Refinement Stability on Diamond-Glued Substrates in VERSF

Level 2 Spectra, the Refutation of Scalar Antichain Localisation, and the Forcing of Cohomological Observables

Keith Taylor VERSF Theoretical Physics Programme

General Reader Abstract

In an earlier paper we showed that when you "zoom in" on the simplest possible building block of a discrete physics substrate — a four-event "causal diamond" — almost all the microscopic detail averages away. Only one trivial pattern (a bulk counting invariant) survives the zooming process. We argued this meant the rich physics of the continuous world can't live in arbitrary microscopic structure, but must live in something more constrained — possibly on the boundaries or "interfaces" between regions, rather than in the bulk.

This paper tests that argument by computing what happens when you glue causal diamonds together to make slightly richer substrates. The result is twofold. First, the boundary picture survives the test: across every glued substrate we examine, no new stable patterns appear beyond the trivial counting one. Second — and this is the more interesting result — even patterns that *are* localised on the boundary interfaces fail to survive zooming if they are simply scalar values (a single number per event). The boundary picture therefore can't be carried by ordinary scalar fields; it has to be carried by something **relational** — a quantity that lives on the *connections* between events rather than on the events themselves. This forces the next paper in the programme to work with structured objects (edge-supported transport, with a built-in equivalence-by-gradients) rather than with scalar fields. The mathematical name for the resulting structure is *cohomology*, and we end by showing that this isn't a choice we made for aesthetic reasons — it's the smallest non-trivial object that the previous results leave room for. The substrate forces it on us.

Abstract

We extend the admissible coarse-graining analysis of VERSF from Level 1 substrates (isolated d-dimensional causal diamonds) to Level 2 substrates obtained by gluing finite collections of causal diamonds along compatible faces. The previous Level 1 paper proved that the linearised refinement operator L_R^(d) on the d-dimensional causal diamond has spectrum {1, (d−1)/d, …, 0} with binomial multiplicities, a one-dimensional eigenvalue-1 subspace spanned by the constant counting mode, and exact extinction of the corner-alternating mode. The interface

ontology that paper proposed depends on a specific structural conjecture about Level 2 substrates: that gluing modes have eigenvalue strictly less than 1.

We test that conjecture by explicit computation. We define the degree-weighted lazy random walk $W = \frac{1}{2}(I + D^{-1}A)$ as the natural Level 2 generalisation of the Level 1 operator, and compute its spectrum on four representative substrates: two diamonds glued at a shared vertex (n = 7), two diamonds sharing a middle antichain (n = 6), three diamonds in a chain (n = 10), and a branching three-diamond substrate (n = 9). In every case we display the full spectrum and identify eigenvectors of interest.

The results establish three structural theorems:

(T1) One-dimensional stable subspace. Across all four substrates, the eigenvalue-1 subspace of W is one-dimensional, spanned by the constant counting mode. No gluing mode has eigenvalue 1. The interface ontology's load-bearing dependency is therefore confirmed at the smallest cases.

(T2) Bipartiteness from causal level structure. The covering-pair graph of any admissible glued substrate is bipartite, because the partial order forces a level function and covering pairs change level by exactly one. As a consequence, the smallest eigenvalue of $D^{-1}A$ is -1 in every case, and the corner-alternating mode (the bipartition indicator function) is in the kernel of W. Exact extinction therefore extends from Level 1 to Level 2 by a structural argument rather than a per-substrate check.

(T3) Antichain-supported scalars are refinement-irrelevant. On the two-diamond antichain-glued substrate, the explicit scalar mode $v = (0, 0, 1, -1, 0, 0)$ supported entirely on the shared antichain {b, c} is an exact eigenvector of W with eigenvalue 1/2. This refutes the naive reading of Conjecture 1a — antichain support of arbitrary scalar functions is **not** sufficient for refinement stability. The interface conjecture must be sharpened.

We also report a fourth result, now with closed-form structure: in chains of k vertex-glued diamonds, the second-largest eigenvalue obeys the **exact formula**

$$\lambda_2(k) = \cos^2(\pi/(4k)) = (1 + \cos(\pi/(2k)))/2,$$

with asymptotic $\lambda_2(k) \approx 1 - \pi^2/(16k^2)$. This closed form is not fitted — it follows from a symmetric-sector reduction of W on the chain to a path-graph spectrum. These are the **near-marginal gluing modes** anticipated in the strong-overlap analysis: not exact refinement invariants, but long-lived substrate-memory structures that may carry physical content at finite refinement depth.

The combined result is that scalar bulk and scalar antichain content are both refinement-irrelevant on glued substrates. The interface ontology survives, but **must be reformulated** for structured (non-scalar) observables.

Section 11 carries out that reformulation. The scalar refutations of Sections 5–8 *force* the minimal non-trivial observable sector: it is

$\mathscr{H}^1\_R(\Lambda) = \{$ edge transport fields $\omega : E^{(1)}(\Lambda) \to \mathbb{R} \} / \{ \omega \sim \omega + d\varphi \}$,

the first cohomology group $H^1(G(\Lambda))$ of the covering-pair graph — equivalently, edge variables modulo scalar gradients. This is not an imported construction: the space of scalar gradients Im(d) is parametrised by vertex scalars φ, which lie in a refinement-trivial sector (Proposition 1). The choice of representative ω vs ω + dφ within a cohomology class therefore amounts to a choice within trivialised data; any non-trivial relational observable must be a class in the quotient. The structured sector grows with substrate complexity (dim $\mathscr{H}^1\_R$ = 1, 2, 3, 3, 3 on the five substrates we analyse; 17 on a single 4-dimensional causal diamond), in sharp contrast to the always-1-dimensional scalar sector. The refinement pullback Γ* descends naturally to a well-defined map on cohomology, identifying the next paper's testable empirical question: compute Γ* on H^1 and determine its spectrum.

A striking byproduct (Section 11.7): gauge equivalence, conventionally postulated as a symmetry of nature, here emerges as a substrate-level necessity. Two edge transport fields ω and ω + dφ define the same cohomology class by construction; the substrate-level justification for treating this as physical equivalence is that the parameter φ labelling different representatives is drawn from a refinement-trivial sector. Different gauge choices therefore differ only by data the refinement flow identifies as trivial. The cohomology quotient is gauge-redundancy by force-of-circumstance, not by stipulation.

Table of Contents

---

1. Introduction and Relation to the Level 1 Result

The previous paper, *Admissible Coarse-Graining and Continuum Emergence in VERSF*, established Theorem 1: on the d-dimensional causal diamond $\Lambda^{(d)}$, the linearised refinement operator

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

— the lazy random walk on the d-cube graph Q_d — has spectrum $\{1, (d-1)/d, \ldots, 0\}$ with binomial multiplicities. The eigenvalue-1 subspace is one-dimensional in every d, spanned by the constant counting mode. The corner-alternating mode is in the kernel.

That result has a sharp consequence: non-trivial continuum content cannot be supported by bulk modes of any single d-dimensional causal diamond. The interface ontology — that continuum physics is carried by structures localised on codimension-1 antichains — was offered as the natural reading of this fact, but it depends on a structural conjecture stated explicitly in §6.6 of that paper:

Dependency note. The interface-localisation reading depends on the conjecture that gluing modes on Level 2 substrates have eigenvalue strictly less than 1. A Level 2 substrate exhibiting eigenvalue-1 gluing modes would generate new refinement-stable bulk observables and weaken the interface ontology accordingly.

That dependency has remained untested. This paper tests it.

The structure of the test is straightforward: define the appropriate Level 2 operator (Section 2), prove a structural lemma showing why admissible Level 2 substrates inherit the bipartiteness that drives Level 1's exact extinction (Section 3), then compute the spectrum of W on four representative diamond-glued substrates (Sections 4–7), examine what the results say about antichain-supported scalar modes (Section 8) and the refinement of the Interface Persistence Conjectures (Section 9), and finally — most consequentially — derive the minimal structured observable sector forced by the scalar refutations (Section 11).

The methodology mirrors the discipline of the Level 1 paper. Where v1 → v7 of that paper kept the empirical core (the d-cube spectrum, the Flux-BCB refutation) ahead of the interpretive frame, this paper does the same: every claim in Sections 8–11 is anchored to a specific computed spectrum in Sections 4–7. The §11 derivation in particular — that $\mathscr{H}^{\wedge}\_R(\Lambda)$ is the unique minimal candidate for structured observables — is forced by the scalar results, not chosen for aesthetic reasons.

---

2. The Degree-Weighted Refinement Operator

A Level 2 substrate is a finite admissible partially ordered set Λ obtained by gluing diamond patches U_α along compatible faces. The events of Λ inherit covering relations from the patches, and the **covering-pair graph** G(Λ) has vertices = events and edges = covering pairs.

For the d-dimensional causal diamond, G(Λ^(d)) is the d-cube graph Q_d, which is d-regular. For glued substrates this regularity is generically lost — events at the gluing locus have higher degree than events in the interior of a patch. The Level 1 operator (1/(2d))(d · I + A) does not transfer.

2.1 The Lazy Random Walk

The correct Level 2 generalisation is the **lazy random walk** on G(Λ):

$$W = \tfrac{1}{2}(I + D^{-1}A)$$

where A is the adjacency matrix of G(Λ) and D is the diagonal matrix of vertex degrees. $D^{-1}A$ is the standard random-walk transition matrix; the laziness factor ½ makes the walk aperiodic when G is bipartite, ensuring that exact-extinction phenomena from Level 1 propagate.

W has three properties immediately relevant to the Level 2 spectrum problem:

(P1) Stochasticity on the constant mode. W · 1 = 1, so the all-ones vector is always an eigenvector with eigenvalue 1.

(P2) Reduction to Level 1 on d-regular substrates. When $D = d \cdot I$, we have $D^{-1}A = (1/d)A$, so $W = \frac{1}{2}(I + (1/d)A) = (d \cdot I + A)/(2d) = L_R^{(d)}$, recovering the Level 1 operator exactly. The Level 2 operator extends the Level 1 operator continuously.

(P3) Degree-corrected averaging at the junction. At a high-degree gluing vertex, each neighbour contributes proportionally less weight to the projection, so the gluing locus does not artificially dominate the random-walk dynamics.

2.2 Why the Lazy Random Walk

The non-lazy walk $M = D^{-1}A$ has the same eigenvalue-1 subspace as W (spanned by 1 on connected graphs), but on a bipartite G it also has eigenvalue -1, attained by the bipartition indicator function. The non-laziness means modes of eigenvalue ≈ -1 *oscillate* under iteration rather than decaying — a feature of M but not what we want for an RG-like flow. The laziness term $\frac{1}{2}(I + \cdot)$ maps the M-spectrum from $[-1, 1]$ to $[0, 1]$, so all modes decay monotonically under iteration. The corner-alternating mode (M-eigenvalue -1) is sent to W-eigenvalue 0 — exact extinction in one step.

This is the same construction that gave the corner-alternating extinction on the d-cube in the Level 1 paper. We are simply applying it to a more general graph.

2.3 General Proposition for Connected Substrates

Proposition 1 (Constant-Mode Uniqueness on Connected Glued Substrates). Let Λ be a finite admissible glued substrate with connected covering-pair graph $G(\Lambda)$. Then the eigenvalue-1 subspace of $W = \frac{1}{2}(I + D^{-1}A)$ is one-dimensional, spanned by the constant mode 1.

Proof. $D^{-1}A$ is the transition matrix of a random walk on $G(\Lambda)$. By the Perron–Frobenius theorem applied to the irreducible stochastic matrix $D^{-1}A$ (irreducibility follows from connectedness of G), the eigenvalue 1 of $D^{-1}A$ is simple and the corresponding eigenvector is the all-ones vector 1 (up to scaling). The lazy walk $W = \frac{1}{2}(I + D^{-1}A)$ inherits this: $W \cdot 1 = 1$ and any other eigenvector v of W with eigenvalue 1 would satisfy $D^{-1}A\, v = v$, contradicting simplicity.

Proposition 1 is structurally important: it tells us in advance that no glued substrate, however its diamonds are arranged, can produce a multi-dimensional eigenvalue-1 *scalar* subspace. The Level 1 result extends to Level 2 for scalar modes by graph-theoretic reasons alone.

What Proposition 1 does **not** tell us is what happens to *non-scalar* modes — bilinear forms on covering pairs, antichain-supported transport, edge fluctuations — which are the natural candidates for interface-localised observables. Those require explicit computation, and they will turn out to behave differently. Sections 4–8 carry that out.

---

3. A Structural Theorem: Bipartiteness from Causal Level Structure

Theorem 1 (Bipartiteness of Admissible Glued Substrates). Let Λ be a finite admissible diamond-glued substrate. Then the covering-pair graph $G(\Lambda)$ is bipartite.

Proof. Admissibility requires that Λ is a partially ordered set with a well-defined level function $\ell : E(\Lambda) \to \mathbb{Z}$ such that covering relations $e \lessdot e'$ satisfy $\ell(e') = \ell(e) + 1$. On a single diamond patch, ℓ is the Hamming weight of the binary label. On glued substrates, ℓ is inherited from the patches; gluing two events identifies them only if they share the same level (admissibility condition (i) on the overlap).

Bipartition: partition $E(\Lambda)$ by parity of ℓ. Every covering pair has endpoints of opposite parity (since ℓ changes by 1). Hence $G(\Lambda)$ is bipartite with bipartition by Hamming-weight parity.

Corollary 1 (Corner-Alternating Extinction on Level 2). For any connected admissible Level 2 substrate, the bipartition-indicator vector — the vector p with $p(e) = +1$ if $\ell(e)$ is even and -1 if $\ell(e)$ is odd — is an eigenvector of W with eigenvalue 0.

Proof. On a bipartite d'-regular graph, the bipartition indicator p is the eigenvector of $D^{-1}A$ with eigenvalue -1 (the smallest possible eigenvalue of any transition matrix). For non-regular bipartite graphs, the same holds when $D^{-1}A$ is interpreted entry-wise: $(D^{-1}A \cdot p)(e) = -p(e)$ because every neighbour of e has opposite parity to e, so the average of p over neighbours is $-p(e)$. Hence $W \cdot p = \tfrac{1}{2}(I + D^{-1}A) \cdot p = \tfrac{1}{2}(p - p) = 0$.

Together Proposition 1 and Corollary 1 give a structural pre-result: **on every connected admissible Level 2 substrate, the spectrum of W contains 1 (multiplicity 1, constant mode) and 0 (containing the bipartition-indicator parity vector).** What varies between substrates is the structure of the intermediate spectrum — the eigenvalues strictly between 0 and 1 — and the multiplicities of 0 and other extremal eigenvalues.

We now examine that intermediate structure on four representative substrates.

---

4. Example I — Two Diamonds Glued at a Vertex (Series)

4.1 The Substrate

Two diamonds glued at a single shared vertex. Patch U_1 has events {a, b, c, d}; patch U_2 has events {d, b′, c′, d′}; the event d is the shared "top of U_1 = bottom of U_2" vertex. Total: $n = 7$ events.

Covering pairs:

- From U_1: (a, b), (a, c), (b, d), (c, d)
- From U_2: (d, b′), (d, c′), (b′, d′), (c′, d′)

Total: 8 covering pairs.

Degree sequence: every vertex has degree 2 except the shared vertex d, which has degree 4.

4.2 Spectrum

Direct computation (the explicit 7×7 matrix W is given in Appendix A.1):

Eigenvalue	Multiplicity
1	1
$(2 + \sqrt{2})/4 \approx 0.8536$	1
1/2	3
$(2 - \sqrt{2})/4 \approx 0.1464$	1
0	1

Trace check: $1 + 0.8536 + 3 \cdot (0.5) + 0.1464 + 0 = 3.5 = 7 \cdot (1/2) = \mathrm{Tr}(W)$. ✓

4.3 Interpretation

The eigenvalue-1 subspace is one-dimensional, spanned by the constant mode $(1, 1, 1, 1, 1, 1, 1)/\sqrt{7}$ (Proposition 1). The eigenvalue-0 subspace is also one-dimensional, spanned by the level-parity vector (+1, −1, −1, +1, −1, −1, +1) (Corollary 1) — verified directly: levels are $\ell(a)=0$, $\ell(b)=\ell(c)=1$, $\ell(d)=2$, $\ell(b')=\ell(c')=3$, $\ell(d')=4$, with parity pattern (0,1,1,0,1,1,0) → bipartition (+,−,−,+,−,−,+).

The novel content is the eigenvalues between 0 and 1. The bulk of the eigenvectors decay at rate 1/2 (as on a single diamond), while two "long-range" modes — at eigenvalues $(2 \pm \sqrt{2})/4$ — appear. These are the **gluing modes**: they have non-trivial support at the junction vertex d and represent fluctuations that propagate across the shared vertex between patches. The slower of the two (eigenvalue ≈ 0.85) is the slowest non-trivial decay; it represents the slowest mixing between the two patches.

Both gluing modes have eigenvalue strictly less than 1. The dependency note is confirmed for this substrate. The two "long-range" eigenvalues $(2 \pm \sqrt{2})/4$ are also exactly $\cos^2(\pi/8)$ and $\sin^2(\pi/8)$ — the k = 2 case of the closed-form chain spectrum derived in Section 6, since Example I is precisely a chain of two diamonds.

---

5. Example II — Two Diamonds Sharing a Middle Antichain

5.1 The Substrate

Two diamonds sharing their middle antichain {b, c}. Patch U_1 has events $\{a_1, b, c, d_1\}$; patch U_2 has events $\{a_2, b, c, d_2\}$. The middle antichain {b, c} is identified between the two patches. Total: n = 6 events.

Covering pairs:

- From U_1: (a_1, b), (a_1, c), (b, d_1), (c, d_1)
- From U_2: (a_2, b), (a_2, c), (b, d_2), (c, d_2)

Total: 8 covering pairs.

Degree sequence: a_1, a_2, d_1, d_2 each have degree 2; b and c each have degree 4.

This is the most important test substrate because it is the case where Conjecture 1a is most directly testable: the shared antichain {b, c} is a codimension-1 interface, and if scalar antichain functions are refinement-stable, they should appear here.

5.2 Spectrum

The explicit 6×6 matrix W is given in Appendix A.2. The spectrum is:

Eigenvalue	**Multiplicity**
1	1
1/2	4
0	1

Trace check: $1 + 4\cdot(1/2) + 0 = 3 = 6 \cdot (1/2) = \mathrm{Tr}(W)$. ✓

5.3 Interpretation

The eigenvalue-1 subspace is the global constant. The eigenvalue-0 vector is the level-parity bipartition (+, +, −, −, +, +). The eigenvalue-1/2 subspace is **four-dimensional** — and this is where the action is.

The four-dimensional eigenvalue-1/2 subspace is rich enough that it admits eigenvectors supported in various ways. Of particular interest: it contains the explicit vector

$v = (0, 0, 1, -1, 0, 0)$

which is supported **entirely on the shared antichain {b, c}**. Direct verification:

$W \cdot v = (0, 0, 1/2, -1/2, 0, 0) = (1/2) \cdot v.$

This vector has eigenvalue 1/2, **not 1**. It is therefore refinement-irrelevant — it decays geometrically under iteration of L_R at rate (1/2)^n. Antichain support of arbitrary scalar functions is insufficient for refinement stability.

The same eigenvector can be written more interpretively: it is the **antichain-localised parity mode** that distinguishes b from c within the shared antichain. The natural "geometric" reading of Conjecture 1a — that interface-localised functions survive — fails for this mode at the scalar level. This is the central negative result of the paper, with consequences pursued in Section 8.

---

6. Example III — Three Diamonds in a Chain: Exact Closed Form for the Near-Marginal Spectrum

6.1 The Substrate

Three diamonds joined in series, each sharing a single vertex with the next. Total n = 10 events.

Vertex labels: a (bottom of D_1), b_1, c_1 (middle of D_1), $d_1 = a_2$ (top of D_1 = bottom of D_2), b_2, c_2 (middle of D_2), $d_2 = a_3$ (top of D_2 = bottom of D_3), b_3, c_3 (middle of D_3), d_3 (top of D_3).

Covering pairs: 12 total. Degree sequence: most vertices have degree 2; the two junction vertices d_1, d_2 have degree 4.

6.2 Spectrum

Eigenvalue	Approximate value
1	1.0000
0.9330	(slowest gluing mode)
0.75	
0.5	4-fold
0.25	
0.067	
0	

6.3 The Near-Marginal Mode: Exact Closed Form

The second-largest eigenvalue $\lambda_2 \approx 0.933$ is the slowest non-trivial decay rate on this substrate. The corresponding eigenvector has support across the entire chain with values running monotonically from positive at the bottom to negative at the top, with a single zero crossing at the central antichain $\{b_2, c_2\}$ (level 3 in the reduced path-graph indexing) — the **fundamental** long-range mixing mode.

The eigenvalue is not approximate. The chain of k vertex-glued diamonds admits a symmetric-sector reduction: within each diamond, the two middle vertices b_i and c_i appear symmetrically in W and decouple into a symmetric combination $(b_i + c_i)/\sqrt{2}$ and an antisymmetric combination $(b_i - c_i)/\sqrt{2}$. The antisymmetric combinations live in the eigenvalue-1/2 subspace (as the antichain-localised modes of Example II). The symmetric combinations, together with the junction vertices and the chain endpoints, form a $(2k + 1)$-dimensional reduced operator whose spectrum is exactly that of the lazy random walk on a path graph of length 2k.

Proposition 2 (Closed Form for the Chain Spectrum). Let $\Lambda_{chain}(k)$ be the chain of k vertex-glued d = 2 diamonds. Under the projection to symmetric within-diamond modes, W reduces exactly to the lazy random walk on the path graph P_{2k+1}. The non-trivial eigenvalues are

$$\lambda_j(k) = \cos^2(j\pi/(4k)) = (1 + \cos(j\pi/(2k)))/2, \text{ for } j = 0, 1, \ldots, 2k,$$

with j = 0 giving the constant mode (eigenvalue 1) and j = 2k giving the kernel mode (eigenvalue 0). The remaining eigenvalues at 1/2 with multiplicity k come from the antisymmetric (within-diamond) sector.

Proof. Define the symmetric subspace $S \subset \mathbb{R}^{|E_{chain}(k)|}$ as the subspace of vectors taking equal values on the two middle vertices b_i, c_i of each diamond. S has dimension $2k + 1$, with basis $(v_0, v_1, \ldots, v_{2k})$ indexed by level: v_ℓ takes the value 1 on every vertex at level ℓ and 0 elsewhere.

The orthogonal complement $S^\perp \subset \mathbb{R}^{|E_{chain}(k)|}$ is k-dimensional, with basis $(w_1, \ldots, w_k)$ where w_i takes the value +1 at b_i, −1 at c_i, and 0 elsewhere. Each w_i is supported entirely on a single diamond's antichain $\{b_i, c_i\}$, and by the same calculation as Example II it is an exact 1/2-eigenvector of W: $W \cdot w_i = (1/2) \cdot w_i$.

The action of W preserves S (and hence preserves $S^\perp$). Direct computation gives the reduced operator $W|_S$ on the level basis. For k = 2 (the chain-of-two-diamonds case, n = 7), the reduced operator in basis (level 0, 1, 2, 3, 4) is:

```
          ℓ=0     ℓ=1     ℓ=2     ℓ=3     ℓ=4
ℓ=0:   [ 1/2     1/2      0       0       0   ]
ℓ=1:   [ 1/4     1/2     1/4      0       0   ]
ℓ=2:   [  0      1/4     1/2     1/4      0   ]
ℓ=3:   [  0       0      1/4     1/2     1/4  ]
ℓ=4:   [  0       0       0      1/2     1/2  ]
```

This is precisely the matrix of the lazy random walk $W_{P_5} = \frac{1}{2}(I + D^{-1}A)$ on the path graph P_5 — endpoints (levels 0 and 4) have degree 1 and contribute ½ on the diagonal plus ½ to their single neighbour; interior levels have degree 2 and contribute ½ on the diagonal plus ¼ to each of two neighbours.

The same construction gives W|S = W{P_{2k+1}} for general k. The reason is not that high-degree junction vertices are "averaged away" but a sharper structural fact about matrix elements in the level-indicator basis:

Lemma (Degree Cancellation at Junctions). In the level-indicator basis (v_0, …, v_{2k}) of the symmetric subspace S, the matrix element of W between levels ℓ and ℓ ± 1 is exactly 1/4 for every interior level transition, regardless of whether the source vertex is a degree-2 non-junction vertex or a degree-4 junction vertex.

Proof. For any source vertex u at interior level ℓ, the matrix element (W|S){ℓ±1, ℓ} computed in the indicator basis equals (1/2) · (1/deg(u)) · (number of edges from u to vertices at level ℓ ± 1), where the prefactor 1/2 is the lazy-walk factor.

At a non-junction interior vertex of degree 2 (e.g., a middle vertex b or c of a single diamond): there is exactly 1 edge to the level above and 1 edge to the level below. Matrix contribution to each adjacent level: (1/2) · (1/2) · 1 = 1/4.

At a junction vertex of degree 4 (e.g., a top-of-one-diamond = bottom-of-next vertex d_i): there are exactly 2 edges to the level above (to the two middle vertices of the upper diamond) and 2 edges to the level below (to the two middle vertices of the lower diamond). Matrix contribution to each adjacent level: (1/2) · (1/4) · 2 = 1/4.

The matrix elements agree exactly. The degree-4 junction does not produce a heavier coupling, because the doubled edge count is precisely cancelled by the halved per-edge weight under $D^{-1}A$. The same calculation holds at any admissible junction of arbitrary degree, because admissibility forces junction vertices to connect to equal numbers of vertices in adjacent levels (one b-c pair above, one below). The reduced operator W|S *is therefore the lazy walk on* P{2k+1} as a structural consequence, not as an averaging artefact.

The lazy random walk W_{P_n} on P_n has known eigenvalues: $D^{-1}A$_{P_n} has spectrum cos(jπ/(n−1)) for j = 0, …, n−1, hence W_{P_n} has spectrum (1 + cos(jπ/(n−1)))/2 = $\cos^2$(jπ/(2(n−1))). Substituting n = 2k + 1 gives

$\lambda_j(k) = \cos^2(j\pi/(4k))$, j = 0, 1, …, 2k.

Combining with the k antisymmetric 1/2-eigenvectors gives the full chain spectrum: the (2k + 1) symmetric eigenvalues from S, plus k copies of 1/2 from S^⊥, accounting for all 3k + 1 dimensions of the chain.

The exact form $\lambda_2(k) = \cos^2(\pi/(4k))$ reproduces every digit of the numerical spectrum:

k	$\lambda_2(k)$ computed	$\cos^2(\pi/(4k))$	Match
1	0.5000	$\cos^2(\pi/4) = 1/2$	✓
2	0.8536	$\cos^2(\pi/8) \approx 0.8536$	✓
3	0.9330	$\cos^2(\pi/12) \approx 0.9330$	✓

k	$\lambda_2(k)$ computed	$\cos^2(\pi/(4k))$	Match
4	0.9619	$\cos^2(\pi/16) \approx 0.9619$	✓
5	0.9755	$\cos^2(\pi/20) \approx 0.9755$	✓
8	0.9904	$\cos^2(\pi/32) \approx 0.9904$	✓
10	0.9938	$\cos^2(\pi/40) \approx 0.9938$	✓
20	0.9985	$\cos^2(\pi/80) \approx 0.9985$	✓

Expanding $\cos^2$ near 0:

$\lambda_2(k) = 1 - \sin^2(\pi/(4k)) \approx 1 - \pi^2/(16k^2)$ for large k.

Remark 1 (Mixing Time on Chain Substrates). The slowest decay rate scales as $\pi^2/(16k^2)$, so the **mixing time** — the number of refinement steps required for the slowest non-trivial mode to decay to a fraction ε of its initial magnitude — scales as

$\tau\_mix(k, \varepsilon) = \log(1/\varepsilon) / \log(1/\lambda_2(k)) \approx (16k^2/\pi^2) \cdot \log(1/\varepsilon)$

for large k. This is the standard diffusive scaling $\tau \sim k^2$ for random walks on path-like graphs, and it gives the "coherence band" structure of the wider VERSF programme a parametric mathematical content: long chains support modes with lifetimes scaling quadratically in chain length, with the proportionality constant $16/\pi^2$ determined exactly.

6.4 What the Closed Form Tells Us

The exact formula $\lambda_2(k) = \cos^2(\pi/(4k))$ elevates Section 6 from "empirical observation with a fitted asymptotic" to "exact result derived from a structural reduction." Three consequences worth flagging:

(i) The near-marginal regime is parametrically controlled. For any specified survival threshold $\varepsilon > 0$ and refinement depth n, the minimum chain length k needed for the mode to survive is $k \gtrsim (\pi/4)\sqrt{(n/\log(1/\varepsilon))}$ (rearranging Remark 1). Metastable substrate-memory structures are not just qualitatively possible; their lifetimes are computable.

(ii) The reduction identifies a clean structural object. The chain spectrum factors into (a) a symmetric sector whose spectrum is exactly that of the lazy walk on a path graph of length 2k, and (b) an antisymmetric sector consisting of k independent 1/2-eigenvectors, one per diamond. The path graph reduction is the carrier of all interesting structure on the chain; the antisymmetric sector is trivial. This factorisation is what the next paper can build on.

(iii) Such modes are not exact continuum observables but are physically distinguished. A continuum observable requires $\lambda = 1$ exactly. But physical structures that survive *finite* refinement depth — say $n = 10^6$ steps — only require $\lambda \approx 1$, not $\lambda = 1$. For $k = 1000$, $\lambda_2 \approx 1 - 6\times10^{-7}$, and $\lambda_2\text{^}(10^6) \approx 1/2$: the mode survives a million refinement steps before halving. This is

qualitatively different from a strict eigenvalue-1 mode, but it is also qualitatively different from a 1/2-eigenvector that decays in two steps.

We do not pursue this direction further in this paper, but note that it gives precise mathematical content to the "coherence band" structure of the wider VERSF programme: long-lived gluing modes at λ slightly less than 1 are the substrate-level objects whose effective continuum behaviour could match phenomenology before the asymptotic limit is reached. The closed form $\lambda_2(k) = \cos^2(\pi/(4k))$ makes this regime parametrically tractable.

---

7. Example IV — Branching Substrate (Higher-Degree Junction)

7.1 The Substrate

A "T-shaped" substrate: one base diamond D_1 with two diamonds D_2, D_3 above it, where D_2 and D_3 share one middle vertex with each other (and both sit above D_1's top). Total n = 9 events.

Vertex labels: a (bottom of D_1), b, c (middle of D_1), d (top of D_1), e (other middle of D_2), f (middle vertex shared between D_2 and D_3), g (top of D_2), h (other middle of D_3), i (top of D_3).

Covering pairs: 11 total. Degree sequence: most vertices degree 2; f has degree 3; d has degree 5 (it is the bottom of both D_2 and D_3, with 2 covering pairs into each, plus 2 from D_1).

7.2 Spectrum

Eigenvalue	Approximate value	Multiplicity
1	1.0000	1
0.8655		1
0.75		1
0.5624		1
0.5		1
0.4376		1
0.25		1
0.1345		1
0		1

Trace check: 1 + 0.8655 + 0.75 + 0.5624 + 0.5 + 0.4376 + 0.25 + 0.1345 + 0 = 4.5 = 9 · (1/2) = Tr(W). ✓

7.3 Interpretation

The branching structure produces a richer non-degenerate spectrum than the chain case. The high-degree junction at d (degree 5) generates a slowest gluing mode at $\lambda \approx 0.866$ — similar magnitude to the chain-of-three case, but with a different eigenvector profile.

The branching is the smallest non-trivial test of non-uniform topology, and the result confirms what the chain and antichain examples already established: every gluing mode has eigenvalue strictly less than 1, and the eigenvalue-1 subspace remains one-dimensional. The structural conclusion is robust.

---

8. Antichain-Supported Scalars Are Not Refinement-Stable

The most consequential single result of this paper is the explicit construction in Section 5: the scalar mode

$v = (0, 0, 1, -1, 0, 0)$

on the two-diamond antichain-glued substrate is supported entirely on the shared antichain {b, c}, yet has eigenvalue 1/2 under W. It is refinement-irrelevant.

This **refutes the naive scalar reading of Conjecture 1a**:

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

The naive reading takes this to mean *any* scalar function supported on a maximal antichain is refinement-stable. Section 5 shows otherwise: at least one such function is exactly an eigenvalue-1/2 mode, decaying geometrically.

The conjecture as written is therefore either (a) false, or (b) was always meant in a stronger structured sense than its literal reading suggests. We argue (b): the conjecture survives if we read it as a statement about **structured observables on antichains**, not arbitrary scalar functions.

8.1 The Refutation Is Even Sharper Than It Appears

The antisymmetric antichain mode $v_{asym} = (0, 0, 1, -1, 0, 0)$ is an exact eigenvector of W with eigenvalue 1/2. The corresponding **symmetric** antichain mode

$v_{sym} = (0, 0, 1, 1, 0, 0)$

admits an exact decomposition into the two extremal eigenmodes of W. Writing $\mathbb{1} = (1, 1, 1, 1, 1, 1)$ for the constant mode (eigenvalue 1) and $p = (+1, +1, -1, -1, +1, +1)$ for the level-parity bipartition vector (eigenvalue 0, Corollary 1):

$v_{sym} = \tfrac{1}{2} \cdot \mathbb{1} + (-\tfrac{1}{2}) \cdot p.$

Since $W \cdot \mathbb{1} = \mathbb{1}$ and $W \cdot p = 0$, repeated application of W gives

$W \cdot v_{sym} = \frac{1}{2} \cdot \mathbb{1}$, $W^n \cdot v_{sym} = \frac{1}{2} \cdot \mathbb{1}$ for all $n \geq 1$.

So v_sym lands **exactly** on the constant eigenspace after a single refinement step, then sits there forever. There are no intermediate decaying components — v_sym contains no 1/2-eigenspace contribution at all. It is precisely a 50/50 mix of the two extremal eigenmodes 1 and 0 of W.

Scalar antichain-supported modes therefore fail to be refinement-stable for **two structurally distinct and equally sharp reasons**:

- **Symmetric case** (v_sym): exact 50/50 mix of the constant mode and the parity mode. Under W, the parity component is annihilated in one step, leaving exactly $(1/2)\cdot\mathbb{1}$. The mode loses its antichain localisation completely in one step.
- **Antisymmetric case** (v_asym): is an exact eigenvector with eigenvalue 1/2. Antichain localisation is preserved under refinement, but the magnitude decays geometrically as $(1/2)^n$.

Neither is refinement-stable. The refutation of the naive scalar reading is therefore complete: no scalar function supported on the shared antichain — symmetric, antisymmetric, or any combination — is in the eigenvalue-1 subspace of W. Moreover, the symmetric component is *not even orthogonal* to the eigenvalue-1 subspace; it carries a definite projection onto it.

8.2 Where The Structure Must Live

The distinction is the same one that emerged in the Flux-BCB analysis of the previous paper (§4.3.2): scalar vertex modes and bilinear-on-covering-pair modes are structurally different objects, with different spectral behaviour under refinement. The Flux observable Φ was a bilinear functional of ρ; its decay rate of 1/2 placed it in the eigenvalue-1/2 subspace of the bilinear refinement action. By a parallel that is unlikely to be coincidence, the antichain-localised scalar $(0, 0, 1, -1, 0, 0)$ is *also* a 1/2-eigenvector here.

This convergence — bilinear forms and antichain-supported scalar parities both ending up at eigenvalue 1/2 — suggests that the refinement-stable subspace, if it exists for structured observables, lives at a *different level of structure* than either of these.

The natural question is therefore not "which of several structured-observable candidates to pick," but rather: what is the *smallest* observable structure that the scalar refutations of Sections 5, 6, 8 actually leave room for? Section 11 answers this directly. The smallest non-trivial candidate is *forced* — it is the quotient of edge transport fields by scalar gradients, equivalently the first cohomology of the covering-pair graph. We develop the derivation there.

---

9. Sharpened Interface Persistence Conjectures

The Interface Persistence Conjectures from the previous paper now receive their first round of empirical sharpening.

9.1 Conjecture 1c (Bulk Decay) — *Confirmed for Scalars at Level 2*

The previous paper stated:

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

The original conjecture did not restrict to scalar modes. For scalar modes on connected admissible glued substrates, Conjecture 1c is **proved** by Proposition 1 of Section 2. On any such substrate, the unique eigenvalue-1 scalar mode is the global constant; every non-constant scalar mode has $|\lambda| < 1$, regardless of whether it is bulk-supported or antichain-supported. The structured-observable version of Conjecture 1c — for bilinear, cohomological, or higher-tensor modes — remains open and is the natural next test.

9.2 Conjecture 1a (Antichain Support) — *Sharpened*

The previous paper stated:

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

Section 5 refutes the naive scalar reading. The sharpened version is:

Refined Conjecture 1a. Non-trivial refinement-stable continuum observables, if they exist beyond the counting invariant, are not arbitrary scalar functions on antichains. They are **structured antichain-supported observables** — connection-like 1-forms on covering pairs, cohomological classes modulo scalar gradients, or Wilson-loop-like edge products — whose definition depends on admissible relations *across* the antichain rather than on point values *on* it.

This is testable. The next step is to define the relevant structured observables explicitly and compute their refinement behaviour. We do not do this here; it is the agenda of the next paper.

9.3 Conjecture 1b (Minimum-Triangle Selection) — *Not Yet Tested*

The minimum-triangle-degree selection principle was not testable at the substrates of this paper, because none of them carries multiple maximal antichains of different triangle degrees. A Level 2 substrate with several antichains of different triangle degrees — likely requiring a higher-dimensional glued construction — would be needed.

9.4 The Near-Marginal Refinement

Section 6 establishes that long chains of glued diamonds support modes with eigenvalues arbitrarily close to 1. The closed form $\lambda_2(k) = \cos^2(\pi/(4k))$ gives this regime a parametric

description rather than just a qualitative observation. This adds a new structural element to the Interface Persistence programme:

Near-Marginal Mode Theorem (Chain Substrates). On the chain of k vertex-glued d = 2 diamonds, the slowest non-trivial decay rate of W is exactly $\lambda_2(k) = \cos^2(\pi/(4k))$, with asymptotic $1 - \pi^2/(16k^2)$. The corresponding eigenvector has support across the entire chain with one sign change at the geometric centre — the fundamental long-wavelength mode of the symmetric-sector reduction.

These modes are not refinement-invariant in the strict sense but are quasi-stable at any finite refinement depth scaling polynomially with substrate size. They are the natural carriers of metastable substrate-coherence structures and may correspond physically to the coherence-band phenomena of the wider VERSF programme.

This is a positive finding. It identifies a regime — between exact stability and rapid decay — where physical structure can plausibly live even in the strict Interface Persistence reading. Whether this regime carries the gauge-theoretic objects sketched in Section 10 is the next question.

---

10. Implications for the Gauge-Structure Paper

The gauge-structure paper was the next major target in the programme. The Level 2 results sharpen what that paper must do:

(i) Scalar interface modes are insufficient. Section 5 establishes that scalar antichain support does not produce refinement-stable observables. The gauge paper cannot rest on the picture "scalar functions on antichains carry gauge structure." It must work with structured observables — connections, holonomies, fluxes.

(ii) The structured-observable sector is identified and forced. The picture is no longer "maybe 1-forms, maybe cohomology, maybe Wilson loops." Section 11 derives — from the scalar refutations alone — the unique minimal candidate: $\mathscr{H}^1\_R(\Lambda) = H^1(G(\Lambda))$, edge transport fields modulo scalar gradients. The next paper's first task is therefore not to *choose* an observable class but to compute the refinement behaviour of this specific one.

(iii) The near-marginal regime is now in scope. Section 6 shows that long-lived non-stable modes exist. The gauge-structure paper need not insist on exact refinement invariance for every gauge-theoretic object: some of the physics may live in the $\lambda \approx 1$ regime, where coherence is metastable rather than eternal.

(iv) The bipartite structure is universal — and stronger than the K-S analogue. Theorem 1 shows that all admissible glued substrates are bipartite by causal level structure. This is a strong constraint: any gauge-theoretic construction on a Level 2 substrate inherits a natural "even/odd" bipartition by Hamming-weight parity. The same bipartition appears in lattice-gauge

formulations (staggered fermions, the Kogut–Susskind construction). The structural difference is important: in K–S the bipartition is *imposed* by hand on a chosen hypercubic lattice; here it is *forced* by causal admissibility, with no choice involved. The next paper can therefore exploit staggered-fermion-style constructions without the usual lattice-artefact worries — the bipartition is not a regularisation choice but a substrate-level theorem.

The reframing of the gauge paper's central thesis is:

Not "interfaces carry gauge fields," but "since scalar refinement-stable sectors remain trivial on Level 1 and Level 2 substrates, the next admissible candidate for non-trivial continuum content is structured transport on antichain-supported interfaces — and the bipartite structure of admissible substrates provides a natural framework for such transport."

This is a stronger and more focused claim. The Level 2 paper makes the gauge paper sharper, not premature.

The next section makes the structural force-of-circumstance explicit: it derives the minimal non-trivial observable sector directly from the scalar refutations of Sections 5, 6, 8.

---

11. The Minimal Structured Observable Sector

Sections 5–8 established that scalar refinement-stable observables on Level 2 substrates are essentially trivial. Specifically:

- **(R1)** The only refinement-stable scalar mode is the constant counting mode (Proposition 1).
- **(R2)** Scalar antichain-supported modes — both symmetric and antisymmetric — fail to be refinement-stable. Antisymmetric modes are exact eigenvectors with eigenvalue 1/2; symmetric modes decompose into the constant mode plus the parity mode (eigenvalues 1 and 0), losing antichain localisation in a single step (Section 8.1).
- **(R3)** The space of scalar gradients $\mathrm{Im}(d) \subset C^1(\Lambda)$ is therefore *parametrised by refinement-trivial data*. Any element $d\varphi$ is constructed from a vertex scalar φ drawn from a one-dimensional refinement-stable subspace (the constant) plus a refinement-irrelevant remainder; the choice of representative $d\varphi$ within $\mathrm{Im}(d)$ is correspondingly a choice within trivialised data.

These three results, taken together, do not just *suggest* that non-trivial continuum content must come from structured observables. They **force the minimal non-trivial sector**, and identify exactly what it is. We make this explicit now.

11.1 Edge Variables: The Smallest Non-Vertex Observable

Define a 1-cochain (an **edge variable**) on the substrate as a function

$\omega : E^{(1)}(\Lambda) \to \mathbb{R}$

where $E^{(1)}(\Lambda)$ is the set of directed covering pairs. The value $\omega(e \to e')$ is interpreted as *directed transport* across the admissible refinement edge from e to e′.

This is the smallest relational observable that is *not* reducible to vertex values. A vertex function $\varphi : V(\Lambda) \to \mathbb{R}$ assigns one number per event; an edge function ω assigns one number per ordered pair of adjacent events. Edge variables therefore live in a strictly larger space than vertex variables, with $|E^{(1)}|$ degrees of freedom rather than $|V|$.

This is not yet electromagnetism. It is the minimal structural step away from the scalar sector that the refutations of Sections 5–8 force us to take.

11.2 Scalar Gradient Equivalence

By (R3), the space of scalar gradients Im(d) is parametrised entirely by refinement-trivial scalar data. Two edge variables differing only by a gradient differ only by a choice within that trivialised parameter space; the substrate refinement flow cannot distinguish them. Define the equivalence relation

$\omega \sim \omega + d\varphi$

where $d : C^0(\Lambda) \to C^1(\Lambda)$ is the **coboundary operator**

$(d\varphi)(e \to e') = \varphi(e') - \varphi(e)$.

This equivalence is the substrate-level statement that *adding refinement-trivial information should not change physical content*. The quotient by this equivalence is exactly the structural object that the scalar refutations leave room for.

11.3 The Minimal Structured Sector

Define the **structured-observable sector**:

$\mathscr{H}^1\_R(\Lambda) = \{$ edge transport fields $\omega : E^{(1)}(\Lambda) \to \mathbb{R} \} / \{ \omega \sim \omega + d\varphi \}$

This is the first cohomology group of the covering-pair graph $G(\Lambda)$, denoted $H^1(G(\Lambda))$. By the discrete Euler characteristic, its dimension is

$\dim \mathscr{H}^1\_R(\Lambda) = |E^{(1)}(\Lambda)| - |V(\Lambda)| + c(\Lambda)$,

where $c(\Lambda)$ is the number of connected components of $G(\Lambda)$ (typically 1).

11.4 The Dimension Count Across Substrates

We compute $\mathscr{H}^1\_R$ explicitly for each substrate from the spectral analysis:

| Substrate | $\|V\|$ | $\|E^{(1)}\|$ | dim $\mathscr{H}$_R | |---|---|---|---| | Single diamond ($d = 2$, $\Lambda^{(2)}$) | 4 | 4 | 1 | | Example I: series glue | 7 | 8 | 2 | | Example II: antichain share | 6 | 8 | 3 | | Example III: chain of three | 10 | 12 | 3 | | Example IV: branching | 9 | 11 | 3 | | d-dim diamond $\Lambda^{(d)}$, $d = 4$ | 16 | 32 | 17 | | d-dim diamond $\Lambda^{(d)}$, general | 2^d | $d \cdot 2^{d-1}$ | $(d-2) \cdot 2^{d-1} + 1$ |

The contrast with the scalar sector is sharp. The eigenvalue-1 scalar subspace is **always one-dimensional**, on every substrate (Proposition 1). The structured sector $\mathscr{H}$_R is **always at least as large as the loop count of the substrate** — it grows with substrate complexity, and grows rapidly in higher substrate dimension. On the $d = 4$ causal diamond it is 17-dimensional even on a single patch.

For the antichain-glued Example II (the most physically suggestive case), an explicit basis for $\mathscr{H}$_R can be given. Writing edges in the basis ($a_1 \to b$, $a_1 \to c$, $a_2 \to b$, $a_2 \to c$, $b \to d_1$, $c \to d_1$, $b \to d_2$, $c \to d_2$), and using "→" to indicate forward traversal of a basis edge and "←" to indicate reverse traversal:

- **The a-loop**: the closed walk $a_1 \to b \leftarrow a_2 \to c \leftarrow a_1$ through the two "bottom" events via the shared antichain. Edge vector $\omega_1 = (+1, -1, -1, +1, 0, 0, 0, 0)$.
- **The d-loop**: the closed walk $b \to d_1 \leftarrow c \to d_2 \leftarrow b$ through the two "top" events via the shared antichain. Edge vector $\omega_2 = (0, 0, 0, 0, +1, -1, -1, +1)$.
- **The diagonal**: the closed walk $a_1 \to b \to d_1 \leftarrow c \leftarrow a_1$ crossing the antichain. Edge vector $\omega_3 = (+1, -1, 0, 0, +1, -1, 0, 0)$.

These three classes span $H^1(G(\Lambda\_II)) \cong \mathbb{R}^3$, and each satisfies $d^T \cdot \omega = 0$ (closed). Notice that each cycle is closed *across* the shared antichain $\{b, c\}$, not localised at it — the cohomology classes are inherently transport-like, encoding loop structure rather than point values.

On the single $d = 2$ causal diamond $\Lambda^{(2)}$, dim $\mathscr{H}$_R = 1: the unique non-trivial structured observable is the **plaquette holonomy** around the diamond's square face — the closed walk $a \to b \to d \leftarrow c \leftarrow a$. This is the simplest possible substrate-level analogue of a Wilson loop, and it provides a concrete picture for the abstract H^1 classes on richer substrates. The lattice-gauge analogy of §10(iv) is therefore not just suggestive: on $\Lambda^{(2)}$ the structured sector *is* the plaquette holonomy, full stop.

11.5 The Pullback: How Refinement Acts on $\mathscr{H}$_R

Two notions of "refinement"

Before defining the pullback we make explicit a distinction that the rest of this section relies on. The paper uses two different notions of refinement, applied at different levels of structure:

- **W**: the lazy random walk on a *fixed* substrate Λ. This is the linearised per-substrate refinement operator of Sections 2–7. Its iteration W^n produces the spectral decay analysed throughout the paper. Proposition 1 establishes that W trivialises scalars: the only refinement-stable scalar mode is the constant.

- **Γ***: the *cross-substrate* pullback Γ* : C^k(Λ_{n+1}) → C^k(Λ_n) induced by the substrate-refinement embedding Λ_n ↪ Λ_{n+1} (midpoint insertion). This maps cochains on a finer substrate back to cochains on the coarser substrate it refines.

These are distinct constructions. W acts within a single substrate's cochain space; Γ* moves between substrate cochain spaces. The argument that "scalar gradient data is refinement-trivial" passes through both: W establishes scalar triviality on each substrate, and Γ* preserves gradient structure across substrates (we show this in a moment). Their compatibility — that W contracts the scalar sector while Γ* preserves the gradient subspace it generates — is what makes the cohomology quotient by Im(d) the appropriate object: the quotient is exactly the structure that both notions of refinement treat consistently.

Construction of Γ*

The refinement map Γ : Λ_n → Λ_{n+1} inserts midpoints, so each edge (e, e′) of Λ_n is subdivided into a path (e, m, e′) in Λ_{n+1}. The natural **pullback** of edge cochains under refinement is

$$\Gamma^*\omega(e \to e') := \omega(e \to m) + \omega(m \to e'),$$

i.e., sum the edge variable along the subdivided path. This is the discrete analogue of the pullback of differential 1-forms under refinement of a triangulation.

Γ* commutes with d

Pullback commutes with the coboundary:

$$\Gamma^*(d\varphi)(e \to e') = (d\varphi)(e \to m) + (d\varphi)(m \to e') = (\varphi(m) - \varphi(e)) + (\varphi(e') - \varphi(m)) = \varphi(e') - \varphi(e) = (d\varphi|_{V_n})(e \to e'),$$

so Γ*(Im d_{n+1}) ⊂ Im d_n. The pullback therefore descends to a well-defined map on cohomology:

$$\Gamma^* : H^1(G(\Lambda_{n+1})) \to H^1(G(\Lambda_n)).$$

Cohomology classes pull back to cohomology classes. The structured-observable sector is intrinsically refinement-covariant under Γ*, by construction.

Joint compatibility

We can now state the joint compatibility cleanly. W contracts the scalar sector C^0 on each substrate; Γ* sends gradients to gradients across substrates. The image $\mathrm{Im}(d^0) \subset C^1$ is therefore:

- contracted by W on each substrate (because its parameter space C^0 is contracted), and
- preserved as a subspace by Γ* across substrates (because Γ* maps gradients to gradients).

Both notions of refinement consistently identify $\mathrm{Im}(d^0)$ as the "trivial" subspace of C^1. The cohomology quotient $C^1 / \mathrm{Im}(d^0) = H^1$ is therefore the unique subspace of C^1 that is well-defined under both — and on which the dynamics of either is non-trivial. That joint compatibility is the structural reason cohomology is the right object.

11.6 What This Forces

The construction of §§11.1–11.5 is not an interpretive choice. It is forced — within the standard discrete-geometric structure of cochains and coboundaries — by the scalar refutations of Sections 5–8 in the following sense:

1. **Vertex scalars carry no non-trivial refinement-stable information** (Proposition 1). The largest refinement-stable scalar subspace is one-dimensional.
2. **Therefore the *parameter space* of scalar gradients $d\varphi$ — namely the space of vertex scalars φ — carries no non-trivial refinement-stable information.** Gradients themselves are well-defined edge cochains, but the choice of which gradient (i.e., the choice of φ) specifies only refinement-trivial data.
3. **Any candidate non-trivial observable sector defined within the standard discrete-geometric structure (cochains C^0, C^1 and the coboundary d) must therefore (a) live on objects larger than vertices, and (b) be insensitive to addition of any element of Im(d) — since Im(d) is parametrised by trivialised data.**
4. **The smallest such sector is the quotient $C^1 / \mathrm{Im}(d) = H^1(G(\Lambda))$.**

There is no smaller candidate within this framework. Vertex variables are trivialised; arbitrary edge variables contain a subspace Im(d) parametrised entirely by trivialised data; the cohomology quotient is the unique minimal non-trivial structure compatible with cochain geometry and the scalar refutations.

This is a substrate-level *derivation* of the cohomology construction relative to a specified mathematical framework — not an import from gauge theory. A reader could in principle imagine non-cochain candidate sectors (different equivalence relations on edge variables, or restrictions to particular subspaces), but no such candidate is smaller than H^1 within standard cochain language. Within that framework, the substrate forces the cohomology quotient.

11.7 Implications for Gauge Redundancy

The structured-observable sector $\mathscr{H}^1\_R$ automatically carries a notion of gauge redundancy. Adding a scalar gradient $\omega \to \omega + d\varphi$ leaves the cohomology class unchanged — this is the defining equivalence of H^1. In conventional gauge theory, this is the statement that "physics is invariant under gauge transformations of the connection."

In conventional treatments, this invariance is *postulated*: nature is taken to have gauge symmetry, and physics is required to be invariant under it. Here, the same equivalence emerges as a substrate-level necessity, but for a different and stricter reason:

Substrate origin of gauge equivalence. Two edge transport fields ω and $\omega + d\varphi$ represent the same cohomology class by definition of H^1. The substrate-level justification for treating this equivalence as physical is that the parameter φ — an arbitrary vertex scalar — labels refinement-trivial data (Proposition 1: the only refinement-stable scalar mode is the constant). Different choices of representative ω vs $\omega + d\varphi$ within the same cohomology class therefore differ only by data that the refinement flow itself identifies as trivial. The physical refinement-stable content of an edge transport field is precisely its cohomology class; the choice of representative is a choice of refinement-trivial data, i.e., a *gauge choice*.

This reframing is striking. The conventional reading — "nature has gauge symmetry, physics is invariant under it" — is replaced by a stricter substrate-level reading: the parameter space of would-be physical differences within a cohomology class is exactly the refinement-trivial scalar sector. Gauge redundancy is not chosen for aesthetic or symmetry reasons. It is forced by which degrees of freedom the refinement flow distinguishes.

We do not develop the consequences of this here. The point is structural: the cohomology quotient is the natural carrier of substrate-level transport, and the gauge equivalence is forced by the refinement triviality of the scalar parameter sector rather than postulated as a symmetry of nature.

11.8 What Remains Open

The construction of $\mathscr{H}^1\_R$ is forced; what is *not* yet established is the refinement behaviour of the induced action Γ^* on cohomology classes. There are three possibilities:

(i) *Γ acts as the identity on H^1** — cohomology classes are exactly refinement-stable, and $\mathscr{H}^1\_R$ is the full refinement-stable sector beyond the scalar counting invariant. This would be the cleanest possible outcome.

(ii) *Γ acts as a contraction on H^1** — cohomology classes decay under refinement, in the same way scalar modes did, leaving only the scalar counting invariant. This would mean the refinement programme fails even at the cohomology level, and the interface ontology is more or less empty.

(iii) *Γ decomposes H^1 into refinement-stable and refinement-irrelevant subspaces** — some cohomology classes survive, others decay. This would be analogous to the d-cube spectrum of the previous paper, but for cohomology classes instead of scalar modes.

We do not pursue this here. Computing Γ^* explicitly on the cohomology classes of Examples II–IV is the natural empirical content of the next paper, and it is the test that determines whether the interface ontology has a non-trivial continuum sector at all.

What this paper has done is identify the **right object** for that next computation. The structured sector is no longer "maybe 1-forms, maybe cohomology, maybe Wilson loops"; it is specifically $H^1(G(\Lambda))$, forced by the scalar refutations and equipped with a natural refinement pullback Γ^*.

12. Epistemic Status

In keeping with the discipline of the previous paper:

Proven (within Level 2):

- Proposition 1: The eigenvalue-1 scalar subspace is one-dimensional on any connected admissible glued substrate.
- Theorem 1: All admissible glued substrates have bipartite covering-pair graphs.
- Corollary 1: The bipartition-indicator vector is in the kernel of W on every Level 2 substrate.
- Proposition 2: Exact closed form for the chain-of-k-diamonds spectrum. The slowest non-trivial decay is $\lambda_2(k) = \cos^2(\pi/(4k))$, derived from a clean symmetric-sector reduction of W to the lazy random walk on the path graph P_{2k+1}. The dimension accounting is exact: the (3k+1)-dimensional fluctuation space decomposes as (2k+1) symmetric levels + k antisymmetric within-diamond modes; the symmetric reduction produces P_{2k+1} on the nose, with no degree-induced reweighting at junction vertices.
- The complete spectrum and eigenvector structure on the four test substrates (Sections 4–7), with explicit matrices in Appendix A.
- The antisymmetric antichain-supported scalar mode (0, 0, 1, −1, 0, 0) on the antichain-glued substrate has eigenvalue 1/2 (Section 5, Section 8). Refutes naive scalar reading of Conjecture 1a.
- The symmetric antichain-supported scalar mode (0, 0, 1, 1, 0, 0) admits the exact decomposition $v_{sym} = \frac{1}{2} \cdot \mathbb{1} + (-\frac{1}{2}) \cdot p$, where $\mathbb{1}$ is the constant mode (W-eigenvalue 1) and p is the level-parity bipartition (W-eigenvalue 0); under W it lands exactly on $(1/2) \cdot \mathbb{1}$ in one step and remains there indefinitely (Section 8.1).
- The minimal structured-observable sector $\mathscr{H}_R(\Lambda) = H^1(G(\Lambda))$ is forced by the scalar refutations (Section 11). The pullback Γ^* under refinement descends to a well-defined map on cohomology classes (Section 11.5). Explicit dimensions are computed for all five test substrates (Section 11.4): the structured sector is non-trivial and grows with substrate complexity, in contrast to the always-1-dimensional scalar sector.

Conditional / Empirical:

- The dependency note of the previous paper (§6.6) is confirmed at the smallest cases. Full proof of "no eigenvalue-1 gluing modes on *any* Level 2 substrate" would require generalisation beyond the connectedness argument of Proposition 1 to all admissible-overlap conditions; the test cases here are positive evidence rather than a fully general proof.
- The asymptotic form $\lambda_2(k) \approx 1 - \pi^2/(16k^2)$ for long chains follows from Proposition 2 by Taylor expansion of $\cos^2(\pi/(4k))$ — this is now a derived statement, not a fit.

Conjectural:

- Conjecture 1b (minimum-triangle-degree selection): not testable at the substrates of this paper, awaiting larger constructions.
- Refined Conjecture 1a: structured observables on antichains may be refinement-stable. This is **now sharpened** to a specific testable question — does the induced action Γ^* on $H^1(G(\Lambda))$ preserve, contract, or partially decompose cohomology classes? (Section 11.8 lists the three possibilities; the answer is the empirical content of the next paper.)
- That near-marginal modes correspond physically to coherence-band structures in the wider VERSF programme: a programmatic conjecture, not a mathematical one.

---

13. Conclusion

The Level 1 paper proved that single causal-diamond bulk modes are refinement-irrelevant. The interface ontology was offered as the natural reading, with one explicit dependency: that gluing modes on Level 2 substrates have eigenvalue strictly less than 1.

This paper tests that dependency by explicit computation on four representative Level 2 substrates. The result is unambiguous: on every substrate tested, the eigenvalue-1 scalar subspace is one-dimensional (Proposition 1, structurally guaranteed), every gluing mode has eigenvalue strictly less than 1, the corner-alternating extinction extends to Level 2 via causal bipartiteness (Theorem 1, structurally guaranteed), and antichain-supported scalar modes have eigenvalue 1/2 rather than 1 (Section 5).

The combined result is:

1. **The interface ontology survives** as a programme — no Level 2 substrate generates new eigenvalue-1 scalar modes, so the previous paper's interpretive frame is empirically supported.
2. **But scalar modes are fully exhausted.** Naive scalar antichain support is refuted (Section 5); the symmetric antichain mode decomposes exactly into the constant mode and the parity mode with no surviving antichain localisation (Section 8.1); scalar gradients are therefore refinement-trivial (Section 11.6). Interface Persistence must be reformulated for **structured observables**.
3. **Near-marginal modes have a closed-form structure.** Long-chain glued substrates support quasi-stable modes obeying the exact formula $\lambda_2(k) = \cos^2(\pi/(4k))$, derived from a symmetric-sector reduction to the path-graph spectrum. These are not exact continuum observables but are long-lived enough — with mixing time scaling as k^2 — to plausibly carry metastable physical content. The parametric control makes this regime tractable.
4. **Bipartiteness is universal at Level 2.** Causal level structure forces all admissible glued covering-pair graphs to be bipartite, so the corner-alternating extinction of Level 1 extends automatically. This gives a natural staggered structure that lattice-gauge constructions can exploit.
5. **The minimal structured-observable sector is forced and identified.** Section 11 derives, from the scalar refutations alone, the unique minimal candidate: $\mathscr{H}^1\_R(\Lambda) = H^1(G(\Lambda))$, the first cohomology of the covering-pair graph — equivalently, edge transport

fields modulo scalar gradients. This is not an imported construction. It is the smallest non-trivial object that vertex-scalar triviality leaves room for. The structured sector grows with substrate complexity (dim $\mathscr{H}^1$_R = 1, 2, 3, 3, 3 on the five substrates analysed; 17 on a single 4-dimensional causal diamond), in sharp contrast to the always-1-dimensional scalar sector. The refinement pullback Γ^* descends naturally to a well-defined map on cohomology classes (Section 11.5).

The next paper — *Interface-Supported Gauge Structure in VERSF* — therefore has not only a sharpened mandate but a specific empirical test: compute the spectrum of Γ^* on $\mathscr{H}^1$_R(Λ) for the substrates of this paper. The three possible outcomes (Section 11.8) — identity on H^1, contraction on H^1, or partial decomposition — determine whether the interface ontology has a non-trivial continuum sector at all, and if so, of what dimension.

The substrate-level origin of gauge equivalence is a striking byproduct (Section 11.7). The conventional gauge-theoretic postulate — "physics is invariant under adding a scalar gradient to the connection" — is replaced by a stricter substrate reading: the parameter φ that labels different representatives ω vs $\omega + d\varphi$ within the same cohomology class is itself drawn from a refinement-trivial sector. Different gauge choices therefore differ only by data the refinement flow identifies as trivial. Gauge redundancy is not a postulated symmetry of nature but a structural consequence of *which degrees of freedom the substrate refinement actually distinguishes*.

Spacetime, in this picture, is neither fundamental nor an emergent illusion — and now neither is it carried by scalar bulk or scalar interface structure. It is the refinement-stable sector of admissible substrate dynamics, and what survives the flow is **structured**: relational, transport-like, defined on edges modulo scalar gradients. The Level 2 spectrum has narrowed the search precisely. The cohomology quotient is what comes next.

Appendix A — Explicit Matrices

All matrices below are $W = \frac{1}{2}(I + D^{-1}A)$ for the relevant substrate. Vertex ordering is as specified in the respective section.

A.1 Two Diamonds Glued at a Vertex (Example I)

Vertex order: (a, b, c, d, b′, c′, d′), n = 7.

The transition matrix $D^{-1}A$ (the non-lazy random walk) has rows:

```
a:    [0,   1/2, 1/2, 0,    0,   0,   0  ]
b:    [1/2, 0,   0,   1/2,  0,   0,   0  ]
c:    [1/2, 0,   0,   1/2,  0,   0,   0  ]
d:    [0,   1/4, 1/4, 0,    1/4, 1/4, 0  ]
b′:   [0,   0,   0,   1/2,  0,   0,   1/2]
c′:   [0,   0,   0,   1/2,  0,   0,   1/2]
d′:   [0,   0,   0,   0,    1/2, 1/2, 0  ]
```

Then $W = \frac{1}{2}(I + D^{-1}A)$ has diagonal entries all equal to 1/2 and off-diagonal entries equal to half the corresponding $D^{-1}A$ entry above. Eigenvalues: $\{1, (2+\sqrt{2})/4, 1/2\ (\times 3), (2-\sqrt{2})/4, 0\}$.

A.2 Two Diamonds Sharing an Antichain (Example II)

Vertex order: $(a_1, a_2, b, c, d_1, d_2)$, n = 6.

W explicitly:

	a_1	a_2	b	c	d_1	d_2
a_1:	1/2	0	1/4	1/4	0	0
a_2:	0	1/2	1/4	1/4	0	0
b:	1/8	1/8	1/2	0	1/8	1/8
c:	1/8	1/8	0	1/2	1/8	1/8
d_1:	0	0	1/4	1/4	1/2	0
d_2:	0	0	1/4	1/4	0	1/2

Eigenvalues: $\{1, 1/2\ (\times 4), 0\}$.

A.3 Chain of Three Diamonds (Example III)

Vertex order: $(a, b_1, c_1, d_1, b_2, c_2, d_2, b_3, c_3, d_3)$, n = 10.

The matrix is 10×10; we give the structure. Each diamond contributes a 4×4 block analogous to the Level 1 diamond, with the junction vertices d_1 and d_2 having degree 4 (each is the top of one diamond and the bottom of the next). The full matrix is available in the reproducibility code (Appendix C).

Eigenvalues (to 4 decimal places): {1.0000, 0.9330, 0.7500, 0.5000 (×4), 0.2500, 0.0670, 0.0000}.

A.4 Branching Substrate (Example IV)

Vertex order: (a, b, c, d, e, f, g, h, i), n = 9. The matrix structure has the central vertex d at degree 5 (top of D_1, bottom of both D_2 and D_3) and the shared middle vertex f at degree 3 (middle of D_2 and D_3).

Eigenvalues (to 4 decimal places): {1.0000, 0.8655, 0.7500, 0.5624, 0.5000, 0.4376, 0.2500, 0.1345, 0.0000}.

Appendix B — Eigenvector Tables

B.1 Example II Eigenvectors

For vertex order (a_1, a_2, b, c, d_1, d_2):

λ	Eigenvector	Support pattern
1	(1, 1, 1, 1, 1, 1)/√6	Constant (counting invariant)
1/2	(0, 0, 1, −1, 0, 0)/√2	**Antichain-localised parity** (eigenvalue 1/2, not 1)
1/2	(1, −1, 0, 0, 1, −1)/2	Bulk parity within levels
1/2	(1, 1, 0, 0, −1, −1)/2	Top/bottom asymmetry
1/2	(1, −1, 0, 0, −1, 1)/2	Cross-pattern
0	(1, 1, −1, −1, 1, 1)/√6	Level-parity bipartition

The four-dimensional 1/2-eigenspace admits many bases; the one above is chosen for interpretability. Key observation: the antichain-localised mode (0, 0, 1, −1, 0, 0) is *exactly* an eigenvector — refinement maps it to itself scaled by 1/2, with no leakage to other vertices.

B.2 Example III Slowest Gluing Mode

For the chain of three diamonds with vertex order (a, b_1, c_1, d_1, b_2, c_2, d_2, b_3, c_3, d_3) and levels (0, 1, 1, 2, 3, 3, 4, 5, 5, 6), the eigenvector at $\lambda_2 = \cos^2(\pi/12) \approx 0.9330$ has the explicit profile

$v_{\lambda_2} \approx (+0.426, +0.369, +0.369, +0.213, +0.000, +0.000, -0.213, -0.369, -0.369, -0.426)$

— monotonically decreasing from +0.426 at level 0 to −0.426 at level 6, with the zero crossing at the central antichain {b_2, c_2} (level 3 — the midpoint of the reduced path P_7). There is **one sign change**, located precisely at the geometric centre. This is the **fundamental** (j = 1) mode of the symmetric-sector reduction to a path graph of length 6.

Level-averaged profile (suppressing within-level symmetric/antisymmetric structure):

Level ℓ	**0**	**1**	**2**	**3**	**4**	**5**	**6**
Average value	+0.426	+0.369	+0.213	0.000	−0.213	−0.369	−0.426
Ratio to ℓ=0	1	√3/2	1/2	0	−1/2	−√3/2	−1

The ratios match the closed-form prediction exactly: the j = 1 eigenvector of $D^{-1}A$ on the path graph $P_{2k+1} = P_7$ has profile

$v_{j=1}(\ell) = \cos(\pi\ell/(2k)) = \cos(\pi\ell/6)$, for $\ell = 0, 1, \ldots, 6$,

giving (1, √3/2, 1/2, 0, −1/2, −√3/2, −1). The lazy random walk $W = \frac{1}{2}(I + D^{-1}A)$ shares this eigenvector with eigenvalue $\frac{1}{2}(1 + \cos(\pi/6)) = \cos^2(\pi/12) \approx 0.9330$. (Direct verification on the simpler P_5 case: for k = 2, the j = 1 $D^{-1}A$-eigenvector is $\cos(\pi\ell/4)$ = (1, √2/2, 0, −√2/2, −1), with eigenvalue $\cos(\pi/4) = \sqrt{2}/2$, giving W-eigenvalue $(1 + \sqrt{2}/2)/2 = \cos^2(\pi/8) = (2 + \sqrt{2})/4$ — exactly the slowest non-trivial eigenvalue of Example I.)

For comparison, the **third** eigenvalue $\lambda_3 = \cos^2(\pi/6) = 3/4$ corresponds to the j = 2 path-graph eigenvector with profile $\cos(2\pi\ell/6) = \cos(\pi\ell/3)$, giving (1, 1/2, −1/2, −1, −1/2, 1/2, 1) at levels 0 through 6 — two sign changes, as expected for the second harmonic. The closed form correctly predicts the full chain spectrum $\lambda\_j(3) = \cos^2(j\pi/12)$ for j = 0, …, 6, giving $\{1, \cos^2(\pi/12), 3/4, 1/2, 1/4, \cos^2(5\pi/12), 0\} \approx \{1, 0.9330, 0.75, 0.5, 0.25, 0.0670, 0\}$ for the symmetric sector, plus three additional 1/2-eigenvectors from the antisymmetric within-diamond sector.

B.2.1 Independent Verification on Larger Chains

To rule out any chance that the $\cos(\pi\ell/(2k))$ match at k = 3 is a low-dimensional accident, we verify the same closed-form profile on longer chains. Level-averaged j = 1 eigenvector ratios (normalised to $\ell = 0 \to 1$) versus the predicted values $\cos(\pi\ell/(2k))$:

k = 4 (chain of four diamonds, n = 13, reduced path P_9, predicted $\lambda_2 = \cos^2(\pi/16) \approx 0.9619$):

ℓ	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Computed	1.0000	0.9239	0.7071	0.3827	0.0000	−0.3827	−0.7071	−0.9239	−1.0000
$\cos(\pi\ell/8)$	1.0000	0.9239	0.7071	0.3827	0.0000	−0.3827	−0.7071	−0.9239	−1.0000

k = 5 (chain of five diamonds, n = 16, reduced path P_{11}, predicted $\lambda_2 = \cos^2(\pi/20) \approx 0.9755$):

ℓ	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Computed	1.0000	0.9511	0.8090	0.5878	0.3090	0.0000	−0.3090	−0.5878	−0.8090	−0.9511	−1.0000
$\cos(\pi\ell/10)$	1.0000	0.9511	0.8090	0.5878	0.3090	0.0000	−0.3090	−0.5878	−0.8090	−0.9511	−1.0000

Both chains match the closed form to machine precision across every level. The structural prediction of Proposition 2 — that the slowest non-trivial eigenvector is exactly $\cos(\pi\ell/(2k))$ regardless of chain length — is therefore confirmed across k = 2, 3, 4, 5, and (by Proposition 2's proof) for all $k \geq 1$.

Appendix C — Reproducibility Code

The following Python code reproduces all computations in this paper. Save as `level2_spectrum.py`:

```
import numpy as np

def lazy_walk(A):
    """W = (1/2)(I + D^{-1} A) — the lazy random walk on adjacency A."""
    D = np.diag(A.sum(axis=1))
    return 0.5 * (np.eye(len(A)) + np.linalg.inv(D) @ A)
```

```
def adjacency_from_edges(n, edges):
    A = np.zeros((n, n))
    for i, j in edges:
        A[i, j] = 1; A[j, i] = 1
    return A

# Example I: series glue
A1 = adjacency_from_edges(7,
[(0,1),(0,2),(1,3),(2,3),(3,4),(3,5),(4,6),(5,6)])
print("Example I eigenvalues:", sorted(np.linalg.eigvals(lazy_walk(A1)).real,
reverse=True))

# Example II: antichain glue
A2 = adjacency_from_edges(6,
[(0,2),(0,3),(1,2),(1,3),(2,4),(3,4),(2,5),(3,5)])
print("Example II eigenvalues:",
sorted(np.linalg.eigvals(lazy_walk(A2)).real, reverse=True))

# Verify the antisymmetric antichain-localised mode is an exact 1/2-
eigenvector
W2 = lazy_walk(A2)
v_asym = np.array([0, 0, 1, -1, 0, 0], dtype=float)
print("W·v_asym =", W2 @ v_asym, "  should equal (1/2)v =", 0.5*v_asym)

# Verify the symmetric antichain mode leaks to the constant
v_sym = np.array([0, 0, 1, 1, 0, 0], dtype=float)
print("W·v_sym =", W2 @ v_sym, "  → leaks to (1/2)·constant, not antichain-
localised")

# Example III: chain of three
A3 = adjacency_from_edges(10, [(0,1),(0,2),(1,3),(2,3),
                               (3,4),(3,5),(4,6),(5,6),
                               (6,7),(6,8),(7,9),(8,9)])
print("Example III eigenvalues:",
sorted(np.linalg.eigvals(lazy_walk(A3)).real, reverse=True))

# Verify the closed form λ₂(k) = cos²(π/(4k)) for chains of length k
def chain(k):
    n = 4 + 3*(k-1)
    edges = []
    for i in range(k):
        a = 0 if i == 0 else 3 + 3*(i-1)
        b, c, d = a + 1, a + 2, a + 3
        edges += [(a,b),(a,c),(b,d),(c,d)]
    return adjacency_from_edges(n, edges)

print("\nVerifying Proposition 2: λ₂(k) = cos²(π/(4k))")
for k in [1, 2, 3, 4, 5, 8, 10, 20]:
    eigs = sorted(np.linalg.eigvals(lazy_walk(chain(k))).real, reverse=True)
    closed_form = np.cos(np.pi/(4*k))**2
    print(f"  k={k:2d}: computed λ₂ = {eigs[1]:.8f}, cos²(π/(4k)) =
{closed_form:.8f}, match = {np.isclose(eigs[1], closed_form)}")
```

Running this code reproduces every numerical claim in this paper, including the closed-form verification of Proposition 2.

References

[1] Taylor, K. *Admissible Coarse-Graining and Continuum Emergence in VERSF.* VERSF Theoretical Physics Programme. (The Level 1 paper; this paper extends its Theorem 1 to Level 2 substrates and tests its §6.6 dependency note.)

[2] Lovász, L. (1993). "Random walks on graphs: A survey." In *Combinatorics, Paul Erdős is Eighty*, Vol. 2, 1–46. Bolyai Society Mathematical Studies. (Standard reference for the spectral theory of lazy random walks on bipartite graphs and the Perron–Frobenius simplicity argument used in Proposition 1.)

[3] Biggs, N. (1993). *Algebraic Graph Theory*, 2nd ed. Cambridge University Press. (Cited in the previous paper for the d-cube spectrum; here for the general bipartite-graph extremal-eigenvalue result.)