

Deriving Closure-Coherence from the $K = 7$ Closure Algebra

Spectral Gap, Local Admissibility Filtering, and the Origin of E4a–E4c in VERSF

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

The previous paper introduced explicit refinement dynamics into the VERSF geometry programme. It showed that *if* refinement updates obey a small set of coherence conditions — called E4a, E4b, and E4c — then nearby substrate regions refine in stable, proportional ways. That stability propagates upward into cone stability and ultimately produces a smooth Lorentzian continuum geometry.

But those coherence conditions were still assumptions.

This paper derives E4a–E4c from explicit graph-theoretic and spectral-locality conditions naturally induced by the $K = 7$ closure algebra.

The central idea is that refinement in the VERSF substrate is not free to evolve arbitrarily. Each local closure state can transition only to a limited set of admissible children, and those admissible transitions form a finite constrained graph with cycles. The structure of this graph naturally filters inconsistent refinement paths. Nearby substrate regions therefore cannot drift apart chaotically under refinement, because incompatible updates are continuously rejected by the admissibility structure.

The paper introduces the **closure transition operator**

$$T : \mathcal{K} \rightarrow \mathcal{P}(\mathcal{K}),$$

which assigns to each closure state its admissible refinement children, and a **position-indexed admissibility filter**

$$T_{\text{coh}}^{\{\ell, x\}} : \mathcal{K} \rightarrow \mathcal{P}(\mathcal{K}),$$

which further restricts to children compatible with the local refinement neighbourhood.

The key claim is that repeated admissibility filtering produces a *spectral-gap-type contraction effect* on closure mismatch: refinement systematically suppresses incoherent divergence.

The main result derives the three dynamical coherence conditions of the previous paper —

- E4a — pointwise contractivity,
- E4b — position-coherence,
- E4c — uniform admissibility gap,

— in **two stages**.

First stage. E4a–E4c follow from six structural hypotheses (S0–S5) on the closure-transition graph. Three of these (S0 cyclic connectivity, S1 finiteness, S2 bounded branching) follow from the $K = 7$ wheel architecture: S1 and S2 from finiteness directly, and S0 — the *foundational topological precondition* — from the hexagonal closure structure of the $K = 7$ wheel (Lemma 5.0). The remaining three (S3 filter contractivity, S4 local-propagation Lipschitz response, S5 uniform spectral gap) are quantitative graph-theoretic properties.

Second stage. S3, S4, S5 are themselves *derivable* from more primitive spectral-locality conditions already motivated by the earlier $K = 7$ closure-Hamiltonian programme:

- **S3 (filter contractivity)** follows from **Nullity-1 of the paired closure response operator** M together with a **positive stiffness gap** $\Delta_M > 0$ above the unique null mode — the same spectral structure that powers the gauge-redundancy and persistence results elsewhere in the programme.
- **S4 (position-coherence)** follows from **finite-range locality** of the closure Hamiltonian together with **uniformly Lipschitz local compatibility penalties** — both already established in the closure-Hamiltonian programme.
- **S5 (uniform spectral gap)** follows from **finite-state spectral compactness** of the admissibility filter family: since the catalogue size is fixed at $K = 7$ and no admissibility-filter transition matrix carries an extra persistent mode beyond the unique coherent one, a uniform gap is automatic. ("No extra persistent modes" is exactly what Nullity-1 already asserts.)

The result moves the programme **two levels deeper**:

$K = 7$ closure spectrum + locality \Rightarrow S0–S5 \Rightarrow closure coherence (E4a–E4c) \Rightarrow $\alpha = 1 \Rightarrow$ Lipschitz continuum geometry.

This matters because it reduces the conditional content of the geometry programme in three stages — from abstract dynamical assumptions about refinement (E4a–E4c), to graph-theoretic structural properties of the closure-transition graph (S0–S5), to **primitive spectral-locality conditions already motivated by the $K = 7$ closure programme** (Nullity-1, stiffness gap, finite-range locality, Lipschitz penalties, finite-state compactness).

What remains open is the explicit numerical verification that the $K = 7$ closure catalogue satisfies these spectral-locality conditions — a finite combinatorial and linear-algebraic task on a fully specified algebraic object. The continuum geometry programme is no longer obstructed by abstract dynamical assumptions; it is obstructed only by an explicit catalogue construction and a small set of named external conditions, all of which are now reducible to finite computation on $K = 7$ closure-Hamiltonian data already in scope of the broader programme.

Contents

- [General-Reader Summary](#)
 - [Abstract](#)
 - [1. Introduction](#)
 - [1.1 Relation to Earlier Closure and Gauge-Structure Papers](#)
 - [2. The \$K = 7\$ Closure Transition Graph](#)
 - [3. The Graph-Derived Closure Metric \$\rho\$](#)
 - [4. The Position-Indexed Admissibility Filter](#)
 - [5. Structural Hypotheses on the \$K = 7\$ Algebra \(S0–S5\)](#)
 - [6. Theorem 1 — Pointwise Contractivity \(E4a\) from S2 + S3](#)
 - [6.4 Role of S0 in Theorem 1](#)
 - [6.5 Nullity-1 Interpretation of Contraction](#)
 - [7. Theorem 2 — Position-Coherence \(E4b\) from S4 + Finite Propagation](#)
 - [8. Theorem 3 — Uniform Admissibility Gap \(E4c\) from S5](#)
 - [9. Main Theorem — Closure Coherence from the \$K = 7\$ Algebra](#)
 - [10. Spectral-Gap Interpretation](#)
 - [11. Stability Flow: From Algebra to Lorentzian Geometry](#)
 - [12. Why This Matters](#)
 - [13. Refinement Dynamics as Substrate Renormalisation Flow](#)
 - [13.1 Entropic Spectral Contraction Principle](#)
 - [14. Limitations and Open Problems](#)
 - [15. Conclusion](#)
-

Abstract

The preceding paper, *Explicit Refinement Dynamics and Stability Flow in VERSF*, introduced explicit refinement evolution operators and proved that admissible refinement dynamics satisfying the closure-coherence sub-conditions E4a, E4b, E4c (together with E5, (D'), and ℓ -independence of the closure catalogue) generate bounded refinement distortion R4' with exponent $\alpha = 1$, yielding Lipschitz continuum cone stability.

However, E4a–E4c remained *conditional* dynamical hypotheses. The present paper derives them from explicit graph-theoretic and spectral-locality conditions naturally induced by the $K = 7$ closure algebra underlying admissible TPB refinement dynamics.

We introduce the **closure-transition graph**

$$G_{\mathcal{K}} = (\mathcal{K}, E),$$

whose vertices are admissible closure states and whose directed edges encode admissible refinement transitions. Refinement evolution is modelled through a position-indexed admissibility-filtered transition family

$$T_{\text{coh}}^{\ell, x} : \mathcal{K} \rightarrow \mathcal{P}(\mathcal{K}),$$

together with selectors that produce the position-indexed family $\{ \mathcal{E}_{\ell, x}^K \}$ of the previous paper.

We isolate six structural hypotheses on the closure-transition graph:

- **S0** — cyclic connectivity, $\beta_1(G_{\mathcal{K}}) \geq 1$ (foundational topological precondition);
- **S1** — finite admissibility graph (proven from $K = 7$ finiteness);
- **S2** — bounded branching (proven from $K = 7$ finiteness);
- **S3** — filter contractivity (conditional; load-bearing);
- **S4** — local-propagation Lipschitz response (conditional; load-bearing);
- **S5** — uniform admissibility spectral gap (conditional; load-bearing).

S0 is foundational — without cyclic connectivity, no admissibility filter can produce local separability, and S3–S5 have no work to do. S1 and S2 are automatic from the $K = 7$ catalogue being finite. S3–S5 are the load-bearing quantitative graph-theoretic conditions on $G_{\mathcal{K}}$.

Three theorems then derive:

- Theorem 1: $S2 + S3 \Rightarrow E4a$, with explicit constant $\beta = \beta_{\text{filter}} < 1$ (and with S0 ensuring S3 has non-vacuous content);
- Theorem 2: $S4 + \text{finite propagation} \Rightarrow E4b$, with explicit constant $L = \max(L_{\text{local}}, \rho_{\text{max}} / (2r))$;
- Theorem 3: $S5 \Rightarrow E4c$, with $\beta_{\ell} \leq 1 - \varepsilon_{\text{gap}}$.

The principal result is therefore

$$K = 7 \text{ closure algebra } (S0\text{--}S5) + \text{finite propagation} \Rightarrow E4a\text{--}E4c.$$

Combined with the previous stability-flow paper, this yields the composite implication

$K = 7$ closure algebra \Rightarrow Lipschitz refinement distortion \Rightarrow Lipschitz continuum cone stability \Rightarrow C^k Lorentzian metric realisations, conditional on completion of the E5 environment-register construction (§14), H9, and the continuum-limit hypotheses.

The paper therefore reduces the remaining dynamical assumptions in the VERSF geometry programme from refinement-dynamical conditions (E4a–E4c) to *checkable graph-theoretic properties* (S3–S5) of the explicit $K = 7$ closure algebra.

1. Introduction

The geometry branch of the VERSF programme has progressively shifted from continuum assumptions toward explicit substrate dynamics, in five identifiable stages.

Stage I. *Structural Necessity of Lorentzian Geometry in VERSF* — given a smooth strongly causal continuum with invariant causal cones and conserved transport density, Lorentzian geometry is forced.

Stage II. *Continuum-Limit Regularity and Cone Convergence in VERSF* — admissible TPB refinement sequences converge to a Lorentzian continuum under regularity hypotheses H6, H6', H7, H8, H9.

Stage III. *Deriving Regularity Hypotheses from Admissible TPB Dynamics in VERSF* — substrate-engineering constraints R1–R4' imply H6, H6', H7, H8.

Stage IV. *Explicit Refinement Dynamics and Stability Flow in VERSF* — explicit refinement evolution operators satisfying closure-coherence conditions E4a, E4b, E4c (together with E5, (D'), and ℓ -independence of the catalogue) imply R4' with $\alpha = 1$.

Stage V (this paper). Closure-coherence conditions E4a–E4c are derived from six structural hypotheses S0–S5 on the $K = 7$ closure-transition graph.

The Stage IV paper proved

$$(E4a + E4b + E4c + E5 + (D')) \Rightarrow \alpha = 1.$$

However, E4a–E4c themselves remained *conditional* dynamical assumptions. The present paper derives them from explicit graph-theoretic and spectral-locality conditions naturally induced by the $K = 7$ closure algebra.

The central claim is that admissibility filtering within a finite constrained closure graph naturally suppresses incoherent refinement divergence — and that the suppression rate is controlled by a spectral-gap-style quantity on the graph itself.

The paper therefore moves the programme from

assumed coherence \Rightarrow *Lipschitz geometry*

to

$K = 7$ *closure algebra* \Rightarrow *closure coherence* \Rightarrow *Lipschitz geometry*,

with the conditional content reduced from refinement-dynamical conditions to checkable graph-theoretic structural properties.

Epistemic status. The main theorem chain (Theorems 1, 2, 3) is *proven conditional* on the structural hypotheses S0–S5 plus finite propagation. S0 (cyclic connectivity, $\beta_1(G_{\mathcal{K}}) \geq 1$) is the foundational topological precondition — without it, no admissibility filter can produce local separability at all. S1 (finite catalogue) and S2 (bounded branching) follow automatically from $K = 7$ being a fixed finite catalogue. S3 (filter contractivity), S4 (local-propagation Lipschitz response), and S5 (uniform admissibility spectral gap) are the three *load-bearing quantitative* structural conditions of this paper. S0–S2 and S3–S5 are checkable graph-theoretic properties — once an explicit $K = 7$ catalogue construction is given, verifying S0 reduces to a Betti-number computation and verifying S3–S5 reduces to finite combinatorial computations on $G_{\mathcal{K}}$. The verification itself is an open programme (§14). The paper's contribution is the reduction: *if* S0–S5 hold for the $K = 7$ closure graph, *then* the dynamical coherence conditions E4a–E4c of the Stage IV paper follow automatically.

1.1 Relation to Earlier Closure and Gauge-Structure Papers

The present paper should not be read as introducing an unrelated refinement dynamics layered onto the earlier VERSF closure programme. Rather, it develops the *dynamical refinement consequences* of structures already identified in the $K = 7$ closure, gauge-redundancy, and cohomological-persistence papers. In particular:

- the **Nullity-1 structure** of the paired closure response operator (a unique persistent coherent mode compatible with closure and gauge redundancy, with all other directions stiffness-supported and therefore energetically suppressed);
- the emergence of **gauge redundancy from closure-preserving phase transport** (Maxwell admissibility and the $U(1)$ gauge structure on the $K = 7$ wheel);
- the **refinement persistence of transport sectors** (admissible transport closures survive coarse-graining as bookkeeping sectors of the substrate);
- the **Wilson-type coarse-grained transport structure** (effective continuum gauge actions emerging from BCB/TPB-admissible plaquette dynamics);

all suggest that admissible refinement dynamics should *suppress incoherent mismatch modes* while *preserving a unique coherent transport sector*. The contraction mechanism developed below is the dynamical refinement-flow manifestation of exactly this picture: incoherent modes are filtered out at each refinement step, leaving the persistent coherent sector intact and stable.

The closure-coherence conditions E4a–E4c of the Stage IV paper should therefore be interpreted not as arbitrary regularity assumptions, but as *refinement-flow manifestations of the underlying closure and transport structure* already established elsewhere in the programme. §6.5 below sharpens this into the Nullity-1 interpretation of pointwise contractivity, and §13 (new) interprets refinement evolution as a substrate-level renormalisation flow analogous to the Wilson construction in lattice gauge theory.

Master Structural Principle. The $K = 7$ substrate supports exactly one persistent coherent transport sector. Cyclic admissibility structure enables the trapping of incoherent alternatives, and admissibility filtering exports the associated entropy into inaccessible refinement sectors.

Continuum geometric regularity emerges because repeated refinement preserves only the coherent transport sector while exponentially suppressing all incoherent modes.

This single principle organises the rest of the paper. The Nullity-1 spectrum of the closure response operator (§6.5), the $\beta_1(G_{\mathcal{K}}) \geq 1$ condition on the closure-transition graph (§5, S0), the spectral gap ε_{gap} of the admissibility filter (§10), and the per-step $k_B \ln 2$ entropy export of the substrate-RG flow (§13) are all scale-resolved aspects of the same architectural commitment: one persistent coherent sector, all alternatives trapped, entropy exported to inaccessible registers, geometry stable under iterated refinement.

2. The $K = 7$ Closure Transition Graph

Let

$$\mathcal{K} = \{ \kappa_1, \dots, \kappa_K \}, K = 7,$$

denote the fixed finite closure catalogue introduced in the Stage IV paper, ℓ -independent across refinement levels.

Define the **closure-transition graph**

$$G_{\mathcal{K}} = (\mathcal{K}, E),$$

with directed edge set

$$E \subseteq \mathcal{K} \times \mathcal{K},$$

where $(\kappa_i, \kappa_j) \in E$ iff κ_j is an admissible refinement child of κ_i — that is, iff κ_j arises as a level- $(\ell+1)$ closure state in some position whose level- ℓ parent closure is κ_i , consistently with

- transport preservation (E2 of Stage IV),
- local compatibility constraints,
- finite propagation,
- admissible cone inheritance through Φ ,
- BCB preservation (E3 of Stage IV).

Define the **admissible-child set** as the out-neighbourhood

$$T(\kappa) = \{ \kappa' \in \mathcal{K} : (\kappa, \kappa') \in E \}.$$

Refinement evolution proceeds not through arbitrary closure transitions but through admissibility-filtered branching inside $G_{\mathcal{K}}$.

Physical content of $G_{\mathcal{K}}$. The closure-transition graph should not be interpreted as an abstract combinatorial construction detached from the earlier VERSF substrate programme. The vertices of $G_{\mathcal{K}}$ represent admissible local closure configurations compatible with

- **closure saturation** (BCB constraint balance at the local closure level);
- **paired transport structure** (the paired closure response operator and its Nullity-1 spectrum);
- **gauge redundancy** (closure-preserving phase transport, $U(1)$ emergence on the $K = 7$ wheel);
- **local admissibility** (the substrate's local-compatibility constraints derived in the Stage III paper);
- **refinement-persistent transport compatibility** (cohomological persistence of transport sectors under coarse-graining).

The edge structure of $G_{\mathcal{K}}$ therefore inherits its physical meaning from the earlier closure-Hamiltonian and transport-sector constructions: an edge (κ_i, κ_j) is admissible iff κ_j arises as a refinement child of κ_i in a way that preserves all of the above structures. Admissible refinement evolution corresponds to dynamically allowed transitions between locally closure-consistent substrate configurations.

This interpretation is important because the contraction mechanism derived below is *not imposed externally* on an arbitrary graph; it emerges from repeated admissibility filtering inside a *physically constrained* closure-transition structure. The closure programme's substrate primitives — closure saturation, paired transport, gauge redundancy, cohomological persistence — collectively determine which transitions $G_{\mathcal{K}}$ admits, and the contraction mechanism is the dynamical refinement-flow consequence of this physical structure.

Architectural remark. The $K = 7$ closure-transition graph is in principle a *fully specified combinatorial object*. It has at most $7^2 = 49$ directed edges (since $|\mathcal{K}| = 7$), including potential self-loops. Self-loops $(\kappa, \kappa) \in E$ correspond to refinement steps in which the closure state persists at the next level — these are admissible under the substrate primitives (they trivially satisfy the closure-saturation and transport-preservation conditions) and are counted in the 49-edge bound. $G_{\mathcal{K}}$'s structure is determined entirely by the admissibility relation derived from the substrate primitives above. An explicit enumeration — listing all 49 candidate edges and verifying, for each, whether the substrate primitives admit it — would render $G_{\mathcal{K}}$ fully explicit. The existence of such an explicit enumeration is what gives the present paper its leverage: the structural hypotheses S0–S5 below become finite combinatorial properties of an explicit graph, rather than abstract conditions on undetermined dynamics.

3. The Graph-Derived Closure Metric ρ

The Stage IV paper required a metric $\rho : \mathcal{K} \times \mathcal{K} \rightarrow \mathbb{R}_{\geq 0}$ on the closure catalogue, and (per the round-3 referee discussion) noted that ρ should be derived from the closure graph itself rather than chosen arbitrarily.

We make this concrete. Define

$$\rho : \mathcal{K} \times \mathcal{K} \rightarrow \mathbb{R}_{\geq 0}$$

as a graph-derived metric satisfying:

- **(M1) Symmetry and definiteness.** $\rho(\kappa, \kappa) = 0$, $\rho(\kappa, \kappa') > 0$ for $\kappa \neq \kappa'$, $\rho(\kappa, \kappa') = \rho(\kappa', \kappa)$.
- **(M2) Triangle inequality.** $\rho(\kappa, \kappa'') \leq \rho(\kappa, \kappa') + \rho(\kappa', \kappa'')$.
- **(M3) Bounded extremes.** $\rho_{\min} := \min_{\{\kappa \neq \kappa'\}} \rho(\kappa, \kappa')$, $\rho_{\max} := \max_{\{\kappa, \kappa'\}} \rho(\kappa, \kappa') < \infty$. Positivity $\rho_{\min} > 0$ follows automatically from (M1) on the finite catalogue \mathcal{K} (positive minimum of finitely many positive numbers); it is not an independent assumption.
- **(M4) Graph-derived structure.** ρ is computable from $G_{\mathcal{K}}$ alone — for example, as the symmetrised shortest-path distance on (\mathcal{K}, E) weighted by an admissibility-incompatibility weight, or as a Hamming distance over the admissibility-constraint signature of each closure state.

(M1)–(M3) are inherited from any metric structure; (M4) is what ties ρ to substrate structure rather than letting it be an arbitrary parameter.

Examples of admissible ρ -constructions.

(a) *Shortest-path distance.* For $(\kappa, \kappa') \in \mathcal{K} \times \mathcal{K}$, let $\rho(\kappa, \kappa')$ be the length of the shortest undirected path between κ and κ' in the symmetrised graph $(\mathcal{K}, E \cup E^{\text{op}})$. Then $\rho_{\min} = 1$ (for adjacent vertices) and $\rho_{\max} \leq 6$ (for the most distant vertex pair in a 7-vertex graph).

(b) *Admissibility-Hamming distance.* For each $\kappa \in \mathcal{K}$, let $\sigma(\kappa) \in \{0, 1\}^N$ denote the binary signature of which of N admissibility constraints κ satisfies. Define $\rho(\kappa, \kappa') = |\{i : \sigma(\kappa)_i \neq \sigma(\kappa')_i\}|$. Then $\rho_{\min} \geq 1$ (since distinct κ have distinct signatures by definition of admissibility) and $\rho_{\max} \leq N$.

The specific choice of construction may affect the numerical value of the constants β , L , ϵ_{gap} below but not the structural results. The Stage IV paper's $L_{\Phi} \leq d_{\{H, \max\}} / \rho_{\min}$ bound retains its content under either construction.

4. The Position-Indexed Admissibility Filter

The position-indexed family $\{\mathcal{E}_{\{\ell, x\}^K}\}_{x \in P_{\ell}}$ of the Stage IV paper is now realised explicitly via an **admissibility filter**.

For each level ℓ and substrate site $x \in P_{\ell}$, define the **coherent-child set**

$$T_{\text{coh}}^{\{\ell, x\}}(\kappa) \subseteq T(\kappa)$$

as the subset of admissible children that are compatible with the local refinement neighbourhood of x — that is, children whose closure data is consistent, under the substrate's local-compatibility constraints, with the closure data of substrate sites within the propagation radius r of x .

The position-indexed closure-update map $\mathcal{E}_{\{\ell, x\}^K} : \mathcal{K} \rightarrow \mathcal{K}$ of the Stage IV paper is then a single-valued **selector** from $T_{\text{coh}}^{\{\ell, x\}}$:

$$\mathcal{E}_{\{\ell, x\}^K}(\kappa) \in T_{\text{coh}}^{\{\ell, x\}}(\kappa).$$

We do not specify the selection rule — it may be deterministic (canonical-choice by graph ordering), data-dependent (chosen to optimise some local admissibility score), or stochastic (uniform sampling, with all results understood almost surely). The structural results below depend only on the admissibility filter $T_{\text{coh}}^{\{\ell, x\}}$, not on the selection rule.

Two further definitions. The **filter image** at position x of a closure pair (κ, κ') is the joint image

$$T_{\text{coh}}^{\{\ell, x\}}(\kappa) \times T_{\text{coh}}^{\{\ell, x\}}(\kappa'),$$

and its **diameter in ρ** is

$$\text{diam}_{\rho}(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, x\}}(\kappa')) := \max \{ \rho(\kappa_c, \kappa_{c'}) : \kappa_c \in T_{\text{coh}}^{\{\ell, x\}}(\kappa), \kappa_{c'} \in T_{\text{coh}}^{\{\ell, x\}}(\kappa') \}.$$

This is the largest possible mismatch between filtered children of κ and filtered children of κ' at position x . Pointwise contractivity (E4a) will follow from bounding this diameter by $\beta \cdot \rho(\kappa, \kappa')$.

5. Structural Hypotheses on the $K = 7$ Algebra (S0–S5)

We isolate six structural hypotheses on the closure-transition graph $G_{\mathcal{K}}$ and its admissibility filter. The first, S0, is the *foundational topological precondition*; without it, no admissibility filter can produce local separability, and S3–S5 below have no work to do.

S0 — Cyclic connectivity of $G_{\mathcal{K}}$ (foundational; topological precondition)

The first Betti number of the closure-transition graph satisfies

$$\beta_1(G_{\mathcal{K}}) \geq 1.$$

Equivalently, $G_{\mathcal{K}}$ contains at least one independent cycle.

Status of S0. S0 is foundational: it is the topological precondition without which admissibility filtering cannot produce local separability at all. The necessity is established by the topological-threshold result for fact formation: on any acyclic transition graph ($\beta_1 = 0$), no bounded local

modification of the dynamics can render previously recombinable alternatives non-recombinable while leaving the graph complement unchanged. This holds independently of how rich or contractive the admissibility filter is — without cyclic connectivity, the alternate-routing required for local trapping does not exist. Consequently, S3–S5 below are vacuous if S0 fails.

For the $K = 7$ closure algebra, S0 is essentially automatic. The first Betti number is a property of the underlying *undirected* graph (the cycle space): $G_{\mathcal{K}}$ is directed, so S0 refers to β_1 of the underlying undirected graph on \mathcal{K} , equivalently of $(\mathcal{K}, E \cup E^{\text{op}})$, which is what matters topologically for the alternate-routing argument of §1.1 and the Fact-Formation paper. For the undirected graph on 7 vertices with c connected components,

$$\beta_1 = |E_{\text{undir}}| - 7 + c.$$

Any *connected* closure-transition graph ($c = 1$) with $|E_{\text{undir}}| \geq 7$ therefore satisfies $\beta_1 \geq 1$ — well below the substrate's expected admissibility richness, since the $K = 7$ wheel architecture (Lemma 5.0 below) already supplies six boundary–boundary edges plus six boundary–hub edges, yielding $|E_{\text{undir}}| = 12$ and $\beta_1 = 12 - 7 + 1 = 6$ on the connected wheel — comfortably exceeding the threshold $\beta_1 \geq 1$. Concretely, the underlying undirected graph is the standard wheel graph W_6 on vertex set $V = \{ \kappa_{\text{h}}, \kappa_{\{b_1\}}, \dots, \kappa_{\{b_6\}} \}$ with edge set

$$E_{\text{undir}} = \{ \{ \kappa_{\text{h}}, \kappa_{\{b_i\}} \} : i = 1, \dots, 6 \} \cup \{ \{ \kappa_{\{b_i\}}, \kappa_{\{b_{i+1}\}} \} : i = 1, \dots, 6, \text{ indices mod } 6 \}.$$

(Lemma 5.0 thereby delivers not merely existence of a cycle but exactly *six independent cycles* on W_6 : the six boundary–boundary–hub triangles $\{ \kappa_{\text{h}}, \kappa_{\{b_i\}}, \kappa_{\{b_{i+1}\}} \}$ form a cycle-space basis, spanning the full 6-dimensional cycle space of W_6 .) S0 is listed as a separate hypothesis nonetheless, because the topological-threshold argument identifies it as the *load-bearing* structural condition for the very possibility of local separability — a condition that S1, S2 (mere finiteness and boundedness) do not entail.

Lemma 5.0 (Cyclic Connectivity from $K = 7$ Wheel Closure). If the $K = 7$ closure catalogue \mathcal{K} decomposes as six boundary states plus one hub state, and admissible transitions preserve boundary–hub closure compatibility, then the closure-transition graph $G_{\mathcal{K}}$ contains at least one independent cycle:

$$\beta_1(G_{\mathcal{K}}) \geq 1.$$

Hence S0 follows from the $K = 7$ wheel architecture.

Proof sketch. Label the $K = 7$ closure states as κ_{h} (hub) and $\kappa_{\{b_i\}}$ for $i = 1, \dots, 6$ (the six boundary states with cyclic indexing modulo 6). The closure-saturation condition at each boundary state $\kappa_{\{b_i\}}$ requires admissible compatibility with its two cyclic neighbours $\kappa_{\{b_{i\pm 1}\}}$ on the boundary ring — this is the *hexagonal closure commitment* of the $K = 7$ wheel architecture, established algebraically in the earlier closure-Hamiltonian programme (see the σ -duality wheel-construction paper and the BCB closure paper for the explicit construction). Each such compatibility induces an undirected edge $\{ \kappa_{\{b_i\}}, \kappa_{\{b_{i+1}\}} \} \in E_{\text{undir}}$. The

six boundary–boundary edges form an undirected 6-cycle on $\{ \kappa_{\{b_1\}}, \dots, \kappa_{\{b_6\}} \}$, contributing $\beta_i \geq 1$ to the underlying undirected graph regardless of the hub state's edge structure.

The hub state κ_h does not replace the boundary ring; it *anchors* it by providing the central closure-saturation condition with respect to which boundary compatibility is defined. A tree on \mathcal{K} would have only $|\mathcal{K}| - 1 = 6$ edges and no cycle, forcing boundary admissibility to reduce to a sequential dependence (a path) rather than the hexagonal closure-commitment ring — which is incompatible with the $K = 7$ wheel's intrinsic algebraic structure. ■

This lemma converts S0 from a structural hypothesis into a *derived consequence* of the $K = 7$ wheel architecture already established in the closure programme.

S1 — Finite admissibility graph (proven from $K = 7$)

$$|\mathcal{K}| = 7 < \infty, |E| \leq 49 < \infty.$$

The closure catalogue and its admissibility-transition edge set are finite. This is immediate from the $K = 7$ architecture.

S2 — Bounded branching (proven from $K = 7$)

$$\Delta_{\max} := \max_{\{\kappa \in \mathcal{K}\}} |T(\kappa)| \leq K = 7.$$

Every closure state has at most K admissible children. This follows from $T(\kappa) \subseteq \mathcal{K}$ and $|\mathcal{K}| = K$.

S3 — Filter contractivity (conditional; load-bearing)

There exists a constant $\beta_{\text{filter}} \in [0, 1)$ such that, for every level ℓ and position $x \in P_{\ell}$ and every $(\kappa, \kappa') \in \mathcal{K} \times \mathcal{K}$,

$$\text{diam}_{\rho}(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, x\}}(\kappa')) \leq \beta_{\text{filter}} \cdot \rho(\kappa, \kappa').$$

The admissibility filter restricts to children whose pairwise mismatch is *strictly smaller* than that of their parents, by a uniform factor $\beta_{\text{filter}} < 1$.

Status of S3. S3 is the principal quantitative structural hypothesis of this paper. Its physical content is that the admissibility filter is *consistency-preserving in mismatch*: incompatible refinements of mismatched parents are filtered out by local-compatibility constraints, leaving only children whose mismatch reflects the residual compatible alternatives. S3 presupposes S0 — without cyclic connectivity, the admissibility filter cannot produce diameter-strict-contraction on alternatives with support outside the modification ball. A derivation of S3 from explicit $K = 7$ catalogue topology is the open programme of §14.

Lemma 5.3 (Filter Contractivity from Nullity-1 and Stiffness Gap). Assume:

1. the paired closure response operator M on the local closure-state space has exactly one null mode ($\text{nullity}(M) = 1$), with the unique null direction corresponding to the admissible coherent transport mode;
2. all non-null modes of M have eigenvalue at least $\Delta_M > 0$ — a positive stiffness gap above the null mode;
3. the admissibility filter $T_{\text{coh}}^{\{\ell, x\}}$ removes refinement children whose mismatch energy exceeds a fixed local tolerance τ ;
4. the closure metric ρ is comparable to the quadratic mismatch energy induced by M — i.e., there exist constants $c_1, c_2 > 0$ with $c_1 \cdot \langle v, M v \rangle \leq \rho(\kappa, \kappa')^2 \leq c_2 \cdot \langle v, M v \rangle$ for v the local mismatch vector between κ and κ' .

Then there exist $c > 0$ and $\beta_{\text{filter}} < 1$ such that

$$\beta_{\text{filter}} \leq \sqrt{(c_2/c_1)} \cdot e^{-c \cdot \Delta_M}, \text{ and } \text{diam}_{\rho}(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, x\}}(\kappa')) \leq \beta_{\text{filter}} \cdot \rho(\kappa, \kappa'),$$

with $\beta_{\text{filter}} < 1$ requiring $\Delta_M > (1/2c) \cdot \log(c_2/c_1)$ — automatically satisfied when the metric comparability is tight (c_2/c_1 close to 1) and Δ_M is bounded below.

Hence S3 follows.

Proof sketch. Decompose any local mismatch vector v as $v = v_{\parallel} + v_{\perp}$, where v_{\parallel} lies in the rank-1 null space of M (the gauge-compatible coherent transport direction) and v_{\perp} lies in the stiffness-supported orthogonal complement. The null component v_{\parallel} corresponds to gauge-equivalent variation and does not contribute to physical closure mismatch — it is unobservable as a closure-mismatch direction. The orthogonal component v_{\perp} lies in the stiffness-supported subspace and carries quadratic mismatch energy

$$\langle v_{\perp}, M v_{\perp} \rangle \geq \Delta_M \cdot \|v_{\perp}\|^2.$$

Refinement-step contraction semigroup. Refinement evolution acts on the local mismatch vector by a one-step contraction semigroup $e^{-c \cdot M}$ for some refinement-step scale $c > 0$ set by the lattice spacing. On the rank-1 null subspace of M this is the identity (null modes are preserved by gauge); on the stiffness-supported complement, $e^{-c \cdot M}$ has operator norm at most $e^{-c \cdot \Delta_M} < 1$. Admissibility filtering (assumption 3) then enforces that surviving children lie in the image of this contraction — specifically, the mismatch-energy tolerance τ is set parent-relative, $\tau = \tau_0 \cdot \langle v_{\perp}^{\text{parent}}, M v_{\perp}^{\text{parent}} \rangle$ for some structural constant $\tau_0 \leq e^{-2c \cdot \Delta_M}$, so that surviving children satisfy

$$\langle v_{\perp}^{\text{child}}, M v_{\perp}^{\text{child}} \rangle \leq e^{-2c \cdot \Delta_M} \cdot \langle v_{\perp}^{\text{parent}}, M v_{\perp}^{\text{parent}} \rangle.$$

Transfer to ρ via two-sided comparability. By assumption 4 (two-sided metric comparability $c_1 \cdot \langle v, M v \rangle \leq \rho(\kappa, \kappa')^2 \leq c_2 \cdot \langle v, M v \rangle$), the parent ρ -distance lower-bounds parent energy and the child energy upper-bounds child ρ -distance:

$$\rho_{\text{child}}^2 \leq c_2 \cdot \langle v_{\perp}^{\text{child}}, M v_{\perp}^{\text{child}} \rangle \leq c_2 \cdot e^{\{-2c \cdot \Delta_M\}} \cdot \langle v_{\perp}^{\text{parent}}, M v_{\perp}^{\text{parent}} \rangle \leq (c_2 / c_1) \cdot e^{\{-2c \cdot \Delta_M\}} \cdot \rho_{\text{parent}}^2.$$

Taking square roots,

$$\rho_{\text{child}} \leq \sqrt{(c_2/c_1)} \cdot e^{\{-c \cdot \Delta_M\}} \cdot \rho_{\text{parent}}.$$

Setting $\beta_{\text{filter}} := \sqrt{(c_2/c_1)} \cdot e^{\{-c \cdot \Delta_M\}}$, the diameter bound follows. Since \mathcal{K} is finite ($|\mathcal{K}| = 7$), the supremum contraction factor over parent pairs is achieved on a finite set; it is strictly less than one whenever $\Delta_M > (1/2c) \cdot \log(c_2/c_1)$, which holds automatically for tight comparability and bounded-below Δ_M . ■

This lemma converts S3 from a structural hypothesis into a *conditional consequence* of the $K = 7$ closure-Hamiltonian spectrum: S3 holds whenever the paired closure response operator has Nullity-1 with a positive stiffness gap, the refinement-step evolution is realised as a contraction semigroup under M with step scale c , and the admissibility filter enforces parent-relative mismatch-energy tolerance — all conditions already motivated by the closure-Hamiltonian programme. The contraction rate β_{filter} acquires an explicit functional form $\beta_{\text{filter}} \leq \sqrt{(c_2/c_1)} \cdot e^{\{-c \cdot \Delta_M\}}$, tying the refinement-flow contraction rate directly to the static stiffness gap of the closure response operator (up to a $\sqrt{(c_2/c_1)}$ factor from metric comparability that is $O(1)$ for any reasonable ρ -construction).

S4 — Local-propagation Lipschitz response (conditional; load-bearing)

There exists a constant $L_{\text{local}} < \infty$ such that for every level ℓ , every pair $x, y \in P_{\ell}$ with $d_{\ell}(x, y) \leq 2r$ (where r is the finite-propagation radius), and every $\kappa \in \mathcal{K}$, the Hausdorff distance between filter images at x and at y satisfies

$$d_H(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, y\}}(\kappa)) \leq L_{\text{local}} \cdot d_{\ell}(x, y),$$

where d_H denotes Hausdorff distance with respect to ρ .

The filter image varies Lipschitz-continuously in substrate position over overlapping propagation neighbourhoods.

Status of S4. S4 is the position-coherence analogue of S3. It says that close substrate positions yield close filter images — a property determined by how local-compatibility constraints propagate across the substrate. A derivation from explicit substrate-neighbourhood structure is open (§14).

S5 — Uniform admissibility spectral gap (conditional; load-bearing)

There exists $\varepsilon_{\text{gap}} > 0$ such that, for every refinement level ℓ , the filter contractivity rate $\beta_{\text{filter}}(\ell)$ of S3 satisfies

$$\beta_{\text{filter}}(\ell) \leq 1 - \varepsilon_{\text{gap}}.$$

Equivalently, the admissibility-rejection rate of the closure-transition graph is bounded below by $\varepsilon_{\text{gap}} > 0$ uniformly across refinement levels. This is a *spectral-gap-style* condition on the transition operator on $G_{\mathcal{K}}$ (§10 develops the spectral interpretation).

Status of S5. S5 prevents β_{filter} from degenerating to 1 as the refinement mesh shrinks. Without S5, the geometric contraction of Theorem 1 would fail uniformly across levels, and the iterated bound of the Stage IV paper would not deliver continuum-limit Lipschitz stability. A derivation from explicit $K = 7$ graph spectral analysis is open (§10, §14).

Lemma 5.5 (Uniform Gap from Finite-State Compactness). Because $|\mathcal{K}| = 7$, the admissibility-filter transition matrices $\hat{T}^{\ell, \mathbf{x}}$ are 7×7 row-stochastic matrices. Assume:

1. the family $\{ \hat{T}^{\ell, \mathbf{x}} : \ell \in \mathbb{N}, \mathbf{x} \in P_{\ell} \}$ takes values in a finite set \mathcal{T} of such matrices, or more generally in a compact subset of the space of 7×7 row-stochastic matrices;
2. **(non-determinism on the boundary cycle, replacing primitivity)** for every $\hat{T}^{\ell, \mathbf{x}}$ in the family, there exists at least one boundary state $\kappa_{\mathbf{b}_i}$ (in the $K = 7$ wheel labelling of §5 Lemma 5.0) on which the filter image has cardinality at least 2:

$$|T_{\text{coh}}^{\ell, \mathbf{x}}(\kappa_{\mathbf{b}_i})| \geq 2;$$

3. consequently, no matrix in the family carries a non-trivial invariant subspace with eigenvalue of modulus one beyond the constant null mode — i.e., **no extra persistent modes** beyond the unique coherent one.

Then there exists $\varepsilon_{\text{gap}} > 0$ such that

$$\varepsilon_{\text{gap}} = \inf_{\hat{T} \in \mathcal{T}} (1 - |\lambda_2(\hat{T})|) > 0,$$

where λ_2 denotes the second-largest eigenvalue in absolute value (excluding the trivial eigenvalue 1 of constants). Hence S5 follows.

Tension with S0 and resolution. S0 (cyclic connectivity) and "no extra unit-modulus modes" (assumption 3) are in apparent tension. A pure rotation on the $K = 7$ boundary 6-cycle — each boundary state transitioning deterministically to the next $\kappa_{\mathbf{b}_i} \rightarrow \kappa_{\mathbf{b}_{i+1}}$ — would be irreducible but *not* primitive, and would carry all six 6th-roots of unity as eigenvalues, including $|\lambda| = 1$ for $\lambda \neq 1$. Such a matrix is consistent with S0 (the cycle is there) but fails assumption 3. Assumption 2 resolves this: admissibility filtering on the $K = 7$ wheel is *not* a pure deterministic rotation; the filter branches on at least one boundary state, mixing closure compatibility across the cycle rather than rotating rigidly. This branching destroys the deterministic 6-cycle eigenvalues other than 1, leaving 1 as the unique unit-modulus eigenvalue. Assumption 2 is therefore both *weaker than primitivity* (it does not require some power of \hat{T} to be strictly positive everywhere) and *sufficient* to break the unit-modulus eigenvalues that S0 alone would otherwise permit.

Hub-boundary admissibility (additional assumption). We additionally assume **hub-boundary admissibility**: for every $\hat{T}^{\ell, \mathbf{x}}$, the hub state $\kappa_{\mathbf{h}}$ is connected to at least one boundary state in

both directions of admissible transition. This makes $\hat{T}^{\{\ell, x\}}$ irreducible on all of \mathcal{K} (not just on the boundary cycle). The $K = 7$ wheel architecture supplies this by construction: the hub κ_h participates in the closure-saturation condition of every boundary state $\kappa_{\{b_i\}}$, so admissible transitions $\kappa_h \leftrightarrow \kappa_{\{b_i\}}$ are non-empty by design.

Proof sketch. By S0 and the $K = 7$ wheel architecture, together with hub–boundary admissibility, the stochastic matrix $\hat{T}^{\{\ell, x\}}$ is *irreducible* on all of \mathcal{K} : every closure state can be reached from every other via admissible transitions. By the Perron–Frobenius theorem for irreducible stochastic matrices, the eigenvalue 1 has multiplicity 1, all other eigenvalues lie in the closed unit disk, and the unit-modulus eigenvalues are exactly the h -th roots of unity, where $h \geq 1$ is the *period* of $\hat{T}^{\{\ell, x\}}$.

Assumption 2 forces *aperiodicity* ($h = 1$). The branching boundary state $\kappa_{\{b_i\}}$ with $|T_{\text{coh}}^{\{\ell, x\}}(\kappa_{\{b_i\}})| \geq 2$, combined with hub–boundary admissibility, ensures that the set of return-path lengths to $\kappa_{\{b_i\}}$ contains at least two coprime elements — for example, a 6-step return around the boundary cycle and a shorter hub-mediated return $\kappa_{\{b_i\}} \rightarrow \kappa_h \rightarrow \kappa_{\{b_j\}} \rightarrow \kappa_{\{b_i\}}$ of length 3 (for $\kappa_{\{b_j\}}$ reachable from the hub by hub–boundary admissibility and adjacent to $\kappa_{\{b_i\}}$ on the boundary), with $\text{gcd}(6, 3) = 3$; combined with a further hub-mediated return of length 2 ($\kappa_{\{b_i\}} \rightarrow \kappa_h \rightarrow \kappa_{\{b_i\}}$, available when the hub self-return is admissible) or 4, this gives a return-length set with overall gcd 1. The period $h = \text{gcd}(\text{return-lengths})$ is therefore 1: the chain is aperiodic. Aperiodicity collapses the unit-modulus spectrum to $\{1\}$, giving $|\lambda_2(\hat{T}^{\{\ell, x\}})| < 1$ strictly.

The eigenvalue map $\hat{T} \mapsto \lambda_2(\hat{T})$ is continuous on the compact space of 7×7 row-stochastic matrices (eigenvalues depend continuously on matrix entries, and reordering preserves continuity of the magnitude-ordered second eigenvalue). The image \mathcal{T} is either finite (with minimum gap achieved directly) or compact (with infimum attained by continuity on a compact domain). In either case, the infimum is strictly positive, since it is attained at some $\hat{T} \in \mathcal{T}$ with $|\lambda_2(\hat{T})| < 1$. Define ε_{gap} as this infimum. ■

This lemma converts S5 from a structural hypothesis into a *conditional consequence* of three minimal architectural conditions: (i) the admissibility-filter family is finite or compact — automatic when the filter is generated by a finite admissibility predicate on a finite catalogue; (ii) the filter is non-deterministic on at least one boundary state, which destroys deterministic-rotation roots-of-unity that S0 alone would otherwise permit; and (iii) consequently no extra persistent modes beyond Nullity-1 — *exactly* what the Nullity-1 spectral structure already asserts in the static description. The uniform spectral gap therefore reduces to the substrate's Nullity-1 architecture plus finite-state structure plus boundary-filter branching, all of which are established or plausibly verifiable consequences of the $K = 7$ closure programme.

6. Theorem 1 — Pointwise Contractivity (E4a) from S2 + S3

Statement. Assume S2 (bounded branching) and S3 (filter contractivity). Then for every selector

$$\mathcal{E}_{\{\ell, x\}^K}(\kappa) \in T_{\text{coh}}^{\{\ell, x\}}(\kappa),$$

the position-indexed family $\{ \mathcal{E}_{\{\ell, x\}^K} \}$ satisfies E4a:

$$\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, x\}^K}(\kappa')) \leq \beta_{\text{filter}} \cdot \rho(\kappa, \kappa'),$$

with $\beta = \beta_{\text{filter}} < 1$.

Proof. Fix ℓ, x , and $\kappa, \kappa' \in \mathcal{K}$. By definition of the selector,

$$\mathcal{E}_{\{\ell, x\}^K}(\kappa) \in T_{\text{coh}}^{\{\ell, x\}}(\kappa), \mathcal{E}_{\{\ell, x\}^K}(\kappa') \in T_{\text{coh}}^{\{\ell, x\}}(\kappa').$$

By definition of the filter-image diameter,

$$\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, x\}^K}(\kappa')) \leq \text{diam}_{\rho}(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, x\}}(\kappa')).$$

By S3,

$$\text{diam}_{\rho}(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, x\}}(\kappa')) \leq \beta_{\text{filter}} \cdot \rho(\kappa, \kappa').$$

Combining,

$$\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, x\}^K}(\kappa')) \leq \beta_{\text{filter}} \cdot \rho(\kappa, \kappa').$$

By S3, $\beta_{\text{filter}} < 1$, so this is E4a with $\beta = \beta_{\text{filter}}$. ■

Remark 6.1 (selector independence). Theorem 1 holds for *every* selector from $T_{\text{coh}}^{\{\ell, x\}}$ — deterministic, stochastic, data-dependent, or canonical. The contractivity property is a structural feature of the filter, not of any particular selection rule. This independence is what permits the Stage IV paper to leave the construction of the position-indexed family open: any selector from a contractive filter delivers E4a with the same constant.

Selector freedom as refinement-gauge redundancy. The selector freedom in the realisation $\mathcal{E}_{\{\ell, x\}^K} \in T_{\text{coh}}^{\{\ell, x\}}(\cdot)$ should be interpreted as the *refinement-dynamics analogue of closure-preserving gauge redundancy* identified in the Maxwell-admissibility programme. Different selectors correspond to different admissible representatives of the same coherent transport sector, and the contractivity result is selector-invariant in the same sense that physical observables are gauge-invariant: the persistent coherent mode (Nullity-1 direction of M , §6.5) is the gauge-invariant content, and individual selector choices are gauge representatives. Selector independence is therefore the *expected* structural feature — not a surprising one — once the connection to closure-preserving gauge redundancy is recognised.

Remark 6.2 (role of S2). S2 (bounded branching) does not appear explicitly in the proof above, but it is used implicitly: without bounded $|T(\kappa)|$, the filter diameter could in principle be infinite, and S3 would have no meaningful content. With $|T(\kappa)| \leq K$ finite for every κ , the diameter is a maximum over a finite set and S3 is a meaningful structural condition.

Remark 6.3 (admissible-pair scope). Following the Stage IV paper's "Scope of E4a" remark, the contractivity bound applies primarily on **admissibly comparable closure pairs** — pairs that can arise as locally compatible alternatives within a common refinement neighbourhood. Extending ρ to all of $\mathcal{K} \times \mathcal{K}$ with $\rho \geq \rho_{\min}$ on every distinct pair (per the Stage IV referee discussion) preserves the global form of E4a, but the physical content of S3 is contraction on the admissible sector.

6.4 Role of S0 in Theorem 1

S0 (cyclic connectivity, $\beta_1(G_{\mathcal{K}}) \geq 1$) does not appear as an explicit input to the algebraic step of Theorem 1's proof, but it is the *topological precondition* under which S3 has non-vacuous content. Specifically, if $\beta_1(G_{\mathcal{K}}) = 0$, then by the Fact-Formation topological-threshold theorem, admissibility filtering on alternatives with support outside any local modification ball *cannot* produce strict diameter contraction — every k-local modification preserves recombability for such alternatives, leaving the filter diameter equal to the parent ρ -distance. S3 then reduces to the trivial bound $\beta_{\text{filter}} = 1$, which is incompatible with the strict inequality $\beta_{\text{filter}} < 1$ required for Theorem 1. The dependency chain is therefore

$S0 \Rightarrow S3$ has non-vacuous content ($\beta_{\text{filter}} < 1$ is achievable) \Rightarrow Theorem 1 delivers strict contraction.

Equivalently: S3 is *satisfiable with the strict bound $\beta_{\text{filter}} < 1$ only if S0 holds*; without S0, S3 reduces to the trivial bound $\beta_{\text{filter}} = 1$, which is incompatible with the strict inequality required for Theorem 1. So although the present proof step does not name S0 explicitly, the assumption of strict $\beta_{\text{filter}} < 1$ in S3 is itself conditional on S0. This makes S0 a load-bearing input to Theorem 1 even though it appears implicitly through S3's non-vacuity.

6.5 Nullity-1 Interpretation of Contraction

The contraction mechanism above admits a natural interpretation in terms of the **Nullity-1 structure** established in the earlier $K = 7$ closure programme.

The paired closure response operator M studied in the closure-Hamiltonian papers satisfies

$$\text{nullity}(M) = 1,$$

corresponding to a *unique persistent coherent mode* compatible with closure and gauge redundancy. All remaining perturbative directions are stiffness-supported, and are therefore energetically suppressed in the substrate's static structure.

The present refinement contraction mechanism may be interpreted as the *dynamical refinement analogue* of this spectral structure:

- admissibility filtering at each refinement step **suppresses incoherent refinement modes** — exactly the modes that are stiffness-supported in the closure Hamiltonian;

- refinement evolution **preserves only the admissible coherent transport sector** — the unique persistent mode in the Nullity-1 picture;
- repeated refinement therefore **contracts closure mismatch away from the unique persistent coherent mode**, with the contraction rate set by the energetic separation between the persistent mode and the stiffness-suppressed perturbations.

In this interpretation, E4a is not merely an abstract contractivity assumption. It is the refinement-flow manifestation of the substrate's Nullity-1 coherence structure: the static energetic suppression of incoherent modes becomes a dynamical contraction of incoherent mismatch under iterated refinement.

Static–dynamical duality: $\text{nullity}(\mathbf{M}) = 1 \leftrightarrow \beta_1(\mathbf{G}_{\mathcal{K}}) \geq 1$. The Nullity-1 correspondence above is in fact one half of a structural duality between two scale-resolved descriptions of the same $K = 7$ architecture:

Static description (closure response operator \mathbf{M})	Dynamical description (closure-transition graph $\mathbf{G}_{\mathcal{K}}$)
$\text{nullity}(\mathbf{M}) = 1$	$\beta_1(\mathbf{G}_{\mathcal{K}}) \geq 1$ (hypothesis S0)
Unique null mode of \mathbf{M}	Persistent coherent transport sector of $\mathbf{G}_{\mathcal{K}}$
Stiffness-suppressed perturbations	Incoherent modes trapped in cycles of $\mathbf{G}_{\mathcal{K}}$
Spectral gap of \mathbf{M} above the null mode	Spectral gap ε_{gap} of $\hat{T}^{\{\ell, x\}}$ above the constant null mode (S5)
Energetic suppression	Topological inaccessibility (Fact-Formation §5.2)

Under this duality, the Nullity-1 spectrum of \mathbf{M} *lifts* to the cyclic transition structure of $\mathbf{G}_{\mathcal{K}}$: the rank-1 null space of \mathbf{M} becomes the persistent transport sector preserved by iterated refinement, and the stiffness-suppressed perturbations of \mathbf{M} become the cycles where incoherent modes get topologically quarantined. The condition $\beta_1(\mathbf{G}_{\mathcal{K}}) \geq 1$ (S0) is the dynamical-substrate analogue of $\text{nullity}(\mathbf{M}) \geq 1$ in the static-response description; both are the structural precondition for a unique coherent sector to survive at the next level of description.

The contraction rate β_{filter} of S3 — and, via §10, the spectral gap ε_{gap} — should therefore ultimately be derivable from the spectral gap between the null mode of \mathbf{M} and the lowest stiffness-supported mode, with the static *energetic* gap and the dynamical *spectral* gap of the transition operator related by a substrate-rescaling factor determined by the refinement step length. This makes the present S3 a refinement-dynamics counterpart of the Nullity-1 spectral analysis already established for \mathbf{M} .

The static–dynamical duality is conjectural at present, but it is the natural structural origin of S3 and S5, and is the preferred derivation target for future work (§14).

Stage VI framing. This duality, if established, would identify the Nullity-1 structure of the closure response operator and the cyclic structure of the closure-transition graph as two scale-resolved descriptions of a single underlying $K = 7$ admissibility architecture. The contraction rate β_{filter} and the static spectral gap of \mathbf{M} would then be related by a refinement-step rescaling,

providing the substrate-level origin of S3 and S5 in a unified spectral statement on M whose dynamical content is unpacked by the lift to $G_{\mathcal{K}}$. A precise formulation and proof — establishing the lifting map from M -eigenspaces to $G_{\mathcal{K}}$ -modes and computing the substrate-rescaling factor explicitly — is the principal Stage VI target. It would close the loop between the static closure-Hamiltonian programme and the dynamical refinement-flow programme: nullity(M) = 1 in the static picture *becomes* $\beta_1(G_{\mathcal{K}}) \geq 1$ in the dynamical picture under the substrate's natural scale-resolution map.

7. Theorem 2 — Position-Coherence (E4b) from S4 + Finite Propagation

Statement. Assume S4 (local-propagation Lipschitz response) with constant L_{local} , and finite propagation with radius $r > 0$. Then the position-indexed family $\{ \mathcal{E}_{\{\ell, x\}^K} \}$ satisfies E4b:

$$\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, y\}^K}(\kappa)) \leq L \cdot d_{\ell}(x, y),$$

with

$$L := \max(L_{\text{local}}, \rho_{\text{max}} / (2r)).$$

Proof. Fix $\ell, \kappa \in \mathcal{K}$, and $x, y \in P_{\ell}$. Two cases.

Case A: $d_{\ell}(x, y) \leq 2r$. The propagation neighbourhoods $N_r(x)$ and $N_r(y)$ overlap (their centres are within twice the radius). By finite propagation, $T_{\text{coh}}^{\{\ell, x\}}(\kappa)$ and $T_{\text{coh}}^{\{\ell, y\}}(\kappa)$ depend only on closure data inside $N_r(x)$ and $N_r(y)$ respectively, and their Hausdorff distance is bounded by S4:

$$d_H(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, y\}}(\kappa)) \leq L_{\text{local}} \cdot d_{\ell}(x, y).$$

Since the selectors lie in their respective filter images,

$$\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, y\}^K}(\kappa)) \leq d_H(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, y\}}(\kappa)) \leq L_{\text{local}} \cdot d_{\ell}(x, y).$$

Case B: $d_{\ell}(x, y) > 2r$. The propagation neighbourhoods are disjoint, and the filter images can in principle differ by the full catalogue diameter ρ_{max} . The trivial bound

$$\rho(\mathcal{E}_{\{\ell, x\}^K}(\kappa), \mathcal{E}_{\{\ell, y\}^K}(\kappa)) \leq \rho_{\text{max}}$$

then suffices, provided we express it as a Lipschitz bound in d_{ℓ} . Since $d_{\ell}(x, y) > 2r$,

$$\rho_{\text{max}} < \rho_{\text{max}} \cdot d_{\ell}(x, y) / (2r) = (\rho_{\text{max}} / (2r)) \cdot d_{\ell}(x, y).$$

Combining the two cases. Setting $L := \max(L_{\text{local}}, \rho_{\text{max}} / (2r))$,

$$\rho(\mathcal{E}_{\{\ell, x\}}^K(\kappa), \mathcal{E}_{\{\ell, y\}}^K(\kappa)) \leq L \cdot d_{\ell}(x, y)$$

uniformly in x, y . ■

Remark 7.1 (interpretation of the two cases). Case A captures the *local* coherence regime: nearby substrate sites have overlapping filter inputs, and their filter outputs differ in proportion to the local data perturbation. Case B captures the *distant* regime: substrate sites farther than $2r$ apart have independent filter inputs, and the best one can say is the trivial catalogue-diameter bound, rescaled by the distance threshold $2r$. The Lipschitz constant L is the *maximum* of the two regime-specific bounds, with the local Lipschitz constant L_{local} typically dominant when r is small relative to typical substrate distances.

Remark 7.2 (sharpness). For substrate models in which $L_{\text{local}} \cdot 2r < \rho_{\text{max}}$ — that is, when the local Lipschitz regime cannot reach the catalogue diameter within the propagation radius — the dominant contribution to L is $\rho_{\text{max}} / (2r)$. In this regime, *smaller r yields larger L* , reflecting that more finely localised propagation produces sharper jumps in filter image as one crosses the $2r$ threshold. The trade-off is structural: finite-propagation locality is in tension with global Lipschitz smoothness.

Remark 7.3 (locality from the closure-Hamiltonian programme). The position-coherence condition E4b is not an abstract regularity assumption but is *strongly supported by the locality structure already derived* in the closure-Hamiltonian programme. Earlier work in that programme established that

- **microscopic interactions are finite-range** in the substrate (closure-Hamiltonian terms couple only nearby admissible configurations);
- **effective long-range couplings decay exponentially** under coarse-graining (massive-channel suppression with characteristic correlation length ξ);
- **coarse-grained effective actions remain local under refinement** (the Wilson-type effective action retains its locality structure across RG steps).

Consequently, neighbouring refinement regions cannot evolve independently under admissible refinement dynamics. Overlapping refinement neighbourhoods necessarily share substantial admissibility structure and closure constraints, because the substrate's microscopic locality forces this. The Lipschitz-type position-coherence captured by E4b therefore inherits *directly from finite propagation and locality* of the closure substrate dynamics — both of which are established consequences of the closure-Hamiltonian programme rather than independent assumptions of the present paper. The constant L_{local} of S4 should ultimately be derivable from the exponential-decay rate of effective couplings, and the finite-propagation radius r from the lattice spacing of the underlying substrate.

8. Theorem 3 — Uniform Admissibility Gap (E4c) from S5

Statement. Assume S5 (uniform admissibility spectral gap) with constant $\varepsilon_{\text{gap}} > 0$. Then the contractivity rate $\beta_{\text{filter}}(\ell)$ from Theorem 1 satisfies

$$\beta_{\text{filter}}(\ell) \leq 1 - \varepsilon_{\text{gap}} \text{ for every } \ell,$$

which is E4c with $\beta = 1 - \varepsilon_{\text{gap}} < 1$.

Proof. Direct from S5. Theorem 1 establishes that the contractivity rate of $\{ \mathcal{E}_{\{\ell, x\}^K} \}$ at level ℓ equals $\beta_{\text{filter}}(\ell)$. S5 asserts the uniform upper bound $\beta_{\text{filter}}(\ell) \leq 1 - \varepsilon_{\text{gap}}$. Setting $\beta := 1 - \varepsilon_{\text{gap}} < 1$ yields the E4c form of the Stage IV paper. ■

Remark 8.1 (spectral-gap origin). S5 is presented here as a structural hypothesis, but §10 below develops a spectral interpretation in which ε_{gap} arises as the spectral gap of a transition operator on the closure-transition graph $G_{\mathcal{K}}$. Under this interpretation, S5 becomes the statement that the closure-graph transition operator has a strict spectral gap bounded below uniformly in ℓ .

Remark 8.2 (consistency with Stage IV E4c). The Stage IV paper stated E4c in the existence form "there exists $\beta \in [0, 1)$ such that $\beta_{\ell} \leq \beta$ for all ℓ ", explicitly to rule out the borderline $\beta_{\ell} \rightarrow 1$ with $\sup = 1$. Theorem 3 here delivers this in the strongest form: $\beta_{\text{filter}}(\ell) \leq 1 - \varepsilon_{\text{gap}}$ with $\varepsilon_{\text{gap}} > 0$ *bounded below uniformly*, so the contraction rate never approaches unity asymptotically. This is strictly stronger than the Stage IV E4c.

9. Main Theorem — Closure Coherence from the $K = 7$ Algebra

Theorem (Closure Coherence from Graph Structure). Let the $K = 7$ closure algebra satisfy the structural hypotheses S0–S5, and let the substrate refinement carry finite propagation with radius $r > 0$. Then the position-indexed family $\{ \mathcal{E}_{\{\ell, x\}^K} \}$ of the Stage IV paper, realised as any selector from the admissibility filter $T_{\text{coh}}^{\{\ell, x\}}$, satisfies E4a, E4b, and E4c, with explicit constants

$\beta = 1 - \varepsilon_{\text{gap}}$ (from S3 and S5), $L = \max(L_{\text{local}}, \rho_{\text{max}} / (2r))$ (from S4 and finite propagation).

Proof. Direct composition of Theorems 1, 2, and 3.

- Theorem 1 (S2 + S3) gives E4a with rate $\beta_{\text{filter}} < 1$.
- Theorem 3 (S5) sharpens this to $\beta_{\text{filter}} \leq 1 - \varepsilon_{\text{gap}}$ uniformly in ℓ .
- Theorem 2 (S4 + finite propagation) gives E4b with constant $L = \max(L_{\text{local}}, \rho_{\text{max}} / (2r))$.

E4c is the uniformity statement implicit in $\beta = 1 - \varepsilon_{\text{gap}}$ being independent of ℓ . ■

Corollary (Composite implication with Stage IV). Under S0–S5 + finite propagation + (the Stage IV hypotheses E1, E2, E3, E5, (D'), and ℓ -independence of $(\mathcal{K}, \rho, \Phi)$), the explicit refinement dynamics of the Stage IV paper satisfy bounded refinement distortion R4' with $\alpha = 1$, and the continuum cone field is Lipschitz with constant

$$K_\infty = L_\Phi \cdot L \cdot A^2 / (A_- \cdot (1 - \beta)) \text{ (Stage IV Theorem 3),}$$

and substituting $\beta = 1 - \varepsilon_{\text{gap}}$ (Theorem 3 of the present paper) and $L = \max(L_{\text{local}}, \rho_{\text{max}}/(2r))$ (Theorem 2 of the present paper):

$$K_\infty = L_\Phi \cdot \max(L_{\text{local}}, \rho_{\text{max}} / (2r)) \cdot A^2 / (A_- \cdot \varepsilon_{\text{gap}}).$$

The $1/\varepsilon_{\text{gap}}$ factor in the continuum constant therefore traces directly to the spectral gap of $\hat{T}^{\{\ell, x\}}$ via the Stage IV $(1 - \beta)$ denominator with $\beta = 1 - \varepsilon_{\text{gap}}$.

This makes the continuum Lipschitz constant entirely explicit in terms of (i) the closure-graph spectral gap ε_{gap} , (ii) the local-propagation Lipschitz constant L_{local} , (iii) the catalogue diameter ρ_{max} , (iv) the propagation radius r , (v) the closure-to-cone response constant L_Φ , and (vi) the two refinement-distance scalings A, A_- from the Stage IV (D').

Six measurable constants — three from the closure graph ($\varepsilon_{\text{gap}}, L_{\text{local}}, \rho_{\text{max}}$), one from the propagation kernel (r), one from the closure-to-cone map (L_Φ), and the (D') constants (A, A_-) — fully determine the continuum-geometry Lipschitz constant.

10. Spectral-Gap Interpretation

The constant ε_{gap} of S5 is presented above as an abstract structural condition. This section develops its interpretation as a *graph-theoretic spectral gap*, making the connection to standard spectral theory on finite graphs explicit and giving S5 a computable form.

10.1 The closure-transition operator

For each level ℓ and substrate site $x \in P_\ell$, the admissibility filter $T_{\text{coh}}^{\{\ell, x\}} : \mathcal{K} \rightarrow \mathcal{P}(\mathcal{K})$ induces a transition operator $\hat{T}^{\{\ell, x\}} : \mathbb{R}^{\mathcal{K}} \rightarrow \mathbb{R}^{\mathcal{K}}$ acting on real-valued functions on \mathcal{K} by uniform averaging over the filter image:

$$(\hat{T}^{\{\ell, x\}} f)(\kappa) = (1 / |T_{\text{coh}}^{\{\ell, x\}}(\kappa)|) \cdot \sum_{\{\kappa' \in T_{\text{coh}}^{\{\ell, x\}}(\kappa)\}} f(\kappa').$$

This is a stochastic operator: $\hat{T}^{\{\ell, x\}} \cdot 1 = 1$ (constants are preserved), and it has spectral radius 1 with the eigenvalue 1 attained by constants.

10.2 The spectral gap

Let

$\lambda_{-2}^{\ell, x} :=$ second-largest eigenvalue of $\hat{T}^{\ell, x}$ in absolute value

(taken over the orthogonal complement of the constants). The **spectral gap** at level ℓ and position x is

$$\varepsilon_{\text{gap}}^{\ell, x} := 1 - |\lambda_{-2}^{\ell, x}|.$$

The uniform spectral gap is

$$\varepsilon_{\text{gap}} := \inf_{\ell, x} \varepsilon_{\text{gap}}^{\ell, x}.$$

S5 is the assertion $\varepsilon_{\text{gap}} > 0$.

10.3 Spectral gap and filter contractivity

The connection to S3 (filter contractivity) is the following standard contraction estimate.

Proposition 10.1 (L²-distribution contraction from spectral gap). Let $\hat{T}^{\ell, x}$ be the transition operator above with spectral gap satisfying $\varepsilon_{\text{gap}}^{\ell, x} \geq \varepsilon > 0$, and let π_x denote its stationary measure. For each $\kappa \in \mathcal{K}$ let

$$\mu_{\{\kappa, x\}} := (1 / |\mathbf{T}_{\text{coh}}^{\ell, x}(\kappa)|) \cdot \mathbf{1}_{\{\mathbf{T}_{\text{coh}}^{\ell, x}(\kappa)\}}$$

denote the uniform measure on the filter image at position x . Then

$$\|\mu_{\{\kappa, x\}} - \mu_{\{\kappa', x\}}\|_{\{L^2(\pi_x)\}} \leq (1 - \varepsilon) \cdot \|\delta_{\kappa} - \delta_{\{\kappa'\}}\|_{\{L^2(\pi_x)\}},$$

where δ_{κ} is the indicator at κ . This is the standard contraction-in-distribution estimate for a stochastic operator with spectral gap ε .

Proof sketch. The map $\kappa \mapsto \mu_{\{\kappa, x\}}$ is one application of $\hat{T}^{\ell, x}$ to the indicator δ_{κ} . The difference $\mu_{\{\kappa, x\}} - \mu_{\{\kappa', x\}}$ lies in the orthogonal complement of constants in $L^2(\pi_x)$ — the subspace on which $\hat{T}^{\ell, x}$ has spectral radius at most $|\lambda_{-2}| = 1 - \varepsilon$. The bound follows. ■

Lemma 10.2 (ρ -diameter \leftrightarrow L²-distribution comparability). For ρ a graph-derived metric on a finite 7-state catalogue (per §3) and π_x a stationary measure of $\hat{T}^{\ell, x}$ bounded above and below uniformly in (ℓ, x) — i.e., $\pi_x(\kappa) \in [\pi_{\text{min}}, \pi_{\text{max}}]$ for all (ℓ, x, κ) with $\pi_{\text{min}} > 0$ — there exist constants $C_1, C_2 > 0$ (independent of (ℓ, x) but possibly dependent on the ρ -construction) such that

$$\text{diam}_{\rho}(\mathbf{T}_{\text{coh}}^{\ell, x}(\kappa), \mathbf{T}_{\text{coh}}^{\ell, x}(\kappa')) \leq C_1 \cdot \|\mu_{\{\kappa, x\}} - \mu_{\{\kappa', x\}}\|_{\{L^2(\pi_x)\}}, \|\delta_{\kappa} - \delta_{\{\kappa'\}}\|_{\{L^2(\pi_x)\}} \leq C_2 \cdot \rho(\kappa, \kappa').$$

Proof sketch. Both inequalities are finite-combinatorial: ρ is positive on the finite catalogue $\mathcal{K} \times \mathcal{K}$ with $\rho_{\min}, \rho_{\max} > 0$ (§3), and π_x is bounded above and below. C_1 is bounded by $\rho_{\max} \cdot \pi_{\min}^{-1/2}$, while C_2 is bounded by $\pi_{\max}^{1/2}$. The argument uses only finiteness and uniform boundedness of π_x ; it is not spectral. ■

Corollary 10.3 (S3 from spectral gap, via Proposition 10.1 + Lemma 10.2). Combining Proposition 10.1 and Lemma 10.2,

$$\text{diam}_\rho(T_{\text{coh}}^{\{\ell, x\}}(\kappa), T_{\text{coh}}^{\{\ell, x\}}(\kappa')) \leq (C_1 \cdot C_2) \cdot (1 - \varepsilon) \cdot \rho(\kappa, \kappa'),$$

so S3 holds with $\beta_{\text{filter}} \leq (C_1 \cdot C_2) \cdot (1 - \varepsilon_{\text{gap}})$, provided π_x is uniformly bounded.

Remark. This separates what spectral theory delivers rigorously (Proposition 10.1, an L^2 -contraction on distributions) from what requires finite-combinatorial work on the catalogue (Lemma 10.2, the $\rho \leftrightarrow L^2(\pi_x)$ comparability). The combined bound $\beta_{\text{filter}} \leq (C_1 \cdot C_2)(1 - \varepsilon_{\text{gap}})$ is the structural ancestor of the $\beta = 1 - \varepsilon_{\text{gap}}$ factor in the continuum-limit constant K_∞ of §9; the comparability constants $C_1 \cdot C_2$ are $O(1)$ for the shortest-path and admissibility-Hamming constructions of §3, but their explicit values depend on the ρ -construction and on π_{\min}, π_{\max} . Computing these constants from the explicit $K = 7$ catalogue is one of the verification tasks of §14.

10.4 Computational character of ε_{gap}

The spectral gap ε_{gap} is a *computable* quantity once the $K = 7$ closure-transition graph is specified explicitly. Concretely:

1. Enumerate all admissible edges $(\kappa_i, \kappa_j) \in E$ by checking the five admissibility conditions of §2.
2. For each candidate position-dependence pattern, construct the admissibility filter $T_{\text{coh}}^{\{\ell, x\}}$ and its transition matrix $\hat{T}^{\{\ell, x\}}$ (a 7×7 row-stochastic matrix).
3. Compute the second-largest eigenvalue in absolute value. The spectral gap is $1 - |\lambda_2|$.
4. Take the infimum over ℓ and x to obtain ε_{gap} .

This is a finite combinatorial computation. The hypothesis S5 ($\varepsilon_{\text{gap}} > 0$) becomes the *verifiable* statement that the resulting 7×7 stochastic operators have a uniform spectral gap. The "uniform admissibility gap" condition of the Stage IV paper, originally an abstract requirement, is reduced here to a standard property of finite stochastic matrices that can be checked directly.

Architectural remark. This is the deepest reduction the present paper achieves, but the uniformity in (ℓ, x) is still doing real work and the architecture has two regimes.

Position- and level-independent filter (simplest case). If $\hat{T}^{\{\ell, x\}}$ is the same 7×7 matrix for every (ℓ, x) — i.e., the admissibility filter is determined by the closure-graph structure alone, with no position or level dependence — then verification of S5 reduces to a *single* 7×7 eigenvalue computation. This is the cleanest realisation of the substrate-RG reduction: ε_{gap} is one number computed once.

Position- or level-dependent filter (general case). If $\hat{T}^{\ell, x}$ varies non-trivially with (ℓ, x) , then $\varepsilon_{\text{gap}} = \inf_{\ell, x} (1 - |\lambda_2(\hat{T}^{\ell, x})|)$ is the infimum over a parameter family. S5 then requires the additional structural fact that the spectral gap is *bounded below uniformly* in its parameters — i.e., that the family doesn't approach a degenerate matrix in the (ℓ, x) -parameter limit. This uniform-gap condition is itself a finite-dimensional matter when the parameter dependence is bounded (Lemma 5.5's finite-state compactness argument), but should be verified rather than assumed: one must check that the family $\{ \hat{T}^{\ell, x} \}$ is finite or compact and contains no limit points with $|\lambda_2| = 1$.

In either regime, the Stage IV E4c — a *dynamical* condition on the refinement evolution requiring infinite-level analysis — is reduced to a *spectral* condition on a finite stochastic operator (or a compact family of such operators) on 7 states. The conditional content of the geometry programme has been pushed from "verify dynamical contraction across all refinement scales" to "compute eigenvalues of a finite stochastic operator (family) on the $K = 7$ algebra".

11. Stability Flow: From Algebra to Lorentzian Geometry

The Main Theorem composes with the Stage IV chain to give a single substrate-to-geometry stability flow. With the new Lemmas 5.0, 5.3, 5.5 in place, the flow now extends back one further stage to the $K = 7$ closure-Hamiltonian spectrum:

($K = 7$ wheel closure: six boundary + one hub) \rightarrow (Lemma 5.0: $\beta_1(G_{\mathcal{K}}) \geq 1$, i.e. S0) \rightarrow (Nullity-1 of M + positive stiffness gap $\Delta_M > 0$) \rightarrow (Lemma 5.3: filter contractivity $\beta_{\text{filter}} \leq \sqrt{(c_2/c_1)} \cdot e^{-c \cdot \Delta_M}$, i.e. S3) \rightarrow (finite-state spectral compactness + no extra unit-modulus modes) \rightarrow (Lemma 5.5: uniform admissibility spectral gap $\varepsilon_{\text{gap}} > 0$, i.e. S5) \rightarrow (S0–S5: closure-graph structural conditions) + (finite propagation) \rightarrow (Theorems 1, 2, 3: E4a, E4b, E4c) \rightarrow (Stage IV Lemma 4.1: closure-coherence inequality) \rightarrow (Stage IV Theorem 1: Lipschitz closure stability) \rightarrow (Stage IV Theorem 2: Lipschitz cone stability) \rightarrow (Stage IV Theorem 3: Lipschitz refinement distortion, R4' with $\alpha = 1$) \rightarrow (Stage III: H8_Lip from R4' + H5 + R2 + bounded overlap) \rightarrow (Stage II: Lipschitz continuum cone field) \rightarrow (Stage I: C^k Lorentzian metric realisations, $k \geq 2$), conditional on completion of the environment-register construction for E5 (§14), H9, and the continuum-limit hypotheses.

Each arrow is a proven implication conditional on the named inputs. The flow has now reached the $K = 7$ closure-Hamiltonian spectrum at one end and Lorentzian geometry at the other, with no remaining abstract assumptions in the middle other than:

- E5 (cone-refinement compatibility) — now a bookkeeping identity once the environment-register construction is supplied (§14);
- H9 (continuum topology) — external to the substrate-dynamical chain;
- continuum-limit hypotheses — proven in Stage II under H6–H9.

All other conditional content has been reduced to:

- the $K = 7$ closure-Hamiltonian spectrum (Nullity-1, stiffness gap $\Delta_M > 0$) — established in the closure-Hamiltonian programme;
- finite-state spectral compactness — automatic for a finite catalogue with a finitely-generated admissibility filter family;
- finite propagation, BCB, bounded overlap — substrate-architectural conditions;
- (D') — refinement-distance compatibility, partly proven (upper half) and partly conditional (non-collapse lower half).

Strengthened claim chain. Compressing the flow above into its load-bearing arrows:

$K = 7$ wheel closure $\Rightarrow \beta_1(G_{\mathcal{K}}) \geq 1 \Rightarrow$ Nullity-1 + stiffness gap \Rightarrow S3, S5 \Rightarrow E4a, E4c.

(E4b lifts from S4, which descends from finite-range locality + Lipschitz penalties by an independent route.) The full continuum-geometry chain is now driven by the $K = 7$ closure-Hamiltonian spectrum together with substrate locality — the abstract regularity assumptions of earlier stages have been eliminated.

The remaining gap to a fully derived programme is therefore concentrated on (i) explicit verification that the $K = 7$ closure-Hamiltonian has the Nullity-1 + stiffness-gap structure (already in scope of the broader programme), (ii) explicit construction of the admissibility-filter family and verification of finite-state compactness, and (iii) completion of the E5 environment-register construction described in §14.

12. Why This Matters

This paper achieves the next structural reduction in the VERSF geometry programme.

The Stage IV paper showed

closure-coherent refinement dynamics \Rightarrow Lipschitz continuum cone stability.

The present paper shows

$K = 7$ closure-graph structural hypotheses (S0–S5) \Rightarrow closure-coherent refinement dynamics.

Composing the two,

$K = 7$ closure algebra (S0–S5) + Stage IV substrate primitives \Rightarrow Lipschitz continuum cone stability.

The conditional content of the geometry programme has been moved from abstract dynamical assumptions ($\beta < 1$ uniformly in ℓ , position-coherence with some L , etc.) to *graph-theoretic structural properties* of a finite, fully specified algebraic object — properties that can in principle be checked by explicit catalogue construction and finite combinatorial computation.

Three reductions in particular deserve emphasis.

Reduction 1: E4a \rightarrow S3 via filter selection. E4a was a condition on the *dynamical map* $\mathcal{E}_{\ell, x}^K$. S3 is a condition on the *admissibility filter* $T_{\text{coh}}^{\ell, x}$. Theorem 1 shows that *any* selector from a contractive filter satisfies E4a, eliminating the need to specify the selection rule.

Reduction 2: E4b \rightarrow S4 via propagation regime decomposition. E4b was a uniform Lipschitz position-coherence condition. S4 is the local-regime version, with the distant regime handled by the trivial catalogue-diameter bound. The two regimes combine to give a uniform $L = \max(L_{\text{local}}, \rho_{\text{max}}/(2r))$.

Reduction 3: E4c \rightarrow S5 \rightarrow spectral gap of 7×7 matrices. E4c was the uniform admissibility gap, an infinite-level dynamical condition. S5 is its graph-theoretic reformulation. Via Proposition 10.1, ε_{gap} becomes the spectral gap of a finite stochastic operator on 7 states — a quantity computable in $O(1)$ finite linear algebra.

The composite effect: the most delicate dynamical hypothesis of the Stage IV paper (uniform contraction rate across all refinement levels) has been reduced to a standard spectral-radius computation on a finite matrix. This is the level of concreteness that geometry programmes are expected to reach before they can be declared dynamically grounded.

13. Refinement Dynamics as Substrate Renormalisation Flow

The refinement dynamics developed in the present paper may be interpreted as a *substrate-level renormalisation process* rather than a purely combinatorial refinement procedure.

Earlier papers in the programme established that closure-preserving coarse-graining of the substrate generates

- **Wilson-type plaquette actions** on admissible substrate cells;
- **effective transport sectors** that persist across coarse-graining as bookkeeping closures;
- **continuum Maxwell / Yang–Mills structures** under admissible coarse-graining of the $= 7$ gauge sector.

The present refinement-flow framework extends this logic *geometrically*. Refinement evolution acts as a dynamical filtering process on admissible closure sectors:

- **incoherent refinement modes are suppressed** by the admissibility filter at each refinement step;
- **admissible transport sectors persist** through refinement, exactly as they persist through coarse-graining in the Wilson construction;

- **cone structure stabilises** as the filtered evolution converges to a fixed-point closure-assignment pattern up to vanishing refinement error;
- **continuum geometric regularity emerges** under repeated refinement, as the Lipschitz bound of Theorem 1 (and its lift to cone stability via the Stage IV chain) becomes the geometric manifestation of fixed-point stability.

In this interpretation, the emergence of Lipschitz continuum geometry is the *geometric analogue of RG stabilisation* in ordinary lattice field theory:

- in lattice gauge theory, repeated Wilson coarse-graining drives the effective action toward a continuum fixed point with controlled UV regularity;
- in the present framework, repeated admissibility-filtered refinement drives the closure-assignment pattern toward a Lipschitz-stable fixed point with controlled cone-field regularity.

The spectral gap ε_{gap} of S5 plays a role directly analogous to the *gap above the marginal operators* at a Wilson fixed point: it controls how rapidly irrelevant modes are suppressed under flow, and its strict positivity is what permits the continuum limit to be taken.

Quantitative content: entropy export per refinement step. The Wilson analogy can be sharpened from qualitative to quantitative by importing the topological-threshold result for fact formation. Each admissibility-filtering event at a refinement step that locally separates two previously recombinable alternatives is, by the topological-threshold construction, a bit-fixation event in the sense of the Fact-Formation paper. Such an event requires *exactly* $k_B \ln 2$ of entropy export, with the discarded alternative stored as correlation information in the auxiliary register E (Lemma 5.2 of the Fact-Formation paper).

Each refinement step of the substrate-RG flow therefore exports at least

$$\Delta S_{\text{step}} \geq N_{\text{bits}}(\ell) \cdot k_B \ln 2$$

of entropy, where $N_{\text{bits}}(\ell)$ is the number of bit-fixation events at level ℓ . The spectral gap ε_{gap} of S5 controls *how many* such events occur per refinement step: a larger gap means more aggressive filtering, more bits fixed, more entropy exported per RG step. The contraction rate β_{filter} and the entropy-export rate per step are therefore two aspects of the same underlying spectral structure of $\hat{T}^{\{\ell, \mathbf{x}\}}$.

This gives the substrate-RG flow a *thermodynamic accounting* directly: each refinement step is an entropy-producing event with a topologically determined minimum cost, and the continuum-limit construction is well-defined precisely because the cumulative entropy export over all refinement steps converges (a property guaranteed by the geometric decay of mismatch under iterated contraction). The spectral gap ε_{gap} thereby acquires an *entropic interpretation* beyond its purely contractive one: it sets the rate of irreversible information accumulation per RG step.

Epistemic status of the entropy-export accounting. The identification above — that each admissibility-filtering event satisfies the conditions of the Fact-Formation topological-threshold

construction *step-wise* at every refinement level — is *conjectural* at present. It presupposes (i) that $N_bits(\ell)$ is well-defined, i.e., every admissibility-filtering event that locally separates recombinable alternatives is identifiable and countable, and (ii) that the Fact-Formation topological-threshold result applies *step-wise* to admissibility filtering at each refinement level, not merely to the global construction. Both presuppositions require verification: (i) reduces to constructing an explicit accounting of which filter-rejection events constitute bit-fixation events versus pure within-sector rearrangement; (ii) requires checking that the filter's local separability structure matches the Fact-Formation hypotheses at each refinement step rather than only in aggregate. The entropy-export accounting above should therefore be read as the *structural target* for the substrate-RG flow's thermodynamic interpretation, with rigorous verification flagged as a Stage VI task (§14).

Architectural remark. This interpretation unifies the present paper's contraction mechanism with the earlier programme's Wilson coarse-graining results under a single substrate-RG umbrella. The Stage III–V chain can then be re-read as a *forward RG construction*: substrate-architectural constraints (Stage III) provide the bare action; explicit refinement dynamics (Stage IV) define the RG step; closure-graph spectral structure (Stage V) determines the spectral gap controlling the flow. Continuum Lorentzian geometry emerges as the IR fixed point of this flow, with topological fact-formation (Fact-Formation paper) supplying the per-step entropy accounting. This is a closer analogy to standard renormalisation than the original "refinement as combinatorial subdivision" reading, and it suggests that further programme structure may be importable from lattice-field-theory technology (block-spin transformations, decimation transformations, exact RG equations) into the VERSF refinement programme.

13.1 Entropic Spectral Contraction Principle

The spectral gap ε_gap , the entropy-export rate ΔS_step , the refinement contraction rate β_filter , and the continuum regularity constant K_∞ are not separate quantities. They are linked by one chain.

Step 1: spectral gap controls contraction. From Theorem 3 ($S5 \Rightarrow E4c$),

$$\beta_filter \leq 1 - \varepsilon_gap.$$

A larger spectral gap therefore yields a smaller contraction factor:

$$\varepsilon_gap \uparrow \Rightarrow \beta_filter \downarrow.$$

Step 2: contraction sets information removal per step. For a stochastic operator with leading non-trivial eigenvalue $1 - \varepsilon_gap$, the χ^2 -divergence of the iterated distribution to its stationary measure contracts at rate $(1 - \varepsilon_gap)^2$ per step (Poincaré inequality on finite Markov chains):

$$\chi^2(\mu_n \parallel \pi) \leq (1 - \varepsilon_gap)^{2n} \cdot \chi^2(\mu_0 \parallel \pi).$$

Relative entropy is bounded above by $\chi^2/2$ on finite state spaces (standard inequality), so the corresponding per-step relative-entropy decrement satisfies

$$\Delta D_{\text{step}} \gtrsim [1 - (1 - \varepsilon_{\text{gap}})^2] \cdot D_{\text{residual}} \approx 2 \varepsilon_{\text{gap}} \cdot D_{\text{residual}} + \mathcal{O}(\varepsilon_{\text{gap}}^2),$$

where $D_{\text{residual}} := D(\mu_n \parallel \pi)$ is the residual divergence-to-stationarity at the current step. The information removed from the accessible refinement sector per step — measured as the relative entropy decrement — therefore scales as

$$I_{\text{removed}} \approx \varepsilon_{\text{gap}} \cdot D_{\text{residual}},$$

linear in ε_{gap} to leading order. Larger spectral gap \Rightarrow larger per-step information removal.

Step 3: Fact-Formation accounting sets entropy export per step. By the Fact-Formation accounting (Lemma 5.2 there), each bit-fixation event stores $k_B \ln 2$ of correlation in the environment register E and exports an equivalent $k_B \ln 2$ of entropy from the accessible sector. (This is the Fact-Formation entropy-export accounting, not Landauer erasure: discarded alternatives are quarantined as correlation in E rather than destroyed; the $k_B \ln 2$ cost is identical in magnitude but tracks the correlation stored in E .) The entropy exported per refinement step therefore satisfies

$$\Delta S_{\text{step}} \geq k_B \cdot I_{\text{removed}} \approx k_B \cdot \varepsilon_{\text{gap}} \cdot D_{\text{residual}},$$

with the proportionality factor D_{residual} reflecting the current accessible-sector divergence from the stationary distribution. For the substrate-RG flow seeded at non-equilibrium ($D_{\text{residual}} = \mathcal{O}(1)$ per refinement step early in the flow), this gives the leading-order bound

$$\Delta S_{\text{step}} \gtrsim k_B \cdot \varepsilon_{\text{gap}}.$$

The chain so far:

$$\varepsilon_{\text{gap}} \uparrow \Rightarrow \Delta S_{\text{step}} \uparrow \Rightarrow \beta_{\text{filter}} \downarrow \Rightarrow \text{faster refinement contraction.}$$

Step 4: contraction sets continuum regularity. The same contraction rate controls the continuum Lipschitz constant derived in §9:

$$K_{\infty} = L_{\Phi} \cdot L \cdot A^2 / (A_{-} \cdot \varepsilon_{\text{gap}}).$$

A larger spectral gap therefore yields a smaller continuum regularity constant:

$$\varepsilon_{\text{gap}} \uparrow \Rightarrow K_{\infty} \downarrow.$$

Since a smaller K_{∞} means tighter Lipschitz control of the continuum cone field, the full chain closes:

Entropic Spectral Contraction Principle.

$$\varepsilon_{\text{gap}} \uparrow \Rightarrow \Delta S_{\text{step}} \uparrow \Rightarrow \beta_{\text{filter}} \downarrow \Rightarrow K_{\infty} \downarrow \Rightarrow \text{stronger continuum regularity.}$$

In words. A larger spectral gap means the admissibility filter more decisively separates coherent from incoherent refinement modes. That produces more irreversible information fixation per refinement step, exports more entropy into the inaccessible register, contracts closure mismatch faster, and leaves a smoother, more tightly controlled continuum cone field. The four quantities — spectral gap, entropy export rate, refinement contraction rate, continuum regularity constant — are *one quantity wearing four hats*, each appropriate to a different scale of the substrate-RG flow.

Programmatic content. The Entropic Spectral Contraction Principle is the quantitative form of the Master Structural Principle of §1.1. Where the Master Principle states architecturally that "continuum geometric regularity emerges because repeated refinement preserves only the coherent transport sector while exponentially suppressing all incoherent modes", the present Principle gives the *rate* at which suppression occurs, the *entropy cost* of that suppression, and the *resulting regularity constant* of the continuum limit — all controlled by the single quantity ε_{gap} . A $K = 7$ catalogue with a small spectral gap produces a geometry programme that just barely reaches Lipschitz regularity, while a $K = 7$ catalogue with a large spectral gap produces one with tight Lipschitz constants and rapid entropy export per RG step. Either is consistent with the architectural picture; the spectral gap of the $K = 7$ closure graph determines which.

Caveat on the small- ε_{gap} regime. The leading-order expansion $\Delta D_{\text{step}} \approx 2 \varepsilon_{\text{gap}} \cdot D_{\text{residual}}$ (Step 2) is valid only for ε_{gap} small. The $K = 7$ closure-Hamiltonian's Nullity-1 spectral structure with stiffness gap $\Delta_M > 0$ (Lemma 5.3) is expected to deliver an $\mathcal{O}(1)$ gap

$$\varepsilon_{\text{gap}} \approx 1 - \sqrt{(c_2/c_1) \cdot e^{-c \cdot \Delta_M}},$$

bounded away from 0 by an $\mathcal{O}(1)$ margin set by the $K = 7$ closure-Hamiltonian spectrum and the metric-comparability constants. In this regime the leading-order expansion is conservative, the actual per-step relative-entropy decrement is stronger than $2 \varepsilon_{\text{gap}} \cdot D_{\text{residual}}$, and the ΔS_{step} lower bound is correspondingly larger. Computing the actual $K = 7$ spectral gap — i.e., the constants c, c_1, c_2, Δ_M — is among the §14 verification tasks; the Principle above gives the structural relation between the spectral gap and the four downstream quantities regardless of where ε_{gap} turns out to sit on $[0, 1)$.

14. Limitations and Open Problems

Several issues remain.

Explicit $K = 7$ closure-transition graph. The principal open task is the explicit construction of $G_{\mathcal{K}}$ — enumerating its 49 candidate edges and verifying, for each, whether it satisfies the five admissibility conditions of §2. This is a finite combinatorial task. Once $G_{\mathcal{K}}$ is in hand, the verification of S_3, S_4, S_5 becomes a tractable computation.

Verification of S_0 from $G_{\mathcal{K}}$. Given the explicit $K = 7$ closure-transition graph, S_0 reduces to a Betti-number computation on the underlying *undirected* graph: count undirected edges $|E_{\text{undir}}|$,

count connected components c , and verify $\beta_1 = |E_{\text{undir}}| - 7 + c \geq 1$. This is the simplest of the structural verifications and is expected to pass for any substantively populated admissibility graph, in particular for the $K = 7$ wheel architecture invoked in Lemma 5.0 (where the boundary cycle alone supplies $\beta_1 \geq 1$).

Verification of S3 from $G_{\mathcal{K}}$. Given $G_{\mathcal{K}}$, S3 (filter contractivity) requires verifying that the admissibility filter $T_{\text{coh}}^{\{\ell, x\}}$ satisfies $\text{diam}_{\rho} \leq \beta_{\text{filter}} \cdot \rho$ uniformly. This reduces, under reasonable propagation models, to enumerating filter outputs across all parent pairs $(\kappa, \kappa') \in \mathcal{K} \times \mathcal{K}$ (at most 49 pairs) and computing β_{filter} as the supremum of the diameter-to-distance ratio. The comparability constants C_1, C_2 of Lemma 10.2 are additional finite computations on the chosen ρ -construction and stationary measure π_x .

Verification of S4 from substrate-neighbourhood structure. S4 (local-propagation Lipschitz response) requires knowing how the filter $T_{\text{coh}}^{\{\ell, x\}}$ responds to small perturbations of its substrate-neighbourhood input. This requires a model of the local-compatibility constraints — how filter membership depends on neighbouring closure data — which is partly determined by the substrate-architectural primitives of the Stage III paper.

Verification of S5 from spectral analysis. S5 (uniform spectral gap) reduces, via §10, to computing the second-largest eigenvalue of the 7×7 transition operators $\hat{T}^{\{\ell, x\}}$ and verifying it is bounded away from 1 uniformly in (ℓ, x) . For canonical filter constructions this is a finite eigenvalue computation; for position-dependent or stochastic filters, it may require a deeper uniform-spectral-gap argument across parameter space.

Explicit spectral construction. The present paper strongly suggests that the contraction constant β_{filter} and the coherence conditions E4a–E4c arise from the *spectral structure* of the $K = 7$ closure-transition operator itself. The natural next step is therefore a coordinated programme of

1. **explicit construction** of the closure-transition graph $G_{\mathcal{K}}$ (the 49-edge enumeration referenced above);
2. **computation of its graph Laplacian** and the associated transition operators $\hat{T}^{\{\ell, x\}}$;
3. **determination of admissibility spectra** — eigenvalues and eigenvectors of $\hat{T}^{\{\ell, x\}}$, with attention to the spectral gap above the constant null mode;
4. **identification of persistent null modes** — connection to the Nullity-1 structure of the closure response operator (§6.5), and verification that the null mode of $\hat{T}^{\{\ell, x\}}$ aligns with the persistent coherent transport sector identified in the closure-Hamiltonian programme;
5. **numerical measurement of refinement contraction rates** in explicit substrate simulations, with comparison to the spectral-gap prediction $\beta_{\text{filter}} \leq (C_1 \cdot C_2)(1 - \varepsilon_{\text{gap}})$ from Corollary 10.3.

Completion of this programme would convert the present structural contraction argument into an *explicit computational substrate theory*: the continuum Lipschitz constant K_{∞} of §9 would acquire a numerical value, the constants $(\beta, L, \varepsilon_{\text{gap}}, L_{\text{local}}, \rho_{\text{max}}, r)$ would all be measurable, and the substrate-to-geometry chain would become fully quantitative.

Choice of graph-derived ρ -construction. The shortest-path and admissibility-Hamming constructions of §3 are two natural candidates. Different constructions yield different numerical constants (C_1, C_2 in Proposition 10.1, ρ_{\min}, ρ_{\max} , etc.) but preserve the structural results. A canonical-choice argument — selecting the ρ -construction that yields the tightest β_{filter} or the cleanest spectral interpretation — is open.

Selector specification. Theorem 1 is independent of the selection rule from $T_{\text{coh}}^{\{\ell, x\}}$. A canonical or substrate-derived selection rule (e.g., maximum-likelihood under a local-energy functional, or canonical-choice by closure-graph ordering) would make the position-indexed family $\{\mathcal{E}_{\{\ell, x\}}^K\}$ fully specified rather than open.

E5 derivation via environment-register construction. The cone-refinement compatibility axiom E5 remains external to the present derivation chain. The present paper derives E4a–E4c; E5 still requires its own structural derivation. The Stage IV §13 discussion identified two derivation paths and flagged path (ii) — construct R from $\{\mathcal{E}_{\{\ell, x\}}^K\} + \Phi$ — as the architecturally preferred direction for VERSF.

The Fact-Formation paper supplies a concrete technical mechanism for executing path (ii). Its environment-register construction (Lemma 5.2 there) shows that any coarse-graining of bijective dynamics that merges alternatives admits a canonical lift to a fully reversible bijection $U : \Omega \times E \rightarrow \Omega \times E$ on an extended substrate $\Omega \times E$. Applied at each refinement step:

1. At each refinement step, the admissibility filter $T_{\text{coh}}^{\{\ell, x\}}$ produces an environment register E_{ℓ} — the discarded-alternative store guaranteed by the topological-threshold construction.
2. The refinement evolution \mathcal{E}_{ℓ} acts on substrate sites; the cone-refinement functor R can be *defined* as the action induced on the closure-generated cone field $C_{\ell} = \Phi \circ \chi_{\ell}$ by *projecting out the trapped sector* E_{ℓ} at each level.
3. Under this construction, R is no longer an independently-specified abstract functor but a *derived* quantity: $R := (\text{projection out of } E_{\ell}) \circ (\text{evolution by } \mathcal{E}_{\ell}) \circ \Phi$.
4. The cone-refinement compatibility identity $R \circ C_{\ell} = C_{\{\ell+1\}} \circ \mathcal{E}_{\ell}^{\wedge P}$ is then *automatic by construction*, with the error term $\delta_{\ell} \rightarrow 0$ inherited from the geometric decay of accessible-trapped overlap under iterated admissibility filtering (the same geometric decay that powers Theorem 1). Specifically, by Theorem 1, the accessible-trapped overlap contracts at rate $\beta_{\text{filter}} \leq 1 - \varepsilon_{\text{gap}}$ per refinement step, so

$$\delta_{\ell} \leq \delta_0 \cdot (1 - \varepsilon_{\text{gap}})^{\ell} \rightarrow 0 \text{ geometrically as } \ell \rightarrow \infty,$$

with rate set by the same ε_{gap} that controls every other quantity in the Entropic Spectral Contraction chain of §13.1. The remaining Stage VI task is to verify that this geometric decay rate is *sufficient* for the continuum-limit construction's required convergence rate — a comparison between the substrate-side rate $(1 - \varepsilon_{\text{gap}})^{\ell}$ and the continuum-side regularity demands inherited from Stage II.

Thus E5 is not an independent physical postulate once the environment-register construction is supplied. It becomes a **bookkeeping identity**: the abstract cone-refinement functor R is the

induced action of substrate refinement on cone fields after discarded alternatives have been stored in E_ℓ . There is no separate dynamical content in $E5$ beyond the choice to track refinement at the level of (accessible) cone fields rather than at the level of the full $(\Omega \times E)$ -evolution; the identity holds by definition of the projection-out operation.

This is a fully substrate-first construction of R : the closure-transition dynamics \mathcal{E} are primary; the abstract refinement functor R is a derived bookkeeping device that records what \mathcal{E} does to cone fields after the trapped sector is projected out. $E5$ then holds by construction, not by assumption. Completing this construction — making the projection step explicit, verifying the $\delta_\ell \rightarrow 0$ rate, and showing that the resulting R agrees with the abstract refinement functor of the continuum-limit chain — is the natural Stage VI task that the present paper sets up.

Numerical refinement simulations. Direct numerical simulation of admissibility-filtered refinement on explicit $K = 7$ catalogues should compute β_{filter} , L_{local} , ε_{gap} , and verify the predicted Lipschitz scaling of closure-mismatch decay. Such simulations would also test the sharpness of the (β, L) constants and the validity of the $L^2 \leftrightarrow \rho$ comparability assumed in Proposition 10.1.

Generation-resolved structure. The $K = 7$ closure algebra has internal structure tied to the $K = 7$ wheel architecture used elsewhere in VERSF (notably in the σ -duality programme). A generation-resolved version of $S3$ – $S5$ — in which the spectral gap ε_{gap} and contraction rate β_{filter} carry generation labels reflecting the wheel structure — would tie the present paper to the broader substrate-dynamics programme more tightly.

Tightness of the $L = \max(L_{\text{local}}, \rho_{\text{max}}/(2r))$ bound. The Lipschitz constant L of Theorem 2 is a max over two regime-specific bounds. Whether sharper estimates are possible — for example, by interpolating between local Lipschitz response and trivial catalogue-diameter bound — is open.

15. Conclusion

This paper derives closure coherence from explicit graph-theoretic and spectral-locality conditions naturally induced by the $K = 7$ closure algebra underlying admissible TPB refinement dynamics.

The principal claim is that finite admissibility filtering on a closure-transition graph with bounded spectral content naturally suppresses incoherent refinement divergence. The conditional content of the geometry programme has been moved from abstract dynamical assumptions to checkable graph-theoretic structural properties.

The main new objects are

- the closure-transition graph $G_{\mathcal{K}} = (\mathcal{K}, E)$;
- the graph-derived closure metric ρ ;

- the position-indexed admissibility filter $T_{\text{coh}}^{\{\ell, x\}}$;
- the structural hypotheses S0 (cyclic connectivity, $\beta_1(G_{\mathcal{K}}) \geq 1$, foundational topological precondition), S1, S2 (proven from $K = 7$ finiteness) and S3, S4, S5 (load-bearing quantitative graph-theoretic conditions);
- the closure-transition operator $\hat{T}^{\{\ell, x\}}$ and its spectral gap ε_{gap} .

The principal result is that under S0–S5 + finite propagation, the position-indexed family $\{\mathcal{E}_{\ell, x}^K\}$ of the Stage IV paper satisfies E4a, E4b, E4c with explicit constants

$$\beta = 1 - \varepsilon_{\text{gap}}, L = \max(L_{\text{local}}, \rho_{\text{max}} / (2r)).$$

Composed with the Stage IV result, this gives

$K = 7$ closure algebra (S0–S5) + Stage IV substrate primitives $\Rightarrow R4'$ with $\alpha = 1 \Rightarrow$ Lipschitz continuum cone stability $\Rightarrow C^k$ Lorentzian metric realisations, conditional on completion of the E5 environment-register construction (§14), H9, and the continuum-limit hypotheses.

The continuum Lipschitz constant becomes entirely explicit in terms of six measurable substrate primitives:

$$K_{\infty} = L_{\Phi} \cdot \max(L_{\text{local}}, \rho_{\text{max}} / (2r)) \cdot A^2 / (A_{-} \cdot \varepsilon_{\text{gap}}).$$

The geometry programme has thus moved one stage deeper: from assumed regularity (Stage II), to substrate-engineering constraints (Stage III), to explicit refinement flow (Stage IV), to graph-theoretic structural conditions on the closure algebra itself (Stage V). The remaining open content is now concentrated on (i) explicit construction of the $K = 7$ closure-transition graph and verification of S3–S5, (ii) completion of the E5 environment-register construction (§14), making the projection step explicit and verifying that the $\delta_{\ell} \rightarrow 0$ rate is sufficient for the continuum-limit construction, and (iii) verification of H9 and the non-collapse half of (D') for the chosen substrate model.

In each case, the open content is concrete and reducible to finite computation or to derivation from established substrate primitives. The path from $K = 7$ closure algebra to Lorentzian continuum geometry is no longer obstructed by abstract dynamical assumptions; it is obstructed only by the explicit catalogue construction and a small set of named external conditions.

Taken together, the present paper suggests that continuum geometric regularity may ultimately arise from three interconnected substrate principles, composed in a single chain:

closure \Rightarrow gauge redundancy / Nullity-1 persistence \Rightarrow admissible transport coherence \Rightarrow refinement contraction \Rightarrow Lipschitz continuum geometry.

The refinement programme therefore no longer appears disconnected from the earlier $K = 7$ and closure-Hamiltonian constructions. Instead, the continuum geometry programme increasingly appears as the *large-scale refinement-stability manifestation* of the same underlying closure and admissibility structures that govern the substrate itself. The Nullity-1 spectral structure of the

paired closure response operator, the gauge redundancy emerging from closure-preserving phase transport, the cohomological persistence of admissible transport sectors, and the Wilson-type coarse-grained transport structure all reappear in the present paper as the substrate-structural origin of the contraction mechanism that makes Lorentzian continuum geometry possible. The various branches of the VERSF programme are converging on a single architecture in which closure structure, gauge structure, and continuum geometry are different scale-resolved manifestations of the same $K = 7$ admissibility algebra.