Admissible Coarse-Graining and Continuum Emergence in VERSF

Refinement Stability, Causal-Partition Persistence, and the Universality-Class Origin of Continuum Observables

Keith Taylor VERSF Theoretical Physics Programme

Plain-Language Summary

Modern physics describes the world using smooth space, continuous time, and fields spread across spacetime. Many approaches to quantum gravity suggest this smooth picture isn't fundamental — that underneath, reality is built from countless tiny indivisible events rather than from a smooth continuum.

If that's true, an immediate puzzle appears. Why doesn't the world look discrete? Why do the things we measure appear smooth, continuous, and largely indifferent to whatever microscopic details lie beneath?

The answer this paper proposes is that the smooth world isn't built directly from the underlying discrete pieces. It consists of the *patterns* in those pieces that survive when you zoom out — features that stay stable however much you look at finer and finer details. The microscopic mess averages away; certain large-scale patterns persist.

To make this precise we need a clear procedure for zooming, a way to identify which patterns survive and which fade, and a careful account of which conservation laws still hold after zooming and which need slight modification. We work all of this out explicitly on the simplest non-trivial example — a four-point "causal diamond" — and show which patterns are stable, which fade, and at what rate. The diamond turns out to support only one stable pattern (a counting invariant): the spacelike and diagonal modes decay, and the most oscillatory pattern — what we call the "corner-alternating" mode — is wiped out *exactly in one step* rather than fading gradually. The same construction generalises cleanly to higher-dimensional substrates, and we prove that for *every* dimension $d \geq 1$ the same picture holds: one stable counting invariant, the corner-alternating mode exactly extinguished in one step, and a clean hierarchy of intermediate decay rates with binomial-coefficient structure. The narrowness of what survives is a strong constraint, and it points away from bulk substrate structure as the carrier of continuum physics and toward boundary-like "interfaces," which is precisely where the wider VERSF programme locates the geometry that fixes physical constants.

We are careful throughout to separate three kinds of claim: things we have proved, things we have proved given stated assumptions, and things we currently only conjecture. The most important open question is whether the surviving patterns all live on certain "interfaces" in the substrate — sheets or boundaries rather than bulk regions. If so, this work would connect to other parts of the VERSF programme that derive specific physical constants from interface geometry.

We do not claim to derive General Relativity, quantum mechanics, or the Standard Model. The narrower claim is that the smooth continuous world isn't a fundamental ingredient — it's the equivalence class of underlying discrete structures whose large-scale features survive zooming. Two very different microscopic substrates can describe the same physical world, provided what survives the zooming is the same.

Abstract

We develop a constructive theory of admissible coarse-graining within the Void Energy–Regulated Space Framework (VERSF). The framework treats physical reality as emerging from discrete irreversible commitment events organised by admissibility constraints rather than from a fundamentally continuous spacetime manifold.

We define refinement maps $\Gamma : \Lambda_n \to \Lambda_{n+1}$ preserving (i) the substrate partial order, (ii) admissibility closure $A[\rho] = 0$ in the appropriate refinement-flow sense, and (iii) operational distinguishability grounded in a substrate-level resolution functional. We construct Γ explicitly on a finite toy substrate (the causal diamond with iterated midpoint subdivision). The BCB constraint is shown to be monotonically non-increasing under refinement, converging to a refinement-invariant limit $\tilde{Q}_\infty$; this is the substrate-level analogue of a renormalisation-group flow toward a fixed-point coupling, rather than a step-by-step invariance.

The central technical result is **Theorem 1 (Refinement-Stability and Decay on the d-Dimensional Causal Diamond)**: defining the linearised RG operator $L_R^{(d)} = (1/(2d))(d \cdot I + A_{Q_d})$ as the lazy random walk on the covering-pair graph (which we identify as the d-cube graph Q_d), the spectrum on the d-dimensional causal diamond is $\{1, (d-1)/d, (d-2)/d, \ldots, 0\}$ with binomial multiplicities $\{C(d,0), C(d,1), \ldots, C(d,d)\}$, for every $d \geq 1$. The eigenvalue-1 subspace is one-dimensional in every d and is spanned by the constant mode; the corner-alternating mode is in the kernel of $L_R^{(d)}$ in every d — exact extinction in one step, independent of dimension. The $d = 4$ spectrum is explicitly $\{1, 3/4, 1/2, 1/4, 0\}$ with multiplicities $\{1, 4, 6, 4, 1\}$.

We read this not as a deflation of the framework but as a structural argument that *non-trivial continuum content must be carried by interface structure on richer substrates*. The eigenvalue-1 subspace remaining one-dimensional in every dimension — containing only the counting invariant — sharpens the case for the Interface Persistence Conjectures (§9) rather than weakening it. Adding bulk substrate dimensionality does not generate new refinement-stable bulk observables; the interface ontology becomes more pressing, not less.

The 16^{-n} commutator-decay rate cited elsewhere in the VERSF programme is *not* a single-step eigenvalue of $L_R^{(d)}$ on the linear fluctuation space for any d, and must therefore arise either from composite-step counting or from a different operator (e.g., the action of L_R on bilinear or

higher-tensor forms). Identifying the correct route is left open as the appropriate next question for the wider programme.

We also test the conjecture that BCB is conserved as net flux through codimension-1 antichains: a worked computation on the diamond refutes the strict form (the candidate flux observable decays at rate ½ per refinement step, placing it in the same eigenvalue-½ subspace as the spacelike and causal-gradient modes), and identifies the search for a corrected flux observable as a sharp open problem.

We organise the generalisability of these results into an explicit **three-level ladder**: Level 1 (the d-dimensional causal diamond, with spectrum computed for all $d \geq 1$), Level 2 (finite diamond-glued DAGs, treated patchwise), Level 3 (arbitrary admissible finite posets, requiring confluence and interface-localisation proofs). We prove a **Local-to-Global Refinement Extension Proposition** lifting the Level 1 result to Level 2 under stated overlap conditions.

We prove **Causal Stability at Levels A and B** (order stability, partition stability) and are explicit that Level C (full Lorentzian metric stability up to conformal scaling) is *not* established and would require additional structure analogous to Poisson sprinkling in causal-set theory.

We frame continuum emergence as a *universality-class* statement, decompose the Interface Persistence Conjecture into three testable sub-conjectures (antichain support, minimum-triangle selection, bulk decay), and provide an explicit failure-modes table pairing each claim with its precondition and consequence of failure.

Table of Contents

---

1. Introduction and Relation to Prior Work

The central problem of any discrete substrate theory is not discreteness. It is *persistence*. If physical reality is fundamentally discrete, why do large-scale observables appear continuous, stable, and largely independent of microscopic substrate detail?

VERSF approaches this through admissible refinement dynamics: physical observables are identified not with individual substrate states but with structures invariant under a class of admissibility-preserving refinement maps. The framework therefore replaces the conventional ontology

microscopic state → macroscopic world

with

admissible refinement equivalence classes → continuum observables.

This idea is not unique to VERSF, and we want to be explicit about its relation to neighbouring programmes:

- **Causal sets** (Bombelli, Lee, Meyer, Sorkin) [1] — substrate is a partial order plus a measure; continuum Lorentzian geometry is recovered (in part) via Poisson sprinkling, which is essential for avoiding preferred-frame artefacts [2].
- **Causal dynamical triangulations** (Ambjørn, Loll, Jurkiewicz) [3] — substrate is a simplicial complex with a causal foliation; continuum emergence is studied via Monte Carlo simulation and the spectral dimension flows from ≈ 2 at short distances to 4 at large [4].
- **Wilsonian renormalisation** [5, 6] — the canonical model for irrelevant-mode decay under coarse-graining; provides the language of fixed points, relevant/marginal/irrelevant modes, and universality classes [7].

VERSF's specific addition is the **admissibility closure** $A[\rho] = 0$ as a refinement-preserving constraint, together with the BCB and TPB structural primitives that fix the substrate's local content algebra. The claim of this paper is *not* that VERSF supplants any of the above. The claim is that within the admissibility-closed sector, refinement induces a well-defined flow on observables whose fixed structures are the candidates for continuum physical quantities.

---

2. Substrate Primitives and Admissibility Closure

The substrate consists of commitment events

$E = \{e_i\}$

organised into a partially ordered set

$\Lambda = (E, \preceq)$

equipped with a vector-valued commitment density

$\rho : E \to \mathbb{C}^{K+1}$

where K = 7 is the structural rank established by six independent derivation routes within the wider VERSF programme [10].

For this paper to stand alone we need a working form of the admissibility functional A[ρ]. We adopt the following schematic — sufficient for the constructions below and consistent with the fuller treatment in the substrate-dynamics programme [11]:

A[ρ] = 0 ⇔ both of the following hold:

- **(BCB)** Bit Conservation and Balance — content conservation across the substrate, formalised in §4 below as a monotone flow to a fixed-point coupling.
- **(TPB)** Ticks-Per-Bit causal consistency: for every covering pair $e \lessdot e'$ in Λ, $\langle \rho(e') - \rho(e), \eta(e, e') \rangle \in \mathbb{Z}$ where $\eta(e, e')$ is the substrate's local resolution vector at the covering pair.

The set of admissible substrate states is

$\mathcal{A}(\Lambda) = \{\rho : E \to \mathbb{C}^{K+1} \mid A[\rho] = 0\}.$

This is a finite-codimension submanifold of the full configuration space, and the refinement maps defined below act on $\mathcal{A}(\Lambda)$ rather than on arbitrary ρ.

---

3. Admissible Coarse-Graining: Explicit Construction

We construct Γ explicitly on a finite toy substrate.

3.1 Direction Convention

We fix the convention that Γ is a **refinement** map — it takes a coarser substrate Λ_n to a finer Λ_{n+1}. The composition $L_R = P \circ \Gamma$ with the coarse projection P (defined in §6) is the operator on which the renormalisation-style flow acts. "Admissible refinement" refers to Γ; "RG flow" refers to L_R.

3.2 Toy Substrate: The Iterated Causal Diamond

Let Λ_0 be the four-event causal diamond:

$E_0 = \{a, b, c, d\}$, with $a \preceq b \preceq d$, $a \preceq c \preceq d$, and b, c incomparable.

This is the minimal non-trivial substrate carrying both timelike (a, d) and spacelike (b, c) separated pairs.

The refinement Γ acts by **causal midpoint insertion**: for every covering pair (e, e′) in Λ_n, insert a new event m(e, e′) with

$$e \preceq m(e, e') \preceq e'$$

and otherwise comparable to the rest of the order exactly as the midpoint of e and e′ would be. The commitment density is extended by

$$\rho(m(e, e')) = \tfrac{1}{2} (\rho(e) + \rho(e')).$$

3.3 Order Preservation

Order preservation is immediate by construction: new events are inserted *between* existing ordered pairs and the existing order on E_n is preserved as a sub-order of Λ_{n+1}. Spacelike pairs remain spacelike.

3.4 TPB Preservation

If $\langle \rho(e') - \rho(e), \eta(e, e') \rangle \in \mathbb{Z}$ for the original covering pair, then for the two new covering pairs (e, m) and (m, e′) the inner products land in $\mathbb{Z}$ provided the new resolution vectors are chosen as $\eta(e, m) = \eta(m, e') = 2 \cdot \eta(e, e')$. The doubling is the natural consequence of η being measured in *ticks-per-bit*: a finer substrate requires more ticks to span the same physical separation, so the resolution-vector magnitude grows under refinement even though the underlying tick spacing shrinks. This is consistent with the TPB primitive of the wider VERSF programme.

3.5 Distinguishability Grounding

To avoid circularity in any operational definition of distinguishability, we adopt a substrate-level resolution functional defined *prior to* the refinement structure but with explicit refinement-covariant scaling:

$$D_{\Lambda_n}(e, e') = \|\rho(e) - \rho(e')\| + 2^{-n} \cdot \delta(e, e')$$

where δ(e, e′) is the order-theoretic distance (length of the shortest chain joining e and e′, with spacelike pairs handled via the size of the common antichain), and the 2^{-n} factor rescales the discrete chain length by the refinement level. D is refinement-stable in the limit $n \to \infty$ for any pair of events that exists at all sufficiently fine refinement levels; the rescaling is chosen so that δ-doubling under midpoint refinement is exactly compensated by the 2^{-n} factor for such pairs. For pairs involving newly-inserted midpoints — events that exist at level n but not at level n−1 — the statement is restricted to the levels at which both events exist. D is therefore defined directly on substrate states without reference to refinement-stable observables, and is itself refinement-stable up to the convergent rescaling on its domain of definition.

BCB preservation requires its own treatment, given that strict total content is *not* refinement-invariant. We address this in the next section.

---

4. BCB Renormalisation: Density Flow and Fixed-Point Conservation

The naive form of BCB — strict conservation of total substrate content $Q_N = \Sigma_{e \in E_N} \|\rho(e)\|^2$ — is not preserved under refinement. After one midpoint subdivision the number of events grows, and the sum of squared norms changes in a controlled but non-trivial way. We give the corrected formulation.

4.1 The Density Functional and Its Flow

Definition (Substrate Content Density). At refinement level n, the substrate content density is

$$\tilde{Q}_n = (1/|E_n|) \cdot \Sigma_{e \in E_n} \|\rho(e)\|^2.$$

Proposition 1 (BCB as Monotone Flow to a Fixed Point). Under midpoint refinement $\Gamma : \Lambda_n \to \Lambda_{n+1}$ on a substrate whose covering-pair graph is d-regular, the content density satisfies

$$\tilde{Q}_{n+1} \leq \tilde{Q}_n,$$

with equality iff ρ is constant across every covering pair of Λ_n. Iterated refinement therefore produces a monotone non-increasing sequence $\tilde{Q}_0 \geq \tilde{Q}_1 \geq \tilde{Q}_2 \geq \ldots \geq 0$, which converges to a refinement-invariant limit

$$\tilde{Q}_\infty = \lim_{n \to \infty} \tilde{Q}_n.$$

Proof. By convexity of the squared norm, for every covering pair (e, e′),

$$\|½(\rho(e) + \rho(e'))\|^2 \leq ½(\|\rho(e)\|^2 + \|\rho(e')\|^2),$$

with equality iff $\rho(e) = \rho(e')$. Summing over all covering pairs C_n of Λ_n:

$$Q_{n+1} = Q_n + \Sigma_{(e, e') \in C_n} \|½(\rho(e) + \rho(e'))\|^2 \leq Q_n + ½\, \Sigma_{(e, e') \in C_n} (\|\rho(e)\|^2 + \|\rho(e')\|^2) = \Sigma_{e \in E_n} (1 + \deg(e)/2) \|\rho(e)\|^2$$

where deg(e) is the covering-pair degree of e in Λ_n (the number of covering pairs in which e participates).

If the covering-pair graph is d-regular — as it is for the d-dimensional causal diamond, whose covering-pair graph is the d-cube graph Q_d — then $\deg(e) = d$ uniformly, giving

$$Q_{n+1} \leq (1 + d/2)\, Q_n.$$

Under midpoint insertion on a d-regular covering-pair graph, the number of events grows by the same factor: $|E_{n+1}| = |E_n| + |C_n| = |E_n|(1 + d/2)$, since $|C_n| = (d/2)|E_n|$. Dividing through yields

$$\tilde{Q}_{n+1} = Q_{n+1} / |E_{n+1}| \leq (1 + d/2)\, Q_n / [(1 + d/2)\, |E_n|] = \tilde{Q}_n.$$

The two $(1 + d/2)$ factors cancel exactly. The sequence is non-negative and non-increasing, hence convergent.

Equality holds throughout iff the convexity bound is saturated at every covering pair, which requires $\rho(e) = \rho(e')$ for every covering pair — i.e., ρ constant on covering pairs.

Non-regular substrates. On substrates whose covering-pair graph is not regular, the cancellation between Q_{n+1} bound and $|E_{n+1}|$ growth no longer holds uniformly. The right object then is the degree-weighted density $\tilde{Q}_n^{(w)} = (1/|E_n|)\, \Sigma_e\, w(e)\, \|\rho(e)\|^2$ for appropriately chosen w (typically $w(e) = 1/(1 + \deg(e)/2)$), or equivalently the lazy random walk with degree-corrected normalisation $W = \tfrac{1}{2}(I + D^{-1}A)$. We do not pursue the non-regular case here but flag that the monotonicity result extends to it with the appropriate degree-weighted observable.

4.2 The Renormalised Conservation Law

BCB is therefore *not* a step-by-step invariant. The refinement-invariant content is the limit $\tilde{Q}_\infty$ — the substrate-level analogue of a *renormalised coupling* in Wilsonian RG, where the bare coupling flows to a renormalised value at the fixed point of the RG transformation.

Corollary 1 (Renormalised BCB Conservation). Substrate content is not conserved at any finite refinement step; rather, the content density $\tilde{Q}$ flows monotonically to a refinement-invariant limit $\tilde{Q}_\infty$, which plays the role of the renormalised BCB coupling.

This is a weaker but accurate statement than the naive total-content reading. Strict total content is not conserved; density is not conserved at each step; only the asymptotic density $\tilde{Q}_\infty$ is refinement-invariant.

4.3 An Alternative Interpretation: BCB as Flux Through Interfaces

A natural conjecture, motivated by §9, is that the *correct* conserved quantity is not bulk content density at all but a flux through codimension-1 interfaces — connecting BCB conservation directly to the Interface Persistence Conjecture and giving BCB an interface-localised substrate origin.

To make this concrete on the diamond: the maximal antichain {b, c} separates the past event a from the future event d. We define the candidate **flux observable** for antichain Σ as a signed bilinear form

$$\Phi(\Sigma) = \Sigma_{(e, e') \text{ covering}, e' \in \Sigma} \langle \rho(e), \rho(e') \rangle - \Sigma_{(e, e') \text{ covering}, e \in \Sigma} \langle \rho(e), \rho(e') \rangle$$

— covering pairs whose upper endpoint lies on the antichain contribute positively (incoming flux), and covering pairs whose lower endpoint lies on the antichain contribute negatively (outgoing flux). This is the net crossing rate through Σ.

On the diamond with $\Sigma = \{b, c\}$, the four covering pairs split as: (a, b) and (a, c) have their upper endpoint in Σ (incoming), while (b, d) and (c, d) have their lower endpoint in Σ (outgoing). Hence:

$$\Phi(\{b, c\}) = \langle \rho(a), \rho(b) \rangle + \langle \rho(a), \rho(c) \rangle - \langle \rho(b), \rho(d) \rangle - \langle \rho(c), \rho(d) \rangle.$$

Φ is a bilinear functional of ρ, not a quadratic form in $\delta\rho$. Its behaviour under refinement is therefore governed not by L_R or $L_R \otimes L_R$ on quadratic forms, but by the induced action of L_R on bilinear forms $\langle \rho(e), \rho(e') \rangle$ supported on covering pairs.

4.3.1 Worked Computation on the Diamond

We evaluate Φ on representative configurations to test the conjecture.

Constant mode $\rho(a) = \rho(b) = \rho(c) = \rho(d) = r$:

$$\Phi(\{b, c\}) = \langle r, r \rangle + \langle r, r \rangle - \langle r, r \rangle - \langle r, r \rangle = 0.$$

The counting invariant carries zero net interface flux, as expected for a uniform configuration.

Corner-alternating mode $\rho = (1, -1, -1, 1)$:

$$\Phi(\{b, c\}) = (1)(-1) + (1)(-1) - (-1)(1) - (-1)(1) = -1 - 1 + 1 + 1 = 0.$$

The maximally oscillatory mode also carries zero net flux — the symmetric ± pattern produces exact cancellation between incoming and outgoing contributions. The corner-alternating mode therefore demonstrates that strong local oscillation does not by itself imply nonzero net interface flux under the present definition.

Asymmetric configuration $\rho = (1, 1, 0, 0)$:

$$\Phi_0(\{b, c\}) = \langle 1, 1 \rangle + \langle 1, 0 \rangle - \langle 1, 0 \rangle - \langle 0, 0 \rangle = 1.$$

Under one refinement step Γ, the midpoint values are $\rho(m_{ab}) = 1$, $\rho(m_{ac}) = ½$, $\rho(m_{bd}) = ½$, $\rho(m_{cd}) = 0$. The covering pairs crossing {b, c} in Λ_1 are (m_{ab}, b), (m_{ac}, c) incoming and (b, m_{bd}), (c, m_{cd}) outgoing. Therefore

$\Phi_1(\{b, c\}) = \langle 1, 1\rangle + \langle ½, 0\rangle - \langle 1, ½\rangle - \langle 0, 0\rangle = 1 - ½ = ½$.

So $\Phi_1 = (½)\ \Phi_0$ on this configuration. The same ratio appears on other non-symmetric configurations: e.g. $\rho = (1, 1, 0, 2)$ gives midpoints $m_{ab} = 1$, $m_{ac} = ½$, $m_{bd} = 3/2$, $m_{cd} = 1$, hence $\Phi_0 = 1 - 2 = -1$ and $\Phi_1 = 1 - 3/2 = -½$, again a factor ½ below.

4.3.2 Conclusion: The Conjecture as Stated Is Refuted

The candidate flux observable Φ as defined is *not* refinement-stable on the diamond; it decays at rate ½ per step under the bilinear refinement action, placing it in the same eigenvalue-½ subspace as the spacelike and causal-gradient modes of L_R on the linear fluctuation space (Theorem 1). The Flux-BCB Equivalence Conjecture as stated is therefore refuted in its strict form.

This is a useful negative result. It demonstrates that:

(i) Refinement stability and interface flux are *structurally distinct* notions — Φ is a non-trivial bilinear form on covering pairs but is not a refinement invariant under the natural bilinear action of L_R.

(ii) The flux-through-interfaces picture for BCB conservation, if it survives at all, requires a *different* flux observable — perhaps incorporating refinement-covariant scaling of the kind appearing in §3.5, or selecting a different antichain, or using a different inner-product convention.

(iii) Identifying the correct flux observable is the natural next step. The negative result here at least narrows the search: any candidate must satisfy $\Phi_{n+1} = \Phi_n$ rather than $\Phi_{n+1} = (½)\ \Phi_n$ under the bilinear action of L_R.

We do not pursue this further here but note that the question is now sharp: a corrected flux observable, if it exists, will lift §9's Interface Persistence Conjectures 1a–1c from "interface localisation of linear refinement-stable observables" to a stronger "interface localisation of bilinear refinement-stable observables that match BCB." The two routes to BCB-as-interface-flux are therefore both open: (a) finding the correct flux observable, and (b) showing that bulk BCB density flow $\tilde{Q}_\infty$ can be re-expressed as an interface integral in the continuum limit.

---

5. The Generalisability Ladder

The diamond is not claimed to represent the general substrate. It is the base case of an inductive refinement programme, and we organise the generalisation explicitly into three levels.

5.1 The Three Levels

- **Level 1 — Symmetric causal diamond.** The construction of §3 and the spectrum of §6 are *fully explicit* at this level. All claims here are provable by direct computation.
- **Level 2 — Finite diamond-glued DAGs.** Substrates obtained by gluing finitely many causal diamonds along compatible faces, with overlaps preserving order and TPB increments. Midpoint refinement applies patchwise. We prove the local-to-global extension proposition below at this level.
- **Level 3 — Arbitrary admissible finite posets.** The most general setting. Extension from Level 2 requires (a) a confluence theorem ensuring that patchwise refinement is well-defined independent of patch choice, and (b) an interface-localisation result identifying where refinement-stable structure lives in non-diamond-glued substrates. Both are open problems within the programme.

5.2 The Local-to-Global Extension

Proposition 2 (Local-to-Global Refinement Extension). Let Λ be a finite admissible poset admitting a cover $\{U_\alpha\}$ by diamond-like subposets such that for every overlap $U_\alpha \cap U_\beta$:

(i) the induced order on the overlap agrees on both sides, (ii) the $\tilde{Q}$-density agrees on the overlap up to a refinement-invariant tolerance, (iii) the TPB resolution vectors $\eta(e, e')$ agree on the overlap.

Then the midpoint-refinement construction extends patchwise from each U_α to a global refinement $\Gamma : \Lambda \to \Lambda'$. The induced spectrum of L_R on Λ' contains local spectral sectors continuously connected to the isolated d-dimensional causal diamond spectrum $\{1, (d-1)/d, \ldots, 0\}$ (with binomial multiplicities; Theorem 1) in the weak-overlap limit, together with additional eigenvalues from gluing modes that we conjecture lie in $[0, 1]$. Throughout this paper, "diamond-glued DAG" allows arbitrary-dimensional causal diamonds as patches; the $d = 2$ case is the simplest but the construction generalises.

Spectral mixing caveat. When overlap regions between glued diamonds are non-trivial, local eigenmodes generically hybridise. The persistence of isolated spectral sectors therefore depends on the overlap operator norm. In the weak-overlap regime

$\|A_{overlap}\| \ll \|A_{local}\|$

standard perturbation arguments imply that the local spectral structure survives continuously under gluing. Outside this regime — when overlaps are strong relative to local structure — eigenmodes mix and the local-spectrum decomposition is no longer literal but only approximate. A full classification of gluing-mode spectra under arbitrary overlap strength remains part of the Level 3 programme.

Proof sketch. Conditions (i)–(iii) ensure that the local midpoint construction is consistent across overlaps: any midpoint inserted by patch U_α is identical to the one that would be inserted by patch U_β on the same covering pair. The induced global Γ is therefore well-defined. The spectrum of L_R on Λ' decomposes into local d-dimensional causal diamond contributions (Theorem 1: $\{1, (d-1)/d, \ldots, 0\}$ with binomial multiplicities, where d may vary patchwise) plus gluing-mode contributions arising from fluctuations supported on overlap regions. A full classification of gluing-mode eigenvalues requires the confluence theorem of Level 3.

5.3 What This Buys

The ladder moves the paper from "one toy example" to "base case plus extension programme." Level 1 is *proved*. Level 2 is *proved conditional on overlap conditions (i)–(iii)*. Level 3 is *flagged as open*, with the two open problems (confluence, interface localisation) stated precisely.

This structure is essential for honest scope-setting: we do not claim the diamond is representative of physical substrates. We claim it is the base case of a controlled programme, with the first inductive step proved.

---

6. The Linearised RG Operator L_R = P ∘ Γ and Its Spectrum

We now define L_R explicitly and compute its spectrum on the diamond.

6.1 The Coarse Projection P_n

Definition (Coarse Projection P_n). Let $\Gamma_n : \Lambda_n \to \Lambda_{n+1}$ be admissible refinement. The coarse projection

$P_n : \Lambda_{n+1} \to \Lambda_n$

acts on fluctuations $\delta\rho : E_{n+1} \to \mathbb{C}^{K+1}$ by averaging midpoint values back onto their parents:

$(P_n\, \delta\rho)(e) = \frac{1}{2} \Sigma_{\{m : e \text{ is a parent of } m\}}\, \delta\rho(m)$

where the sum runs over all midpoints adjacent to e in Λ_{n+1}, and the ½ weight ensures that each midpoint's content is shared equally between its two parents.

The original event value $\delta\rho(e)$ is *not* preserved by P; it is recovered from the midpoints adjacent to e, each of which carries half of e's value plus half of its other parent's value. This is the natural choice that makes L_R = P ∘ Γ stochastic on the constant mode (eigenvalue 1) and is equivalent to the *lazy random walk* on the covering-pair graph of Λ_n — a principled choice from the theory of Markov operators on graphs.

The L^2-adjoint Γ^T is a natural alternative but produces a non-stochastic operator (constant-mode eigenvalue 2 before normalisation), a different and less clean spectrum, and no exact-extinction direction. Within the one-parameter family of stochastic operators $M_\alpha = \alpha I + (1-\alpha)\, D^{-1}A$ on a d-regular covering-pair graph, the corner-alternating eigenvalue is $\alpha - (1-\alpha) = 2\alpha - 1$, which vanishes iff $\alpha = \frac{1}{2}$. The random-walk choice is therefore forced within this family by the two structural features the spectrum depends on — stochasticity on the constant mode and exact extinction on the corner-alternating mode.

A stronger uniqueness claim — that the random-walk P is unique among *all* stochastic projections commuting with the d-cube automorphism group and sending the corner-alternating mode to zero — would require additional constraints. The relevant algebraic object is the Bose–Mesner algebra of the Hamming scheme H(d, 2) — the (d+1)-dimensional commutative algebra generated by the d-cube graph's distance matrices — so two constraints alone (stochasticity, corner-alternating extinction) leave d−1 degrees of freedom. We do not pursue that here, but flag that the construction is naturally embedded in this larger algebraic structure.

Regularity caveat. The lazy-random-walk identification $W = ½(I + D^{-1}A)$ reduces to $(1/(2d))(dI + A)$ only for d-regular covering-pair graphs. The d-dimensional causal diamond (§6.4) is d-regular, so the identification holds. Level 2 substrates (glued DAGs) need not be regular, and the correct generalisation there is $W = ½(I + D^{-1}A)$ directly, with degree-weighted normalisation per vertex.

6.2 The Linearised RG Operator

Definition (Linearised RG Operator). The linearised RG operator is the composition

$L\_R = P\_n \circ \Gamma\_n$: (fluctuations on Λ_n) → (fluctuations on Λ_n).

Iterated refinement-then-projection produces a flow on a fixed finite-dimensional fluctuation space, on which spectral analysis is well-defined.

The eigenvalues quoted in Appendix A are the spectrum of $L\_R = P \circ \Gamma$ as a fixed-dimensional operator, not of Γ alone (which changes dimension at each step).

6.3 Spectrum on the Diamond

Working with a single component of ρ for clarity, the fluctuation space on Λ_0 is $\mathbb{C}^4$ in the basis (δa, δb, δc, δd). The covering-pair graph of the diamond is the 4-cycle a–b–d–c–a, with adjacency matrix A. One full step of $L\_R = P \circ \Gamma$ takes the form $L\_R = (1/4)(2I + A)$, whose spectrum is (1/4){4, 2, 2, 0} = {1, ½, ½, 0}. Eigenvectors:

- **v_1 = (1, 1, 1, 1)**, eigenvalue **1** — constant mode (counting invariant), refinement-stable.
- **v_2 = (1, 0, 0, −1)**, eigenvalue **½** — causal-gradient mode along a → d, refinement-irrelevant.
- **v_3 = (0, 1, −1, 0)**, eigenvalue **½** — spacelike "checkerboard" mode, refinement-irrelevant.
- **v_4 = (1, −1, −1, 1)**, eigenvalue **0** — corner-alternating mode, killed exactly in one step.

The eigenvalue-1 subspace is one-dimensional. The matrix form and verification of these eigenvalues are given in Appendix A.

The corner-alternating mode being killed exactly — not merely decaying geometrically — is a stronger structural result than geometric decay. The covering-pair graph on the diamond is bipartite with the two halves coupling antisymmetrically in the corner-alternating direction, which produces exact cancellation under the parent-projection P. Whether the same exact extinction obtains for the corner-alternating mode on higher-dimensional substrates is the question taken up next.

6.4 The d-Dimensional Spectrum from the d-Cube Graph

Why the d-cube appears naturally. The emergence of the d-cube graph as the covering-pair graph is not imposed externally; it follows directly from the bitwise causal structure of the d-dimensional causal diamond. Each event is labelled by a binary string of length d, with $x \preceq y$ iff x's bits are pointwise $\leq$ y's bits. Covering relations correspond to changing exactly one bit from 0 to 1. The undirected covering-pair graph therefore connects vertices differing in precisely one coordinate, which is exactly the defining adjacency relation of the Hamming cube Q_d. The spectrum of $L_R^{(d)}$ is therefore inherited from the algebraic structure of the Hamming scheme rather than inserted ad hoc.

The natural generalisation of L_R to the d-dimensional causal diamond is the lazy random walk on Q_d:

$$L_R^{(d)} = (1/(2d)) \cdot (d \cdot I + A_{Q_d})$$

where A_{Q_d} is the d-cube graph adjacency. The d-cube graph has well-known spectrum $\{d, d-2, d-4, \ldots, -d\}$ with binomial multiplicities $\{C(d,0), C(d,1), \ldots, C(d,d)\}$. Therefore $L_R^{(d)}$ has spectrum

$$(1/(2d)) \cdot \{2d, 2d-2, 2d-4, \ldots, 0\} = \{1, (d-1)/d, (d-2)/d, \ldots, 0\}$$

with the same binomial multiplicities. The corner-alternating mode (the all-(±1) eigenvector) is the unique direction with eigenvalue $-d$ in A_{Q_d} and therefore eigenvalue **0** in $L_R^{(d)}$ in every dimension $d \geq 1$ — *exact extinction in one step, independent of d*. The constant mode is the unique direction with eigenvalue $+d$ in A_{Q_d} and eigenvalue 1 in $L_R^{(d)}$, also independent of d.

6.4.1 The d = 4 Spectrum Explicitly

For $d = 4$, A_{Q_4} has spectrum $\{4, 2, 0, -2, -4\}$ with multiplicities $\{1, 4, 6, 4, 1\}$ (the binomial coefficients $C(4, k)$). Therefore

$$\mathrm{spec}(L_R^{(4)}) = \{1, 3/4, 1/2, 1/4, 0\}$$

with the same multiplicities:

Eigenvalue	Multiplicity	Interpretation
1	1	Constant mode (counting invariant)

Eigenvalue	Multiplicity	Interpretation
3/4	4	Single-axis gradients (slowest non-stable decay)
1/2	6	Two-axis modes
1/4	4	Three-axis modes
0	1	Corner-alternating mode (exact extinction in one step)

The eigenvalue-1 subspace remains one-dimensional in $d = 4$, as in $d = 2$. The corner-alternating mode is still extinguished in one step. The intermediate eigenvalues 3/4, 1/2, 1/4 give multiple decay rates with binomial multiplicities — a richer subspace structure than the $d = 2$ case, in which the slowest non-stable mode (rate 3/4 per step, multiplicity 4) corresponds to causal-gradient-like directions in the four substrate axes.

6.4.2 What This Says About the 16^{-n} Rate

The 16^{-n} rate cited elsewhere in the VERSF programme [12] for non-Abelian commutator-defect decay is **not** an eigenvalue of $L\_R^{(4)}$ under this construction. The eigenvalues of $L\_R^{(4)}$ are {1, 3/4, 1/2, 1/4, 0}, and 1/16 is not among them. There are two structurally distinct ways the rate could nevertheless be consistent with this construction:

(a) **Cumulative decay across composite steps.** If the wider programme's "16^{-n} rate" refers to decay across n *coarse-grained* steps where each coarse step is itself a fixed number of refinements, then 16^{-n} could correspond to a single eigenvalue raised to 4n or similar. For instance, four refinements of the eigenvalue-½ subspace give $(½)^{4k} = 16^{-k}$. The match here depends on the precise convention for "one decay step" in [12] — whether it is one refinement of $L\_R^{(4)}$ or a composite step — and a careful comparison against [12]'s construction is needed to settle this.

(b) **Different operator on the substrate.** The 16^{-n} rate could come from a different operator entirely — perhaps $L\_R^{(4)}$ acting on bilinear or higher-tensor forms (where products of eigenvalues become available: $1/16 = (1/4)^2$ or $(1/2)\cdot(1/8)$ etc.), or a still-different construction in the wider programme.

We do not commit to either possibility here. What is established by this section is that the 16^{-n} rate is **not** the single-step eigenvalue of $L\_R^{(d)}$ on the linear fluctuation space for any d, and its origin within VERSF must come from a different construction or a different counting of refinement steps. Identifying that construction is left open and is the appropriate next question for the wider programme to address.

6.5 Theorem 1 (Refinement-Stability and Decay on the d-Dimensional Causal Diamond)

Theorem 1. Let $\Lambda^{(d)}$ be the d-dimensional causal diamond (2^d events labelled by binary strings of length d, with $x \preceq y$ iff x is bitwise $\leq$ y), and let $L\_R^{(d)} = (1/(2d))(d \cdot I + A\_\{Q\_d\})$ be the linearised RG operator on the 2^d-dimensional complex fluctuation space. Then:

(i) **Spectrum.** $L\_R^{(d)}$ has spectrum {1, (d−1)/d, (d−2)/d, …, 0} with binomial multiplicities {C(d,0), C(d,1), …, C(d,d)}.

(ii) **One-dimensional stable subspace.** The eigenvalue-1 subspace is one-dimensional in every d ≥ 1, spanned by the constant mode on the 2^d-event fluctuation space.

(iii) **Corner-alternating extinction.** The corner-alternating mode (the all-(±1) eigenvector of A_{Q_d}, parameterised by parity of binary-string Hamming weight) is in the kernel of $L\_R^{(d)}$ in every d ≥ 1 — extinguished exactly in one step.

(iv) **Decay structure.** All non-constant modes have eigenvalue ≤ (d−1)/d, and iterated refinement drives the fluctuation spectrum onto the one-dimensional constant subspace, with the corner-alternating direction extinguished in a single step and intermediate directions decaying geometrically at rates determined by the binomial structure of the d-cube spectrum.

Proof. The covering-pair graph of $\Lambda^{(d)}$ is the d-cube graph Q_d, since covering pairs in the bitwise partial order correspond to single-bit flips. The d-cube graph has spectrum {d, d−2, …, −d} with binomial multiplicities (standard result on distance-regular graphs and the Hamming scheme; see [9]). Therefore

$$\mathrm{spec}(L\_R^{(d)}) = (1/(2d))(d + \mathrm{spec}(A\_\{Q\_d\})) = (1/(2d))\{2d, 2d-2, \ldots, 0\} = \{1, (d-1)/d, \ldots, 0\}$$

with the same binomial multiplicities. Since Q_d is d-regular, the all-1 vector is the unique eigenvector of A_{Q_d} with eigenvalue +d, giving constant-mode eigenvalue 1 in $L\_R^{(d)}$. Since Q_d is bipartite (with bipartition by Hamming-weight parity), the smallest adjacency eigenvalue is −d, attained uniquely up to scaling by the bipartition indicator function — the all-(±1) parity vector. This gives corner-alternating eigenvalue 0 in $L\_R^{(d)}$. (iv) follows from (i) directly.

6.6 Interpretation of the One-Dimensional Stable Subspace

The eigenvalue-1 subspace being one-dimensional in *every* dimension d ≥ 1 — containing only the counting invariant — is a strong structural result: adding bulk substrate dimensionality does not generate new refinement-stable bulk observables. Non-trivial continuum content must therefore come either from richer substrates (Levels 2 and 3) *or* from substrate structure not captured by bulk modes — i.e., from interface localisation on antichains, which is precisely the empirical content of the Interface Persistence Conjectures (§9).

Dependency note. The interface-localisation reading depends on the conjecture that gluing modes on Level 2 substrates have eigenvalue strictly less than 1 under L_R. A Level 2 substrate exhibiting eigenvalue-1 gluing modes would generate new refinement-stable bulk observables and weaken the interface ontology accordingly. This contingency is added to the failure-modes table in §10.2.

Subject to that dependency, the diamond and d-cube computations make the Interface Persistence Conjectures more pressing, not less. If bulk modes on every d-dimensional causal diamond reduce to counting invariants under refinement, *and* if gluing modes on Level 2 also have eigenvalue < 1, then all non-trivial continuum content must live on antichains — and the Interface Persistence Conjectures 1a–1c follow as the natural classification of which antichains carry it.

---

7. Causal Stability — Three Levels

We organise the causal-stability claim into three levels and are explicit about which are proved.

7.1 The Three Levels

Level A — Order Stability. For all pairs $(e, e') \in E_n \times E_n$, comparability vs incomparability under $\preceq$*n agrees with that under* $\preceq${n+1}.

Level B — Causal-Partition Stability. The partition $\pi_n : E_n \times E_n \to \{\text{timelike, spacelike}\}$ induced by $\preceq_n$ (where "timelike" means comparable and "spacelike" means incomparable) is preserved under Γ.

Level C — Lorentzian Metric Stability. The emergent interval $s^2(e, e')$, defined as the refinement-stable continuum limit of a substrate-level quadratic form, is invariant under admissible refinement up to a conformal rescaling.

Level B is deliberately the *coarse* two-way partition rather than the three-way {timelike, spacelike, null} partition. The reason: under midpoint refinement, what was a null (covering) pair in Λ_n is no longer covering in Λ_{n+1} — a midpoint has been inserted between its endpoints, making it a two-step timelike pair at the next level. The covering relation is not refinement-stable, so the null sub-classification (if defined via covering) cannot be either.

Recovering a refinement-stable notion of null separation requires a separate definition (light-cone boundary in the continuum limit, or minimum-chain-length structure under refinement-covariant rescaling) and is not pursued here.

7.2 What This Paper Proves

Proposition 3 (Causal Stability at Levels A and B). On the iterated d-dimensional causal diamond, midpoint refinement preserves Level A (order stability) and Level B (the {timelike, spacelike} causal partition).

Proof. Refinement inserts new events between covering pairs but does not alter comparability of existing pairs. Timelike pairs in $E_n \times E_n$ remain comparable in Λ_{n+1} (with strengthened chains via the new midpoints). Spacelike pairs remain incomparable, since no new event is inserted that would generate a chain between them. The {timelike, spacelike} classification on $E_n \times E_n$ is therefore unchanged under Γ.

7.3 What This Paper Does *Not* Prove

Open: Level C. The substrate-level quadratic form $s^2 = c_\rho^2 \tau^2 - d^2$ has Lorentzian signature *by construction*. Recovering it as the *unique* refinement-stable quadratic form on the emergent tangent space — rather than positing it — requires:

(i) a substrate-level measure $\mu : E \to \mathbb{R}$ that is refinement-covariant; (ii) a proof that the orbits of μ-preserving substrate automorphisms generate, in the continuum limit, an isometry group conjugate to $O(1, d-1)$ rather than $O(d)$ or some discrete subgroup; (iii) a uniqueness theorem identifying s^2 up to conformal rescaling.

Causal sets achieve (i) via Poisson sprinkling [2]. The substrate analogue of Poisson sprinkling within VERSF is not yet specified. Level C is therefore **open**, and we make no claim of having established it.

This three-level decomposition gives a referee a precise map of what is proved and what is not.

7.4 Scope of the Causal-Stability Claim

To be fully explicit: the present paper establishes refinement stability of *causal order* and the *causal partition* (Levels A and B) only. It does not establish emergent Lorentz symmetry in the strong field-theoretic sense used in relativistic quantum field theory — i.e., we do not prove that the emergent tangent structure admits a representation of the Lorentz group $O(1, d-1)$ acting on quantum fields, nor that physical observables transform covariantly under such a representation.

The claim is therefore weaker but precise: admissible refinement preserves causal structure at the order-theoretic level. Whether this extends uniquely to continuum Lorentz invariance in the field-theoretic sense remains open and would require the full programme outlined in §7.3.

---

8. Continuum Emergence as a Universality Class

We now consider what we regard as the conceptual centre of the paper.

8.1 The Equivalence Relation

Define an equivalence relation on admissible substrate histories:

$\Lambda \sim_R \Lambda' \Leftrightarrow$ *for every refinement-stable observable O,* $\lim_{n\to\infty} O(L_R^n \Lambda) = \lim_{n\to\infty} O(L_R^n \Lambda')$.

The equivalence class $[\Lambda]_R$ contains all substrate histories that share the same continuum-limit values of refinement-stable observables. This is the substrate-level analogue of a *universality class* in statistical mechanics.

8.2 The Universality-Class Statement

Universality-Class Principle. Continuum physics corresponds to an equivalence class $[\Lambda]_R$ rather than to any particular microscopic substrate realisation. Different admissible substrates within the same class produce identical continuum observables; substrates in different classes correspond to physically distinct continuum worlds.

This is not itself a theorem — it is the *interpretive principle* that gives the refinement-stability theorem its physical content. Theorem 1 proves that refinement-stable observables exist and are well-defined; the universality-class principle states what they *mean*.

8.3 Comparison to Statistical Mechanics

The parallel is direct. The Ising model and the φ^4 scalar field flow to the same fixed point under Wilsonian RG and share critical exponents — they lie in the same universality class [7, 8]. Microscopic details differ; macroscopic observables are identical. In VERSF, two admissible substrates with different microscopic ρ structure but identical refinement-stable observables lie in the same admissible refinement class. Microscopic substrate detail differs; emergent continuum physics is identical.

The disanalogy is also instructive: Wilsonian RG operates on a fixed background spacetime; VERSF refinement operates on the substrate from which spacetime itself emerges. The mathematical machinery is similar; the ontological role is different.

---

9. Interface Persistence — Decomposed Conjectures

The Interface Persistence Conjecture decomposes naturally into three sub-conjectures, each of which is independently testable and (we hope) eventually provable.

Conjecture 1a (Antichain Support). Refinement-stable nontrivial observables in $V_\infty \subset \mathcal{A}(\Lambda)$ are supported on maximal antichains of Λ.

Conjecture 1b (Minimum-Triangle Selection). Among maximal antichains, refinement-stable support localises on the antichains of *minimum triangle degree* — where the triangle degree of an antichain $\Sigma \subset \Lambda$ is the number of three-element chains ($e_1 \preceq e_2 \preceq e_3$) intersecting Σ at e_2. This is the substrate analogue of codimension-1 hyperfaces in a smooth manifold: minimum triangle degree corresponds to maximally "flat" antichains carrying minimal causal branching, and is the order-theoretic counterpart of the boundary-curvature-minimising surfaces of differential geometry.

Conjecture 1c (Bulk Decay). Bulk-supported fluctuations — those *not* localised on any maximal antichain — have eigenvalues $|\lambda| < 1$ under L_R.

9.1 Why Decomposition Matters

The decomposed form is *attackable*: each sub-conjecture has its own proof strategy.

- **1a** can be tested on Level 2 substrates (glued DAGs) by computing the spectrum of L_R and checking that eigenvalue-1 eigenvectors have support concentrated on antichain structure.
- **1b** is a selection principle within the antichain support — it adds the geometric content (minimum triangle degree) that connects to the interface-geometry derivation of α^{-1} in the wider VERSF programme.
- **1c** is a *negative* statement that is, in principle, the easiest of the three to prove: one needs to show that bulk modes always carry some oscillation that gets averaged out.

9.2 Why This Matters for the Wider Programme

If all three hold:

1. Bulk substrate degrees of freedom are refinement-irrelevant — the continuum world emerges from boundary-localised admissible structure.
2. Gauge transport, U(1) connections, and the record-current structure of the broader VERSF programme [13, 14] acquire a substrate-level origin as interface-localised refinement-stable observables.
3. The interface-geometry derivation of the fine-structure constant in the wider VERSF programme [15] gets a *refinement-theoretic* justification rather than only a geometric one.

None of 1a, 1b, 1c is established here. All three are flagged as **conjectural** in §10.

---

10. Epistemic Status and Failure Modes

In keeping with the VERSF programme's epistemic-labelling discipline:

10.1 Status Table

Proven (within Level 1, the d-dimensional causal diamond for all $d \geq 1$):

- Order preservation under midpoint refinement.
- TPB preservation under the resolution-doubling rule.
- BCB monotone flow to a fixed point (Proposition 1): $\tilde{Q}$ is non-increasing under refinement and converges to a refinement-invariant limit $\tilde{Q}_\infty$, on substrates with d-regular covering-pair graphs.
- Spectrum of $L\_R^{(d)}$ on the d-dimensional causal diamond: $\{1, (d-1)/d, \ldots, 0\}$ with binomial multiplicities (Theorem 1).

- The eigenvalue-1 subspace is one-dimensional in every d, spanned by the counting invariant.
- The corner-alternating mode is in the kernel of $L\_R^{(d)}$ for every $d \geq 1$ — exact extinction in one step (Theorem 1(iii)).
- The $d = 4$ spectrum $\{1, 3/4, 1/2, 1/4, 0\}$ with multiplicities $\{1, 4, 6, 4, 1\}$ (§6.4.1).
- The 16^{-n} rate is *not* a single-step eigenvalue of $L\_R^{(d)}$ on the linear fluctuation space for any d (§6.4.2).
- **Refutation (within Level 1):** The Flux-BCB Equivalence Conjecture, in its strict form (Φ refinement-stable as a bilinear form), is refuted on the diamond: $\Phi\_\{n+1\} = (½)\ \Phi\_n$ on non-symmetric configurations (§4.3.2).
- Causal Stability at Levels A and B (Proposition 3).

Conditional (true under stated assumptions, not yet proved in full generality):

- Local-to-Global Refinement Extension (Proposition 2) — proved conditional on overlap conditions (i)–(iii) and the weak-overlap regime for the spectral-sector connection.
- Distinguishability functional D refinement-invariance in the limit $n \to \infty$ — established under the 2^{-n} rescaling of δ; full proof of convergence rates omitted.
- BCB monotone flow on non-regular substrates — extends via degree-weighted observables but not pursued here.
- Extension of Theorem 1 from Level 1 to Level 3 — requires the confluence theorem at Level 3.

Conjectural:

- Identification of the 16^{-n} rate within the wider VERSF programme: which operator on which space produces it (§6.4.2).
- Whether a *corrected* flux observable (incorporating refinement-covariant scaling or a different convention) can recover BCB as interface flux (§4.3.2): the strict-form conjecture is refuted, but the underlying picture may still be salvageable in a different form.
- Causal Stability at Level C (full Lorentzian metric stability up to conformal scaling).
- Lorentz symmetry in the strong field-theoretic sense (§7.4).
- Interface Persistence Conjectures 1a, 1b, 1c.
- Uniqueness of the refinement-stable Lorentzian quadratic form.

10.2 Failure Modes Table

Claim	Required condition	Consequence if condition fails
BCB asymptotic conservation	$\tilde{Q}$ monotone, bounded below	$\tilde{Q}$ has no fixed point; renormalised BCB coupling undefined
Distinguishability scaling	2^{-n} rescaling of order-theoretic distance δ	D not refinement-stable; circularity returns

Claim	Required condition	Consequence if condition fails
Spectrum result on $\Lambda^{(d)}$	d-regular covering-pair graph (the d-cube Q_d) + bitwise-covering structure	Spectrum loses d-cube binomial form; falls back to direct computation per substrate
Eigenvalue-1 subspace ≠ trivial	At least the counting invariant survives	No refinement-stable observable; continuum emergence fails at the bulk level
Causal-partition stability	Order-preserving Γ	Causal partition can change under non-order-preserving coarse-graining
Local-to-global extension	Overlap conditions (i)–(iii) on cover	Patchwise refinement becomes ill-defined; Level 2 collapses to Level 1
Continuum richness on glued substrates (too small)	Eigenvalue-1 subspace stays 1D on Level 2	Universality class trivially small — but interface ontology is forced
Continuum richness on glued substrates (too large)	Gluing modes with eigenvalue < 1 (Prop. 2 conjecture)	If Level 2 gluing modes have eigenvalue = 1, new bulk-supported refinement-stable observables appear and interface ontology weakens
16^{-n} rate from this construction	Construction must apply to bilinear/tensor forms, or composite-step counting	If neither route works in [12]'s construction, 16^{-n} has no derivation from $L\_R^{(d)}$ on linear fluctuations
Interface persistence	Stable modes localise on antichains (Conj. 1a)	Interface ontology weakens; gauge structure loses substrate-level origin

This table makes the *contingency structure* of the paper transparent. Each claim has a precondition and a known consequence of that precondition failing.

---

11. What Is and Is Not Claimed

This paper claims:

- A refinement operator Γ can be constructed explicitly on a non-trivial substrate (Level 1), with order preservation, TPB preservation, and refinement-covariant distinguishability all checkable by direct computation.
- BCB is correctly read as a monotone flow to a fixed-point coupling $\tilde{Q}\_\infty$, not as a strict step-by-step invariant.
- The linearised RG operator $L\_R^{(d)}$ on the d-dimensional causal diamond has computed spectrum {1, (d−1)/d, …, 0} with binomial multiplicities, for every d ≥ 1 (Theorem 1). The eigenvalue-1 subspace is one-dimensional in every d, spanned by the counting invariant.
- The corner-alternating mode is in the kernel of $L\_R^{(d)}$ for every d ≥ 1 — extinguished exactly in one step, a structurally stronger result than geometric decay.

- The 16^{-n} commutator-decay rate cited elsewhere in the VERSF programme is *not* a single-step eigenvalue of $L\_R^{(d)}$ on the linear fluctuation space for any d, and must come from a different construction or different step-counting convention (§6.4.2).
- The Flux-BCB Equivalence Conjecture, in its strict form (Φ refinement-stable as a bilinear form on covering pairs), is refuted on the diamond: Φ decays at rate ½ per step under the bilinear refinement action (§4.3.2). Recovering BCB as interface flux requires a corrected flux observable; identifying it is left open.
- The Level 1 result extends to Level 2 (diamond-glued DAGs) under explicit overlap conditions (Proposition 2), with the local spectrum continuously connected to the d-dimensional diamond spectrum in the weak-overlap regime.
- Causal Stability at Levels A and B is preserved under admissible refinement.
- The one-dimensional eigenvalue-1 subspace in every dimension *motivates* the Interface Persistence Conjectures: non-trivial continuum content cannot be supported by bulk modes of any d-dimensional causal diamond, and must therefore live on interfaces of richer substrates.
- Continuum physics is naturally framed as a universality class [Λ]_R.

This paper does not claim:

- A derivation of General Relativity, quantum mechanics, or the Standard Model.
- A derivation of the 16^{-n} commutator-decay rate from the $L\_R^{(d)}$ spectrum on linear fluctuations: the spectrum is {1, (d−1)/d, …, 0} and 1/16 is not among these eigenvalues.
- A non-trivial eigenvalue-1 subspace in any dimension: only the counting invariant survives.
- Causal Stability at Level C (full Lorentzian metric stability).
- Extension of the construction to Level 3 (arbitrary admissible finite posets).
- Proof of Interface Persistence Conjectures 1a, 1b, or 1c.
- That VERSF supplants causal sets, CDT, or Wilsonian RG — only that it adds admissibility closure as a refinement-preserving constraint.

---

12. Conclusion

VERSF proposes that physical reality is not fundamentally continuous, but emerges from admissible irreversible substrate dynamics.

This paper does four things that together turn the framework from gesture into result:

1. **Explicit construction.** Γ is built by midpoint subdivision on the iterated causal diamond, with order preservation, TPB preservation, and refinement-covariant distinguishability all checkable by direct computation.
2. **Explicit linearised RG operator and full d-dimensional spectrum.** $L\_R^{(d)} = (1/(2d))(d \cdot I + A\_\{Q\_d\})$ is defined on the d-dimensional causal diamond for every d ≥ 1. The spectrum is {1, (d−1)/d, …, 0} with binomial multiplicities (Theorem 1), with the

eigenvalue-1 subspace one-dimensional and the corner-alternating mode in the kernel — exact extinction in one step in every dimension.
3. **BCB as flow to a fixed point.** Strict total content is not conserved at any finite refinement step; rather, the content density $\tilde{Q}$ is monotonically non-increasing and converges to a refinement-invariant limit $\tilde{Q}_\infty$, which plays the role of the renormalised BCB coupling.
4. **Generalisability ladder.** Level 1 is proved (in every dimension); Level 2 follows under stated overlap conditions (Proposition 2); Level 3 is open. The d-dimensional causal diamond is the base case, not the claim.

The conceptual core — that continuum physics is the universality class of admissible substrate refinement — survives intact, with one important refinement of its content: in every dimension $d \geq 1$, the eigenvalue-1 subspace on the d-dimensional causal diamond is one-dimensional and contains only the counting invariant. Non-trivial continuum content therefore cannot be supported by bulk modes of any single causal diamond patch, of any dimension. The natural carrier of continuum geometry is interface structure on richer substrates, which is precisely the empirical content of the Interface Persistence Conjectures 1a–1c. The d-cube spectrum computation makes those conjectures more pressing rather than less.

The relation between this construction and the 16^{-n} commutator-decay rate cited elsewhere in the VERSF programme has been clarified: the $d = 4$ spectrum is $\{1, 3/4, 1/2, 1/4, 0\}$ with binomial multiplicities, and 16^{-n} is not an eigenvalue of $L_R^{(4)}$ on the linear fluctuation space. The 16^{-n} rate must therefore arise either from composite-step counting (e.g., $(1/2)^4 = 1/16$ across four refinements of a specific eigensubspace) or from a different operator — perhaps the action of L_R on bilinear or higher-tensor forms, or a still-different construction in the wider programme. Identifying the correct operator and step convention is left open as the appropriate next question for the wider programme.

What This Paper Actually Contributes

The primary contribution of this paper is therefore not a derivation of known physics, but a structural result about admissible substrate dynamics: under repeated admissible refinement, most microscopic bulk structure is washed out, leaving only a narrow class of refinement-stable observables. The paper identifies and classifies that refinement-stable sector for the d-dimensional causal diamond — a one-dimensional eigenvalue-1 subspace spanned by the counting invariant in every dimension $d \geq 1$ — and frames continuum emergence as the persistence class of admissible substrate flows. Adjacent results (the BCB monotone-flow theorem, the Flux-BCB refutation, the Level A/B causal stability) support and constrain this central structural claim.

The Negative Half of the Interface Ontology

The structural result is *one-sided*. The paper proves:

Bulk coarse-graining kills information. Under iterated refinement on the d-dimensional causal diamond, every non-constant bulk fluctuation decays to zero (some at rate $(d-1)/d$ per step, the corner-alternating mode in a single step). Only the counting invariant survives.

But it does not prove the complementary claim:

Interfaces preserve the right information. The codimension-1 antichain structure of admissible substrates carries refinement-stable observables sufficient to encode continuum physics.

The Interface Persistence Conjectures 1a–1c name what this complementary claim would require, but resolving them is beyond the present paper's scope. The asymmetry is structural and worth flagging plainly: this paper establishes the negative half of the interface ontology — that bulk dynamics cannot do the work — and identifies the positive half (interface dynamics is rich enough to do it) as the central open question for the wider VERSF programme. The Flux-BCB refutation in §4.3.2 is an early indication that the positive half is genuinely non-trivial: a natural first guess at an interface-localised conservation law fails on the simplest substrate.

Whether the restricted refinement-stable sector is rich enough to carry the empirical content of continuum physics — General Relativity, quantum mechanics, the Standard Model — therefore remains the appropriate next question. The most promising direction is to compute refinement-stable bilinear and higher-tensor observables localised on codimension-1 antichains of Level 2 substrates, and to test whether their structure connects to the gauge-theoretic and constitutive results of the wider programme.

Spacetime, in this picture, is neither fundamental nor an emergent illusion. It is the refinement-stable sector of admissible substrate dynamics — a universality class of equivalent microscopic realisations, identified not by what it is made of, but by what survives the flow.

Appendix A — Worked Spectrum on the Diamond

We display the matrix form of $L_R = P \circ \Gamma$ on the four-event diamond with one-component ρ and verify the spectrum by direct calculation.

Label events (a, b, c, d) and write $\delta\rho = (\delta a, \delta b, \delta c, \delta d)$. Under one refinement step Γ, four midpoints are inserted with values:

Midpoint	Parent pair	Value
m_{ab}	(a, b)	$\frac{1}{2}(\delta a + \delta b)$
m_{ac}	(a, c)	$\frac{1}{2}(\delta a + \delta c)$
m_{bd}	(b, d)	$\frac{1}{2}(\delta b + \delta d)$
m_{cd}	(c, d)	$\frac{1}{2}(\delta c + \delta d)$

Applying P (definition §6.1, midpoint-only averaging) returns the fluctuation to $\mathbb{C}^4$. For event a, the adjacent midpoints are m_{ab} and m_{ac}, so:

$(P \circ \Gamma)(\delta\rho)(a) = \frac{1}{2}[\frac{1}{2}(\delta a+\delta b)] + \frac{1}{2}[\frac{1}{2}(\delta a+\delta c)] = \frac{1}{4}(\delta a+\delta b) + \frac{1}{4}(\delta a+\delta c) = \frac{1}{2}\delta a + \frac{1}{4}\delta b + \frac{1}{4}\delta c$.

Applying this to every event, one full step of L_R = P ∘ Γ in the basis (δa, δb, δc, δd) is:

L_R = (1/4) ·

2	**1**	**1**	**0**
1	2	0	1
1	0	2	1
0	1	1	2

A.1 Structural Identification

Let A denote the adjacency matrix of the covering-pair graph. The covering pairs of the diamond are (a, b), (a, c), (b, d), (c, d), which form the 4-cycle a–b–d–c–a. Thus

L_R = (1/4)(2I + A)

and spec(L_R) = (1/4)(2 + spec(A)). The 4-cycle adjacency has well-known spectrum {2, 0, 0, −2}, giving

spec(L_R) = (1/4){4, 2, 2, 0} = {1, ½, ½, 0}.

A.2 Eigenvector Verification

Direct application of L_R to the four basis-aligned modes:

- **v_1 = (1, 1, 1, 1).** L_R · v_1 = (1/4)(2+1+1+0, 1+2+0+1, 1+0+2+1, 0+1+1+2) = (1, 1, 1, 1). Eigenvalue **1**.
- **v_2 = (1, 0, 0, −1).** L_R · v_2 = (1/4)(2 − 0, 1 − 1, 1 − 1, 0 − 2) = (½, 0, 0, −½). Eigenvalue ½.
- **v_3 = (0, 1, −1, 0).** L_R · v_3 = (1/4)(1 − 1, 2 − 0, 0 − 2, 1 − 1) = (0, ½, −½, 0). Eigenvalue ½.
- **v_4 = (1, −1, −1, 1).** L_R · v_4 = (1/4)(2 − 1 − 1 + 0, 1 − 2 + 0 + 1, 1 + 0 − 2 + 1, 0 − 1 − 1 + 2) = (0, 0, 0, 0). Eigenvalue **0**.

Trace check: Tr(L_R) = 4 · (2/4) = 2, and 1 + ½ + ½ + 0 = 2 ✓.

A.3 Implications

The one-dimensional eigenvalue-1 subspace contains only the constant mode. The causal-gradient direction v_2 decays as $(½)^n$. The spacelike mode v_3 decays at the same rate. The corner-alternating mode v_4 is in the kernel — extinguished in a single step.

The d-dimensional generalisation is treated in §6.4: for the d-dimensional causal diamond, $L\_R^{(d)} = (1/(2d))(d \cdot I + A\_\{Q\_d\})$ has spectrum {1, (d−1)/d, …, 0} with binomial multiplicities, and the corner-alternating mode is in the kernel of $L\_R^{(d)}$ for every $d \geq 1$. The d = 4 spectrum is explicitly {1, 3/4, 1/2, 1/4, 0} with multiplicities {1, 4, 6, 4, 1}. The 16^{-n} commutator-decay rate cited elsewhere in the VERSF programme [12] is not a single-step eigenvalue of $L\_R^{(d)}$ for any d, and its identification within the wider programme is left open (§6.4.2).

Appendix B — Comparison with Adjacent Programmes

Programme	Substrate	Refinement	Continuum recovery	Status of Lorentz invariance
Causal sets	Partial order + measure	Poisson sprinkling at finer scale	Conformal structure from order; volume from measure	Achieved via Poisson sprinkling
CDT	Causal simplicial complex	Refinement of triangulation	Studied via Monte Carlo; spectral dimension flow	Built in via causal foliation
Wilsonian RG	Lattice field theory	Block-spin / momentum shell	Continuum field theory at fixed point	Built in via Lorentzian action
VERSF (this paper)	Admissible partial order with $\rho : E \to \mathbb{C}^8$	Causal midpoint with density-normalised admissibility preservation	Universality class [Λ]_R of refinement-stable observables under $L\_R = P \circ \Gamma$	**Levels A, B proved at Level 1; Level C (full Lorentz invariance) open**

VERSF's specific contribution to this landscape is the admissibility closure A[ρ] = 0 — read via §4 as a monotone flow to a fixed-point coupling — as a refinement-preserving constraint, which restricts the universality class to substrates carrying BCB and TPB structure. Whether this restriction is sufficient to single out our physical universe is an open question of the wider programme.

References

External literature

[1] Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. D. (1987). "Space-time as a causal set." *Physical Review Letters* **59**, 521–524.

[2] Bombelli, L., Henson, J., and Sorkin, R. D. (2009). "Discreteness without symmetry breaking: A theorem." *Modern Physics Letters A* **24**, 2579–2587.

[3] Ambjørn, J., Jurkiewicz, J., and Loll, R. (2005). "Reconstructing the universe." *Physical Review D* **72**, 064014.

[4] Ambjørn, J., Jurkiewicz, J., and Loll, R. (2005). "Spectral dimension of the universe." *Physical Review Letters* **95**, 171301.

[5] Wilson, K. G. (1971). "Renormalization group and critical phenomena. I. Renormalization group and the Kadanoff scaling picture." *Physical Review B* **4**, 3174–3183.

[6] Wilson, K. G. (1975). "The renormalization group: Critical phenomena and the Kondo problem." *Reviews of Modern Physics* **47**, 773–840.

[7] Goldenfeld, N. (1992). *Lectures on Phase Transitions and the Renormalization Group*. Addison-Wesley / Westview Press.

[8] Cardy, J. (1996). *Scaling and Renormalization in Statistical Physics*. Cambridge University Press.

[9] Biggs, N. (1993). *Algebraic Graph Theory*, 2nd ed. Cambridge University Press. (Cited in Theorem 1 for the d-cube graph spectrum and the smallest-eigenvalue characterisation for d-regular bipartite graphs; see the treatment of distance-regular graphs and the Hamming scheme.)

Additional context (not cited inline but relevant): Henson, J. (2009), "The causal set approach to quantum gravity," in D. Oriti (ed.), *Approaches to Quantum Gravity*, Cambridge University Press; Rovelli, C. (2004), *Quantum Gravity*, Cambridge University Press; Loll, R. (2020), "Quantum gravity from causal dynamical triangulations: a review," *Classical and Quantum Gravity* **37**, 013002.

VERSF programme (Keith Taylor)

References [10]–[15] point to relevant strands of the wider VERSF Theoretical Physics Programme rather than to specific published papers. The full corpus is hosted at versf-eos.com; the topic descriptors below identify which strand of the programme each inline citation refers to. Specific paper titles, version numbers, and DOIs will be added at publication.

[10] Taylor, K. *VERSF Theoretical Physics Programme: K = 7 structural rank derivations.* versf-eos.com.

[11] Taylor, K. *VERSF Theoretical Physics Programme: substrate architecture, BCB/TPB primitives, and the admissibility closure A[ρ] = 0.* versf-eos.com.

[12] Taylor, K. *VERSF Theoretical Physics Programme: non-Abelian commutator decay (source of the 16^{-n} rate cited in §6.4.2).* versf-eos.com.

[13] Taylor, K. *VERSF Theoretical Physics Programme: U(1) gauge transport and Maxwell admissibility from substrate principles.* versf-eos.com.

[14] Taylor, K. *VERSF Theoretical Physics Programme: gravity from tensorial closure and the record-current structure.* versf-eos.com.

[15] Taylor, K. *VERSF Theoretical Physics Programme: derivation of the fine-structure constant from interface geometry.* versf-eos.com.

Related VERSF work (not cited inline but relevant context): The wider programme also develops quantum reconstruction (Born rule and Hilbert-space derivations), commitment dynamics and thermodynamic emergence, cosmological-constant derivation, and constitutive-predictive bridges connecting substrate dynamics to observational predictions. See versf-eos.com for the full corpus.