers to H.17's actual substrate — RESOLVED, yes (Addendum BY): confirmed as both an existence condition and a containment mecha-

nism for seeded partial selection; distant overlap-anchored clusters shown to coexist independently and superpose exactly, supplying a direct candidate mechanism for BW.6, registered as a falsifiable prediction before testing.

- BW.6's open question — why does spatial coherence degrade with substrate scale — SUBSTANTIALLY REFRAMED (Addendum BZ): a mechanism-constructed multi-cluster state at BW's own N≈90,778 reproduces BW's reported 14% global contiguity almost exactly (0.171) while per-component coherence (0.893 mean) exceeds every prior reference level in the record. SUPPORTED for mechanism-constructed states; full adjudication against BW's own literal survivor state remains open pending recovery of its exact parametrization.
- Whether a signed (vs. magnitude) associator reduction explains the standing BT/BU/BV generic-IC discrepancy — RESOLVED, no (Addendum CA): the natural G2-invariant 3-form reduction, tested on both the orbit and lattice substrates, produces no stationary generic-IC partial regime at any tested h. The leading candidate explanation is eliminated; the underlying discrepancy remains open, narrowed to four remaining candidates (CA.4). A genuine new disorder signature (IC-dependent ignition threshold) was logged as a side-finding.
- What causes CA's IC-dependent ignition-threshold disorder — RESOLVED, mechanistically (Addendum CB): every run passes through a sharp activity bottleneck; surviving nucleation sites are significantly overlap-enriched against a degree-matched null (z=+7.29, n=1,867, replicated across 10 seeds), after a more natural global-cubic-form hypothesis was tested and cleanly disconfirmed. Unifies CA's finding with the BX–BZ mechanism and reframes H.17's standing structured-seed gap into a concrete sub-question (see the merged Highest-priority register entry).
- A pre-existing documentation gap (Addenda BM–BW missing from SFT_Addenda_Archive.docx) — FOUND AND LOGGED at v17, REPAIRED at v20 (see Part V §3 and the new v20 section below). BX–CB were appended in full at v17; the BM–BW gap did not compound further and is now fully closed.

6.8 Substantially advanced at v30 (Addenda EH–EL)

- Whether the DS–EG construction family constitutes a genuine selection dynamics over adjacency, the missing ingredient that closed the Trugenberger/Kelly mapping at Addendum L.5 — YES, on the observation that these constructions generate C_ij once, sequentially, rather than revising Postulate 2's fixed-data character (Addendum EH.0). Directly testing grow_hybrid (verified literal recovered source) found a genuine bimodal survival transition — full sustained growth or collapse to N<50, zero intermediate outcomes across 80+ runs — confirmed N-independent (Addendum EH.2) and independently reproduced on a second, differently-coded pre-fix construction along a different control axis, also explaining a previously undocumented Addendum EG.8 finding (Addendum EH.3).
- Whether this substrate escapes Addendum AX.10's own growth-vs-selection trade-off — NO, though it clears one necessary condition (Addendum EI.1–EI.6): triangle count scales linearly with N, not O(1) as on Addendum CV's honest random-regular substrates, but Postulate 3's recovered triangle-term dynamics still gives Addendum BT.2's negative extinct/universal outcome, not BT.3's positive partial selection. Three controlled isolation tests (topology, marginal coupling-value distributions, two-point spatial correlation: Addendum EI.5.1–EI.5.4) each eliminated the obvious explanation, with the topology test also validating the dynamics implementation against BT.3's own native construction style.
- What actually distinguishes grow_hybrid's landscape from an honest uniform substrate — IDENTIFIED via direct basin-counting on the dynamics landscape (Addendum EJ.1–

EJ.4.1): grow_hybrid converges to exactly 2 attractors, confirmed saturated across a 50–500 seed range; an honest i.i.d.-uniform substrate converges to 4 major basins at coarse resolution with finer sub-structure still slowly accumulating at fine resolution. Roughly half this gap is traced to a real, statistically-supported (Spearman $p<0.05$, strengthening under noise-averaging rather than washing out) but partial causal channel — global S^6-coverage discrepancy — via a properly calibrated 24-realization test (Addendum EJ.7); three-point h_ijk homogeneity was tested and eliminated as an alternative explanation.

- The session's own opening contradiction (DZ.3's reported context_size effect on grow_contextual versus a topology-invariance proof developed the same session) — FORMALLY CLOSED, UNRESOLVED (Addendum EK): the recovered description's fixed FIFO boundary queue makes conductance provably independent of context_size; four candidate reconciling mechanisms were tested and none both matches the recovered description and reproduces DZ.3's numbers. Logged as a permanent open discrepancy, sharpening Addendum EG's own caution on DZ/EA's ~ -1 ceiling from "unconfirmed at larger N" to "unreproduced at any N, by four independent mechanisms."
- Whether the smooth/rugged landscape split recurs on a second construction, and what governs it — YES, and TRACED TO A CRITICAL-FLUCTUATION PATTERN (Addendum EL): the naive pre-fix construction splits categorically by construction seed (4 smooth, 3 rugged of 7 checked to convergence, Addendum EL.1–EL.3), and this split's severity peaks specifically near the construction's own internal glue-onset critical boundary and damps on both sides (Addendum EL.4) — the textbook signature of critical fluctuations, now independently confirmed at three levels (growth, dynamics, landscape structure) and on two independently-coded constructions.
- No formal status-tag changes result from any of EH, EI, EJ, EK, or EL. H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 1 remains POSTULATED; Postulate 3 and Addendum AX.10 are unaffected in formal status. Next unused addendum letter: EM.

6.9 Substantially advanced at v29 (Addenda EE–EG)

- Whether CS.4's own internal pointer to CS.5 actually resolves its gating question — NO, checked and corrected (Addendum EE.1): CS.5 established only spatial uniformity, a fact about the coupling's form, not an answer to what v1, v2 physically represent. CS.4 remains open, now for a precisely identified reason. A previously unflagged drift between Postulate 7's own text ("algebraic triads of substrate elements") and BD/CR's realized construction (free internal fields, not substrate-derived quantities) is surfaced, two readings left open (Addendum EE.2).
- Whether the project's own postulate-patching criterion (Addendum CT.1–CT.3) licenses CS.3's proposed v–M coupling — MECHANICALLY YES, but the criterion was never tested against a cross-sector case before (Addendum EE.3–EE.5): every prior application (the triangle term) reached deeper into data already integrated within Postulate 3's own sector; CS.3 would open an entirely new channel between Postulate 7's v and Postulate 1/2's M. A gap in the standard itself, not merely in this application of it. CS.4 remains open.
- Whether CR/BD's already-verified dynamical selection mechanism can supply Postulate 8's separate "which 4 of 7" generator question — NO, but sharpened substantially (Addendum EF.1–EF.8): DK.7–8's chirality-operator construction is confirmed to work for any orthonormal 4-frame in $\mathrm{Im}(\mathbb{O})$, not only a canonical placeholder (EF.1, a useful general fact independent of the rest). An initial counting obstruction is raised and corrected in place — DK.8's own text already implies a clean 4+3 partition, under which $\hat{n}$ is naturally one of Postulate 8's four (EF.2–EF.4). What remains open is narrower than before:

the other three of Postulate 8's four generators must lie in a specific, already-identified 3-dimensional subspace L, which g2 is independently re-derived to show is a genuine flat direction — not fixed by any residual symmetry, nor by any dynamics currently in the theory (EF.5–EF.8). "Which 4 of 7" is now "which basis of L."

- Whether Postulate 4's early Cheeger-type dichotomy (Addendum B.6) holds up against the v28 H.17 thread's newest construction — INCOMPLETE AS STATED, extended (Addendum EG.1–EG.3): DZ/EA's contextual-gluing construction is bounded-degree and non-polynomial-growth with a shrinking (non-O(1)) gap — a genuine third regime outside B.6's two illustrative branches, logged as an extension rather than a retraction. A candidate explanation connecting this to Addendum EB.2's own open geometric-identity puzzle is tested in a controlled model and found to be a transient crossover, not a stable effect (EG.4–EG.7); two attempts to test this on the real construction (source recovered via conversation_search) produced one large-N result under different, honestly-flagged parameters (EG.8–EG.9) and one correctly-halted failed calibration (EG.10–EG.11). DZ/EA's own ~ -1 ceiling is left explicitly unconfirmed at the N range so far tested. See Postulate 4 and H.17's own register entry above for the full account.
- No formal status-tag changes result from any of EE, EF, or EG. AZ.13's $\gamma<0$ contingency (Postulate 7 sense) and Postulate 6/8's "which 4 of 7" question (Postulate 6/8 sense) are both unchanged in formal status; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY; Postulate 4 remains DERIVED. Next unused addendum letter: EH.

6.10 Substantially resolved at v43 (Addenda GO–GQ)

- Whether the Postulate-3 literal-coefficient tie-back (associator-norm, $|\varphi|$, φ^2) can be closed by extending DB.5/DJ's classification technique beyond the alternating trilinear part it was originally confined to — the full space of G2-invariant polynomials on (a,b,c) is classified exactly, for the first time (Addendum GP). At degree 3: φ is the unique G2-invariant trilinear function, full stop (strengthens DB.5). At degree 6: exactly 6-dimensional, explicit basis given; φ^2 and associator-norm^2 both proven independently necessary. RESULT IS NEGATIVE for the hoped-for forcing theorem — no DB.5-style uniqueness exists at this degree; the two candidates are co-equal, not one forced over the other. A genuine structural fact (Gram_det = φ^2 + ¼·associator-norm^2) is independently re-derived via invariant theory; this is Addendum CU.9's own identity, not new — an overclaim of novelty caught and corrected within the same addendum once checked against the archive.
- Whether φ^2 shows the same dynamical kink-avoidance signature associator-norm/$|\varphi|$ were shown to lack, extending DH/DI's protocol to φ^2 for the first time and, separately, running DH's own flagged-but-never-completed multi-triangle generalization test for all three candidates simultaneously — SUGGESTIVE, NOT FORCING (Addendum GO): across seven triangles properly targeted via a two-stage joint-criticality screening (after an initial mean-field-dynamics substitution was caught as the wrong tool and corrected), φ^2's response is tightly bounded (3.6–4.7%) in every live instance while $|\varphi|$'s spans three orders of magnitude including two genuine derivative discontinuities — the expected statistical signature of smooth versus kinked, not a proof. Two in-place corrections: the screening's yield was originally mischaracterized (28 of 60 candidate triangles had genuine tri-candidate life, not 2 of 60 as first stated); the low rate of well-balanced cases is traced to genuine spatial clustering of activation state near threshold (8× excess variance over independent-sampling expectation), a real feature of this class of cascade dynamics, not a screening weakness.
- Whether Addendum EF.7's flat direction ("which basis of L", Postulate 6/8's leftover generator question) can be resolved — YES, RESOLVED AND PROVEN (Addendum GQ), superseding EF.7's own finding in place rather than silently: for BD's dynamically-selected

associative triad and n-hat orthogonal to it, the associator [triad_i,triad_j,n-hat] over the three triad pairs gives an exact orthonormal basis of L (Gram=4I_3), proven directly via Cayley-Dickson doubling of the triad's own quaternion subalgebra — not merely verified numerically — and fully general (depends only on the triad's associativity and n-hat's orthogonality to it, not on how n-hat was itself constructed). Replicated across four independent full dynamical realizations (the real CR.3/BD potentials implemented and optimized from scratch, not a random stand-in). A companion general identity found en route: for any orthonormal triad in the octonion imaginary space, the map x to (phi(t1,t2,x),phi(t1,t3,x),phi(t2,t3,x)) is an exact conformal submersion from the triad's 4-dimensional complement onto R^3 (F4·F4-transpose = c·I_3), with the physically-real associative triad sitting at this identity's exact zero-c degenerate limit — the reason phi gives no information on the real configuration and the associator was needed instead. RECOMMENDATION: Postulate 6/8's "which 4 of 7" entry upgraded from CONJECTURE/OPEN toward DERIVED at this consolidation (applied below).

- Two further corrections, logged in place within the same addendum: Addendum GA.5's proposed link between this question and Addendum BI.7's seed-point (e0) arbitrariness does not hold — for two independent reasons, not one. GQ's resolution of L never references e0 at all; and BI.7 itself was already closed by Addendum BJ (the triangle term's ≈200× scale factor traced to an orbit-parametrization artifact, not a genuine free constant), well before GA.5 was written and cited it as open.
- STATUS-TAG CHANGE, this consolidation: Postulate 6/8's "which 4 of 7" question (Postulate 6/8 sense) upgraded from CONJECTURE/OPEN to DERIVED (Addendum GQ.11–GQ.13, a closed proof, not numerical evidence). No other status-tag changes result from GO, GP, or GQ. Postulate 3's literal-coefficient class remains POSTULATED, UNDERDETERMINED among {associator-norm, $|\varphi|$, φ^2}; AZ.13's $\gamma<0$ contingency (Postulate 7 sense) is unchanged; H.17 remains CONJECTURE/OPEN, PERMANENT ENTRY. Next unused addendum letter: GR.

6.11 Substantially resolved at v18 (Addenda CC–CE) — carried from v18

- Whether a premature classification cutoff (short integration windows misreading transients as arrested states) explains the BT/BU/BV discrepancy — RESOLVED, no (Addendum CC): full trajectories tracked to t=300 across a wide h-sweep show sharp, fast transitions, not slow creep a short window could misclassify.
- Whether a different signed or randomized-sign associator reduction permits generic-IC arrest where the natural G2 3-form (eliminated at CA) does not — RESOLVED, no (Addendum CC): both a sign-only reduction and a fully decorrelated frustration control remain strictly binary from generic ICs.
- What predicts the individual-seed nucleation variance named at BY.2–BY.4/BZ.3/CB.4, previously untested — SUBSTANTIALLY ADVANCED (Addenda CC–CD): a new metric, pairwise triangle-overlap (R_i), predicts individual-seed threshold more strongly than the register's existing T_i, confirmed on both an independent reconstruction and the real by1 substrate. A genuine residual (~32–45%) remains open.
- Whether the octonion associator's specific magnitude values (as opposed to topology alone) drive seeding threshold — RESOLVED, no, given fixed topology (Addenda CC, CE): two independent decorrelation tests (magnitude shuffling; fully independent random weights) find topology predicts threshold at least as well as the real algebra, on both a reconstructed and the real substrate. Whether the algebra is essential to generating the topology itself (as opposed to only weighting it) remains untested — CONJECTURE/OPEN.

- Whether a prior session's exact generating code, rather than only its prose description, can be recovered without a surviving cached file or the architect resupplying material — RESOLVED, yes (Addendum CD): conversation search over this project's own past sessions locates literal construction code (RNG seeds, algorithms), demonstrated on the by1 substrate and independently verified against that session's own contemporaneous sanity figures.
- Whether this recovery technique generalizes to the project's oldest standing discrepancy (BT/BU/BV) — RESOLVED, yes, partially closing the task (Addendum CE): BU/BV's exact lattice substrate and coupling-scale convention recovered and verified to an exact match (5,801/5,801 triangles; associator mean 5.658 vs. 5.66 reported). Three further candidate explanations for the discrepancy itself eliminated (unscaled h; wide h-scan; octonion basis convention, the last via a general Gram-matrix-invariance proof, not only an empirical test) without reproducing the reported result — the underlying discrepancy survives this attempt and remains OPEN, now narrowed to the literal dynamics/IC code specifically.

6.12 Substantially advanced at v21 (Addenda CG–CN) — carried from v21

- The single-site-invariant derivation route for R/S/T's free parameters (μ, s, A_c) — CLOSED (Addendum CG): a general G_2 orbit-transitivity theorem, independently re-derived from scratch this session (octonion algebra and g2 rebuilt directly from defining constraints, not reused from Addendum BG/BO), strictly generalizes Addendum AR's earlier associator-squared-specific negative finding to every possible single-site construction.
- Whether R/S/T's non-locality (Addendum BN) holds on the theory's actual H.17 substrate, not only toy Watts-Strogatz networks — CONFIRMED, yes (Addendum CH), under both generic and localized seeding with long-time convergence verified rather than assumed; one promising-looking localized-seed counter-example was caught mid-session as a slow-expanding transient front, not a genuine attractor, and logged as a false lead rather than silently dropped. A secondary finding narrows Addendum AR's own forced-θmax=2 degeneracy to its original substrate construction specifically — it does not carry over to the k-NN orbit substrate used here.
- Why no derivation route, single-site or global-aggregate, has ever worked for R/S/T — UNDERSTOOD, structurally (Addendum CI): unlike the triangle term's own defining equation, which already contains a genuine Postulate-1 geometric object (the associator) for a coupling constant to normalize against, R/S/T's defining equations contain no Postulate-1 content anywhere in their functional form — a single structural fact subsuming three previously separate negative results (Addenda AM/AR, AR's no-preference finding, and Addendum CG).
- Whether triangle-closure is required for a multi-site correction to escape Addendum AY's no-go theorem at all — RESOLVED, no (Addendum CJ): an open "wedge" motif (same functional form, no requirement that the third edge exist) escapes the extinct/blowup dichotomy exactly as well as the closed triangle, isolated from and then confirmed with real pairwise coupling present. Only the degree-matching relationship between the driving and restoring terms governs bare stabilization, not motif topology.
- Whether triangle-closure is required for the specific local-attractor regime Addendum BH found — INITIALLY OVERCLAIMED, THEN CORRECTED IN PLACE (Addenda CK–CL): BH's own local attractor was exactly reproduced (matching Addendum BN.5's independently-reported profile to three decimal places) via recovered source code and a validated fsolve-based search methodology; the wedge motif was initially found to

support no stable fixed point of any kind at the triangle's own validated point (Addendum CK), but a fuller 2D parameter scan (per CK's own flagged next step) found genuine, robustly stable wedge local attractors elsewhere in parameter space (Addendum CL) — correcting CK.5's overly broad reading. The real, quantified distinguishing property is discoverability, not existence: an approximately 25–35× gap in basin size/discoverability between the two motifs at their own respective best points.

- Whether this discoverability gap is general or an artifact of the one graph tested — SUBSTANTIALLY ADVANCED (Addenda CM–CN): the gap replicates in direction and comparable magnitude across three independently-seeded Watts-Strogatz graphs (Addendum CM, ratios ≈8.5–35×), and replicates in direction — though at a notably smaller magnitude, ≈5× — on the theory's own real octonion-orbit substrate (Addendum CN), after recalibrating the working parameter window specifically for that substrate's denser triangle structure. The substrate-dependence of the gap's magnitude is itself flagged CONJECTURE/OPEN, not resolved.

6.13 Substantially advanced at v23 (Addenda CO–CR) — carried from v23

- Whether BB.1's non-degeneracy criterion for Postulate 8's Route B could emerge generically from Route A's own machinery at a band-touching point, rather than only ever being postulated directly — YES in 2D, NO in 3D under strict bipartiteness (Addendum CO): 186/300 generic and 194/300 octonion-derived hoppings produce a genuine, 100%-non-degenerate band-touching point in 2D (Cl(0,2) structure, confirmed via Sylvester's law); in 3D a 2-band model supplies only 2 real functions of k regardless of dimension, capping rank at 2 of 3 under strict bipartiteness — a dimension-counting fact, not a fluke, confirmed numerically (nodal line, one exact zero eigenvalue). Breaking bipartite purity restores full rank 3 immediately, creating a direct tension with Addendum M.3's graph-level-bipartite reading of Postulate 8.
- Whether Postulate 8's Route A/Route B fork can be sidestepped entirely — rather than forced via internal argument, as BG's six negatives already tried — via topological solitons on Postulate 7's already-derived, fully-broken-SU(3) vacuum manifold (Addendum BD) — RESOLVED, no clean sidestep exists (Addendum CP). Gate 1 (does a site-dependent field theory for Postulate 7 even exist to build a soliton on) — DERIVED, yes: BD's on-site potential extended with a minimal nearest-neighbor coupling supports BD's attractor as a stable local vacuum (48→8 zero-mode collapse under coupling, zero negative modes, robust across regular/irregular/disconnected substrates). Gate 2 (does a hedgehog soliton on that field theory achieve finite-size stabilization) — NEGATIVE: standard Derrick collapse: the theory's existing quartic term is amplitude-type, not the derivative-type term genuine Skyrme stabilization requires, and no non-patching fix was identified.
- Whether CO's 3D negative result actually threatens Route B's viability, or only its specific 2-band test construction — RESOLVED, the latter (Addendum CQ): a genuine 4-component (Weyl-decomposed) representation — two components per sublattice, exactly what a real 3+1D Dirac fermion already requires physically — achieves full rank-3 non-degeneracy in 3D at near-universal rate (199/200 octonion-derived, 195/200 generic baseline), with zero violation of bipartite purity. Separately, whether this 2-component structure derives canonically from Postulate 1's octonion algebra itself (rather than importing Pauli matrices as an external tool) — NEGATIVE, and instructively so: 0 of 8 random G2 rotations preserved a fixed associative triple among the octonion basis's 7 genuine triples — recognized as AZ.13's own obstruction recurring one level up, not a new independent gap.
- Whether AZ.13's standing direction-selection question — open since v11, every prior

treatment (BC, BD, BG, Postulate 7 itself) simply taking a direction as given — admits any genuine mechanism at all — YES, a first candidate found (Addendum CR): a 2-vector "collapse" mechanism under a standard tachyonic-mass potential dynamically breaks G2 to exactly SU(3) (stabilizer dimension 8, confirmed two independent ways), and composes without interference with Postulate 7's own BD triad mechanism for a second stage (combined minimum energy exactly the sum of the two isolated minima; final stabilizer exactly 0, complete breaking G2→SU(3)→{1}). Honestly scoped against Addenda BD.5–BD.6's prior, permanent closure of the strong (symmetry-invariant) version of the same question: this is a genuinely different, weaker sense of "preferred" (vacuum selection, not unbroken symmetry) — a real advance, not a reopening. Honest limitation: the mechanism requires an unforced tachyonic mass sign that nothing in Postulates 1–3 currently forces — the same status the Standard Model's own Higgs potential sign has, an assumption, not a derivation.

6.14 Substantially advanced at v24 (Addenda CW–CZ) — carried from v24

- Whether Postulate 3's second-order dynamics carries a genuine stable-spiral signature (real $\kappa < -\eta^2/4$, otherwise stable — Addendum O's corrected target, superseding Addendum J's original complex-κ reading) anywhere other than the bare-cubic branch O.3–O.4 already searched and found empty — YES, robustly (Addendum CW): ≈99% of a 44,579-point stable-fixed-point sample on the triangle term's original 3-site toy case, nonlinearly confirmed to 4 decimal places, sharply contrasting with O's 0/197 and 0/3485 null results on the branch Addendum AY's theorem forecloses.
- Whether this generalizes beyond the one toy case — YES (Addendum CX–CY): complete (25/25) coverage on two independent, structurally distinct attractors on a real, independently-validated Watts-Strogatz network (Addendum BH/BN/CK's own published fixed point, exactly reproduced via literal source-code recovery), and complete (3000/3000) coverage on the theory's actual N=3,000 octonion-orbit substrate itself under the associator-norm coupling.
- Whether φ (Addendum CU's second non-patching generator) and R's (quintic) genuinely-stable branch show the same signature — YES, both (Addendum CZ): φ on the same real substrate matches the associator-norm's complete (3000/3000) coverage exactly; R's stable branch (Addendum S's A_+ root, rebuilt on an exactly-regular network) shows complete (25/25) coverage, nonlinearly confirmed. A logged correction (CZ.1) accompanies this result: Addendum CY.8 had incorrectly claimed Addendum AY forecloses R/S/T's dominant direction from ever being stable — AY covers the bare cubic only; R/S/T's own correction terms, per Addendum S's two-branch theorem, do genuinely stabilize a branch, and R/S/T's actual disqualification (Addenda CG/CI/CT) is on postulate-patching grounds, not instability.
- Honest synthesis (Addendum CZ.6): across six independent fixed points spanning three mechanisms and both toy and real substrates, the signature is complete or near-complete everywhere it has been checked — read now as a likely-generic property of η=0.8 second-order dynamics in this theory, not a lever for adjudicating Postulate 3's four-candidate underdetermination. Does not change any formal status tag.

6.15 Substantially advanced at v25 (Addenda DA–DI) — carried from v25

- φ vs. the associator-norm literal-coefficient question, pushed on five independent angles — φ proven the unique G2-invariant alternating trilinear form via direct $g2/\Lambda^3(\mathbb{R}^7)$ construction (DB.3–DB.5, a classification theorem); signed φ shown to fail the theory's own

generic-IC protocol outright, narrowing the viable class to {associator-norm, |φ|} (DD); DB.5's reach and the AZ.13/CR "economy of structure" argument both found narrower than first framed and corrected in place (DC, DE); the differentiability distinction between the two candidates shown real and dynamically consequential in general (DH–DI) but non-discriminating between them specifically, after due multi-substrate testing (DG, DI, correcting an initial single-substrate finding at DF). No formal status-tag changes.

6.16 Substantially advanced at v27 (Addenda DL–DR)

- Whether v25's flagged follow-ups on the associator-norm/|φ| kink-propagation finding hold up under closer scrutiny — YES, on all three counts (Addenda DL–DM): the by1 substrate reconstructed bit-exact from recovered primary-source code; a full N=3,000 dynamics scan confirmed genuine kink propagation on the theory's real canonical substrate for the first time; a 20-substrate batch (n=20, fully paired) confirmed DI's null with much greater power (Mann-Whitney p=0.250) — still no discriminating advantage between the two candidates. A striking natural near-degeneracy flagged in |φ|'s survivor set at DL.2 was chased to full resolution at DM: real and generic (confirmed via a structure-blind shuffle null across 10 substrates), but fully explained by a mundane triangle-count/degree coupon-collector confound once matched nulls are built — not a hint of deeper selection physics.
- Whether the associator is itself uniquely forced as a vector-valued object, the way DB.5 proved φ is as a scalar — YES (Addendum DN): via Casimir decomposition of $\Lambda^3(\mathbb{R}^7)$ under g2 (independently re-derived, matching DB.5's $1\oplus7\oplus27$ split exactly), the associator is shown to be the unique (up to scale) G2-equivariant alternating trilinear map into $\mathrm{Im}(\mathbb{O})$ — a companion theorem to DB.5's scalar-case result, closing the residual question of whether some other vector-valued candidate might exist.
- Whether a sharper postulate-patching argument could disqualify |φ| in favor of the associator-norm (is the absolute-value operation an ad hoc dynamical fix, where the associator-norm's non-negativity is "free"?) — NO (Addendum DN): both ||[a,b,c]|| and |φ| are literally the same canonical octonion norm map, restricted to different-rank objects (a 7-dimensional vector vs. the real subalgebra). Logged as a closed negative, not silently dropped.
- Whether a third literal-coefficient candidate exists that avoids the associator-norm/|φ| differentiability question entirely — YES: φ^2 (Addendum DO), smooth everywhere by construction, non-negative by construction, found dynamically viable and — bisected on three independent substrates — wider in relative discoverability window than both established candidates every time (roughly 1.7–1.9×).
- Whether φ^2's advantage survives a genuine stress test (larger N, seed-variance robustness) — YES, decisively (Addendum DP): a fourth substrate at double the prior largest N still shows φ^2 widest (4/4 substrates now); the gap between φ^2's mean relative width and the better established candidate's mean, at fixed substrate across 3 seeds, is roughly 10× the largest seed-to-seed noise observed — not a single-seed artifact.
- Whether φ^2's advantage can be explained analytically rather than left as an observed empirical pattern — YES, in two stages (Addenda DQ–DR). First (DQ): a single-triangle self-consistency closure gives an exact elementary ignition threshold; calibrated on by1 to two fixed percentiles (q_low≈92%, q_high≈58% of the per-node aggregated coupling distribution) and validated out-of-sample, unmodified, on three fresh substrates, landing within 10–25% of the true bisected values every time. Traces φ^2's advantage to its higher statistical skewness relative to the other two candidates. Second (DR): the empirical percentile calibration is superseded by a genuine zero-fitted-constant cascade theory — a self-consistent fixed-point iteration solving each node's actual cubic self-consistency

equation against its real specific triangle neighbors (closed-form via Cardano) — matching the true nonlinear dynamics to within 0.2–7.3 percentage points (vs. DQ's 10–30) on all four substrates × three candidates, and validated at the node level: 83–91% exact active-node overlap with the true ODE-integrated fixed point, with amplitudes on the shared active set matching to 3–4 significant figures. A structural (topology-only) percolation control confirmed the ≈58% threshold is a genuine property of the coupling-value distribution, not an artifact of graph connectivity (which percolates far earlier, ≈20–30%).

- Net effect on Postulate 3's status — UNCHANGED (POSTULATED, UNDERDETERMINED). The viable literal-coefficient class is revised from {associator-norm, $|\varphi|$} to {associator-norm, $|\varphi|$, φ^2} — a genuine addition, not a forced selection. φ^2 is now the most heavily evidenced and mechanistically understood of the three candidates on discoverability grounds, by a margin exceeding any prior candidate comparison in this project's history, without a forcing theorem existing for any of them.

6.17 Substantially advanced at v26 (Addenda DJ, DK)

- Whether a G2-invariant symmetric trilinear form exists on Im($\mathbb{O}$) at all, completing DB.5's alternating-sector classification — RESOLVED, no (Addendum DJ): the $\mathrm{Sym}^3(\mathbb{R}^7)^*$ representation (84-dimensional) has an exactly zero-dimensional G2-invariant subspace. Combined with DB.5, the entire G2-invariant trilinear content on Im($\mathbb{O}$) is the single form φ.
- Whether the g2-classification technique (DB.5, DJ) can be pointed at Postulate 8's Route A/Route B fork's specific-realization sub-question — YES, closing AZ.12/AZ.13's realization gap conditionally (Addendum DK.1–DK.9): given Route B and G2-equivariance, the spinor module is forced to be $S \cong V \oplus \mathbb{R}$ (the only viable 8-dimensional real G2-module), and the octonion-multiplication generator realization on it is forced, unique up to gauge and one discrete chirality choice — shown exactly to be octonion conjugation, an orientation reversal of Im($\mathbb{O}$). Does not touch BG's own six-for-six negative result on Route B's necessity.
- Whether AZ.13's "which 4 of 7" leftover-direction question can be closed by the same technique — YES, upgraded from an empirical search finding to a proof (Addendum DK.10–DK.12): su(3)=Stab_g2(n)'s commutant on the leftover 6-dimensional space is exactly 2-dimensional ($\cong \mathbb{C}$), admitting only trivial idempotents, so no nontrivial invariant subspace of any dimension exists there. Permanently forecloses future algebraic/symmetry-based selection attempts; does not name the four directions.
- Whether Addendum CS.3's tachyonic-sign coupling proposal survives scrutiny of its own functional form — TIGHTENED, legitimacy gate unchanged (Addendum DK.14–DK.17): a scope slip in CS.5 corrected (the coupling needs no site index); the space of G2-invariant v–M couplings classified as exactly 2-dimensional; CS.3's proposed form confirmed to be the natural point within it. CS.4's own gate — whether substrate data may enter CR's potential at all — remains the live blocker; AZ.13's $\gamma<0$ contingency (Postulate 7 sense) is unchanged.
- Whether the found chirality sign (DK.7) is the same operation as Postulate 6's "chirality-parity" leg's physical parity — CLOSED BY DISSOLUTION (Addendum DK.18): checked directly against Addenda AZ.2/BB.5's own text; no standalone parity operator, distinct from the pseudoscalar ω, exists anywhere in the record for it to be identified with. DK.7–DK.8's own content (C is exactly octonion conjugation, conjugating rather than negating chirality) stands unaffected.
- One correction logged: Addendum CS.5's stated reasoning against pursuing CS.3 ("no site index to sum over") does not survive direct comparison with CS.3's own formula

(Addendum DK.14) — logged per the project's no-silent-corrections standard, not silently worked around.

6.18 Substantially resolved at v20 (Addendum CF) — carried from v20

- The standing BT/BU/BV generic-IC parametrization discrepancy, open since Addendum BY.1 and this entry's single highest-leverage item across six addenda — RESOLVED (Addendum CF): BT.4's and BV.3's literal dynamics/IC code recovered via conversation search; cause identified as an initial-condition convention mismatch (generic ICs Uniform(0.3,0.8), not the Uniform(−0.3,0.3) every session since BY.1 had used); reproduced qualitatively on a reconstructed orbit substrate (4/4 seeds partial) and near-exactly on the verified-recovered lattice substrate (5/5 seeds, matching BV.3's own figures at two of five tested points); isolated decisively via a clean two-factor test (IC convention alone necessary and sufficient, dynamics order immaterial).
- Whether Addenda CA–CE's negative and mechanistic findings survive this resolution — RESOLVED, yes, unretracted (Addendum CF.6): they characterize a real, reproducible, honestly-adopted IC regime distinct from BT/BU/BV's own, not a series of failed reconstructions of it.
- Whether the conversation-search recovery technique (first demonstrated at v18, Addendum CD) generalizes to full dynamics/IC code, not only substrate construction — RESOLVED, yes (Addendum CF): the same route that recovered by1's construction code and BU/BV's substrate/coupling-scale convention also recovered BT.4's and BV.3's own integration code, now a three-times-demonstrated methodology.

6.19 Substantially resolved at v20 (documentation-and-consistency consolidation, no new addenda)

- The Addenda BM–BW archive gap, open since it was found and logged at v17 — REPAIRED (no new addendum letter; recovery applied to existing addenda). All thirteen addenda (BM.1–BM.5 through BW.1–BW.7) recovered from the original source session via conversation_search and re-inserted into SFT_Addenda_Archive.docx between Addendum BL and Addendum BX, using the same provenance-recovery technique demonstrated at Addenda CD, CE, and CF.
- Whether the initial BM–BW recovery pass was itself accurate — NO, PARTIALLY, CORRECTED. A second conversation_search pass found the first recovery had omitted BN.6 entirely and misnumbered BN.8's content as "BN.4." Corrected in place at Addendum BN's own header in the archive, logged as a correction rather than silently overwritten — the first documentation-propagation instance in the project's history whose own repair needed a follow-up correction.
- Whether SFT_Addenda_Archive.docx and SFT_Master_Document_v19.docx were genuine binary Word documents — NO, for both. Both had been saved as plain text with a .docx extension (the same bug previously found and fixed once before, per the v16 record, apparently silently reintroduced since). Both rebuilt as genuine binary docx and verified by PDF render.

6.20 Substantially resolved at v15 (Addenda BG–BL) — carried from v15

- Whether the triangle-term Postulate 3 candidate generalizes beyond its original toy case — RESOLVED YES, with real scope limits (Addendum BH).

- Whether Addendum BF's elimination of the octonion associator as h's origin survives given BH's finding — REVISED, not fully reversed (Addendum BI).
- The resulting scale factor left open by Addendum BI — CLOSED (Addendum BJ).
- Robustness of this result to generator choice — TESTED (Addendum BK): 7/10 draws succeed. Superseded at v16 by BO–BQ's fuller, out-of-sample-validated treatment.
- External audit of Addenda BH–BK — CONDUCTED AND ASSESSED (Addendum BL).

6.21 Retired — do not revisit unless new information emerges

- Original first-order gradient-flow master equation as fundamental — RETIRED, Part I.
- Unconditional SU(3) 'Grand Unification' claim — RETIRED, Part I, Defect 3.
- Linear LIV dispersion — RETIRED, superseded by the quadratic result.
- Wilson-term fermion doubling fix, domain-wall fermions, Majorana reality condition on Clifford generators, global Ramanujan-expander stability bonus, naive per-site Wegner bound, imported continuum Lifshitz-tail form, Gromov's theorem as a route to resolving adjacency-before-geometry, prior Γ_n growth-rate claim, real-Pauli-only Clifford ansatz, generic local curvature-driven graph rewiring, the standard σ_i>0 fixed point as a demonstration of emergent unitarity, continuation-based single-parameter tuning for Postulate 3 — all RETIRED or ELIMINATED as previously documented, unaffected by this consolidation.