Refinement Cohomology and the Spectrum of Γ* on Admissible Substrates

Edge Transport, Cohomological Persistence, and the First Non-Trivial Continuum Sector in VERSF

Keith Taylor VERSF Theoretical Physics Programme

General Reader Abstract

Two previous papers established that simple "scalar" patterns — single numerical values assigned to each event in a substrate — cannot carry meaningful continuum physics. Every such pattern either decays under "zooming in" or becomes trivial. The second paper showed that what *might* survive is a richer kind of pattern: not values on events, but values on *connections between events* — what mathematicians call edge transport, modulo a natural equivalence that quotients out the already-trivialised scalar information. The mathematical name for this structure is *cohomology*. The second paper argued that cohomology is the smallest possible structure that could still carry non-trivial physics after the scalar refutations.

This paper performs the test: does cohomology actually survive refinement? The answer is **yes, with a precise scaling**. We show that under the natural refinement pullback, every cohomology class on every substrate we test is preserved as a class, scaled by exactly a factor of 2 per refinement step. The factor 2 reflects that each refinement doubles the edge count along which one would measure transport; it is bookkeeping, not physics. The class itself — the structural object — survives perfectly. This is the first positive emergence result of the programme after a long string of no-go theorems on scalar structure. Cohomological transport is the first non-trivial continuum sector identified in VERSF. The next paper, on substrate gauge structure, can now be written.

Abstract

The previous paper (*Refinement Stability on Diamond-Glued Substrates in VERSF*) proved that scalar refinement-stable observables on Level 1 and Level 2 admissible substrates are essentially trivial: the only such observable is the global constant counting mode. It further derived, from these scalar refutations alone, the minimal non-trivial structured observable sector

$\mathscr{H}^1\_R(\Lambda) := H^1(G(\Lambda)) = C^1(\Lambda) / \mathrm{Im}(d^0)$,

i.e. edge transport fields modulo scalar gradients. The construction was forced: any non-trivial observable defined within the standard cochain framework must mod out by the gradient subspace (which is parametrised by the trivialised scalar sector) and must live on objects larger than vertices; H^1 is the unique minimal candidate.

This paper studies the refinement dynamics of that sector. The cochain pullback

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

induced by midpoint refinement $\Lambda_n \hookrightarrow \Lambda_{n+1}$ commutes with the coboundary d, descends to a well-defined operator on cohomology $\Gamma^* : H^1(G(\Lambda_{n+1})) \to H^1(G(\Lambda_n))$, and pairs canonically with the natural lift $L : H^1(G(\Lambda_n)) \to H^1(G(\Lambda_{n+1}))$ sending each coarse cycle to its subdivided counterpart:

Theorem. *L is an isomorphism, and*

$$\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{H^1(G(\Lambda_n))}.$$

The factor 2 arises structurally — every coarse edge subdivides into two fine edges, and the cochain pullback sums their values. The substantive content is that **L identifies $H^1(G(\Lambda_n))$ with $H^1(G(\Lambda_{n+1}))$ canonically, dim H^1 is preserved exactly, and Γ^* recovers each coarse cohomology class from its fine lift up to the factor 2**. No cohomology classes are created or destroyed by midpoint refinement; the eigenvalues of Γ^* on H^1 (in the L-image basis) are all equal to 2.

This is the first positive emergence result of the VERSF programme after the scalar no-go theorems. The continuum candidate is no longer a scalar field on substrate events but a relational, refinement-persistent structure carried by edge transport modulo scalar gradients. The route to a substrate-derived gauge structure paper is now open: $\mathscr{H}^1_R$ is the carrier, $\Gamma^* = \mathrm{Id}$ on it, and the next paper can impose admissibility, BCB, and TPB constraints on this sector directly.

Table of Contents

---

1. Introduction and Notational Conventions

The previous paper reached two decisive results. The first was negative: on every connected admissible glued substrate tested, the eigenvalue-1 subspace of the linearised refinement operator $W = \tfrac{1}{2}(I + D^{-1}A)$ is one-dimensional, consisting only of the global constant counting mode. Bulk scalar modes decay (Level 1, d-cube spectrum), gluing scalar modes decay (Level 2, every test substrate), and scalar antichain-supported modes decay or collapse to the constant (antisymmetric mode: exact eigenvalue 1/2; symmetric mode: exact decomposition $\tfrac{1}{2}\cdot\mathbb{1} + (-\tfrac{1}{2})\cdot p$, sitting forever on $(1/2)\cdot\mathbb{1}$ after one step). Scalar information beyond the counting invariant is trivialised by refinement.

The second was constructive. Given the scalar refutations, the minimal non-trivial observable sector consistent with standard cochain geometry was *forced* to be

$$\mathscr{H}^1\_R(\Lambda) = C^1(\Lambda) / \mathrm{Im}(d^0) = H^1(G(\Lambda)),$$

the first cohomology of the covering-pair graph. Any candidate sector defined within the cochain framework must (a) live on objects larger than vertices (since vertex variables are trivialised), and (b) be insensitive to addition of any element of $\mathrm{Im}(d^0)$ (since $\mathrm{Im}(d^0)$ is parametrised entirely by the trivialised scalar sector). The cohomology quotient is the unique minimal candidate.

That paper identified the carrier. This paper asks the next question:

Does cohomology survive refinement?

The answer determines the future of the programme. If H^1 survives — if cohomology classes are exactly refinement-stable — then VERSF has a genuine non-scalar continuum carrier, and a substrate-derived gauge-structure paper can be written. If H^1 does not survive — if cohomology classes decay like scalars — then the no-go situation extends to edge transport and the programme must move to higher-dimensional structures (H^2, plaquette cohomology, or something more exotic).

We will show the first outcome holds, with a precise scaling factor. The structural result is *the lift $L : H^1(\Lambda) \to H^1(\Lambda')$ is a canonical isomorphism, and $\Gamma^* \circ L = 2 \cdot Id_{H^1(\Lambda)}$* (Theorem 1). The factor 2 is bookkeeping arising from edge-count doubling under midpoint refinement, not physical decay or growth, and is absorbed by the natural renormalisation $\Gamma^* := (1/2)\Gamma^*$. The substantive content is the canonical identification of H^1 across refinement scales via L.

1.1 Notational Conventions

We adopt the conventions of the two previous papers. A substrate Λ is a finite admissible poset; $G(\Lambda)$ is its covering-pair graph (vertices = events of Λ, edges = covering pairs). The space of k-cochains is

$$C^k(\Lambda) = \{ \omega : E^{(k)}(\Lambda) \to \mathbb{R} \},$$

where $E^{(0)}(\Lambda) = V(\Lambda)$ (vertices) and $E^{(1)}(\Lambda)$ is the set of directed covering pairs. The coboundary operator $d^0 : C^0 \to C^1$ is

$$(d^0\varphi)(e \to e') := \varphi(e') - \varphi(e).$$

The first cohomology is $H^1(G(\Lambda)) = \ker(d^1) / \mathrm{Im}(d^0)$; on a directed graph with no 2-cells specified, $\ker(d^1) = C^1$ and $H^1 = C^1 / \mathrm{Im}(d^0)$. Its dimension is the cycle rank

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

where $c(\Lambda)$ is the number of connected components. For connected Λ this is just $|E| - |V| + 1$.

Throughout, "refinement" means midpoint refinement: the map $\Gamma : \Lambda_n \to \Lambda_{n+1}$ that inserts a midpoint vertex m on every covering pair (e, e') of Λ_n. Each coarse edge $e \to e'$ becomes the two-edge path $e \to m \to e'$ in Λ_{n+1}.

---

2. From Scalar Triviality to the Cohomology Sector

We briefly restate the result of the previous paper that motivates the present analysis.

(R1) *Connectedness ⇒ scalar uniqueness.* On any connected admissible Λ, $\ker(W - I) \cap C^0(\Lambda) = \mathbb{R} \cdot \mathbb{1}$. The only refinement-stable scalar mode is the constant.

(R2) *Antichain modes are exhausted.* The antisymmetric antichain-supported scalar mode $\omega\_asym = (0, 0, 1, -1, 0, 0)$ on the antichain-glued substrate is an exact eigenvector of W with eigenvalue 1/2. The symmetric antichain mode $\omega\_sym = (0, 0, 1, 1, 0, 0)$ admits the exact decomposition $\omega\_sym = \frac{1}{2} \cdot \mathbb{1} + (-\frac{1}{2}) \cdot p$, where $\mathbb{1}$ is the constant mode (eigenvalue 1) and p is the level-parity bipartition (eigenvalue 0). It carries no eigenvalue-1/2 content at all and lands exactly on $(1/2) \cdot \mathbb{1}$ after one step.

(R3) *Im(d^0) is parametrised by trivialised data.* Scalar gradients $d^0\varphi \in C^1(\Lambda)$ are well-defined edge cochains, but the parameter space — the choice of vertex scalar φ — is the trivialised scalar sector of (R1). Different elements of $\mathrm{Im}(d^0)$ differ only by choices within data the refinement flow does not distinguish.

Within the standard cochain framework, any non-trivial refinement-stable observable sector must (a) live on cochains C^k with $k \geq 1$ (since C^0 is trivialised), and (b) be insensitive to addition of any element of $\mathrm{Im}(d^0)$ (since this is trivialised parameter data). The smallest such candidate is

$\mathscr{H}^1\_R(\Lambda) := C^1(\Lambda) / \mathrm{Im}(d^0) = H^1(G(\Lambda))$.

Whether this sector is *actually* refinement-stable is the question of this paper.

---

3. The Refinement Pullback Γ* on Edge Cochains

Let $\Gamma : \Lambda\_n \to \Lambda\_\{n+1\}$ be midpoint refinement. Each coarse edge $e \to e'$ becomes the two-edge fine path $e \to m \to e'$.

3.1 The Natural Lift

The fundamental object pairing $C^1(\Lambda\_n)$ with $C^1(\Lambda\_\{n+1\})$ is not a single map but a pair:

- The **natural lift** $L : C^1(\Lambda\_n) \to C^1(\Lambda\_\{n+1\})$ is defined by $L(\omega)(e\hat{}L) = L(\omega)(e\hat{}R) = \omega(e)$ for each coarse edge $e \in E\hat{}(1)(\Lambda\_n)$ and its two fine halves $e\hat{}L = (e \to m\_e)$ and $e\hat{}R = (m\_e \to e')$. Equivalently: the lift assigns to each fine edge the value of the coarse edge it subdivides.
- The **cochain pullback** $\Gamma^* : C^1(\Lambda\_\{n+1\}) \to C^1(\Lambda\_n)$ goes the other direction, defined in §3.2 below.

The two together form the central machinery of the paper. Theorem 1 will be stated as a relation between them: $\Gamma^* \circ L = 2 \cdot \mathrm{Id}\_\{C^1(\Lambda\_n)\}$.

3.2 Cochain Pullback

The pullback of fine 1-cochains to coarse 1-cochains is

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

for $\omega \in C^1(\Lambda_{n+1})$. This is the **sum-along-subdivided-path** convention. It is the discrete analogue of integrating a 1-form along a refined parametrisation of the same path. The corresponding 0-cochain pullback is restriction $r : C^0(\Lambda_{n+1}) \to C^0(\Lambda_n)$, $\varphi \mapsto \varphi|_{V(\Lambda_n)}$.

Direct calculation: on any lifted cochain,

$(\Gamma^* \circ L)(\omega)(e \to e') = L(\omega)(e \to m) + L(\omega)(m \to e') = \omega(e) + \omega(e) = 2 \cdot \omega(e)$,

so $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{C^1(\Lambda_n)}$. This is the cochain-level statement; the H^1-level version follows once we verify both maps descend to cohomology.

3.3 Γ* Commutes With the Coboundary

For any $\varphi \in C^0(\Lambda_{n+1})$,

$(\Gamma^* d^0_{n+1} \varphi)(e \to e') = (d^0_{n+1} \varphi)(e \to m) + (d^0_{n+1} \varphi)(m \to e') = (\varphi(m) - \varphi(e)) + (\varphi(e') - \varphi(m)) = \varphi(e') - \varphi(e) = (d^0_n (r\,\varphi))(e \to e')$.

So $\Gamma^* \circ d^0_{n+1} = d^0_n \circ r$, i.e. **Γ* sends gradients to gradients**:

$\Gamma^* (\mathrm{Im}\, d^0_{n+1}) \subset \mathrm{Im}\, d^0_n$.

L sends gradients to gradients for a structurally different reason; the verification is given in §3.4.

3.4 Descent to H^1

By §3.3, Γ* descends to a linear map

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

$[\omega] \mapsto [\Gamma^* \omega]$.

The well-definedness is immediate: if $[\omega] = [\omega']$ in $H^1(\Lambda_{n+1})$, then $\omega - \omega' \in \mathrm{Im}\, d^0_{n+1}$, so $\Gamma^*(\omega - \omega') \in \mathrm{Im}\, d^0_n$, hence $[\Gamma^* \omega] = [\Gamma^* \omega']$ in $H^1(\Lambda_n)$.

The lift L also descends to a linear map $L : H^1(G(\Lambda_n)) \to H^1(G(\Lambda_{n+1}))$. To show this we need $L(\mathrm{Im}\, d^0_n) \subset \mathrm{Im}\, d^0_{n+1}$. Given $\varphi \in C^0(\Lambda_n)$, define $\psi \in C^0(\Lambda_{n+1})$ by

$\psi(v) := 2\,\varphi(v)$ for $v \in V(\Lambda_n)$, $\psi(m_e) := \varphi(e) + \varphi(e')$ for the midpoint of $e \to e'$.

Then on each subdivided edge $e \to m \to e'$:

$(d^0_{n+1}\,\psi)(e \to m) = \psi(m) - \psi(e) = (\varphi(e) + \varphi(e')) - 2\,\varphi(e) = \varphi(e') - \varphi(e) = (d^0_n\,\varphi)(e)$,
$(d^0_{n+1}\,\psi)(m \to e') = \psi(e') - \psi(m) = 2\,\varphi(e') - (\varphi(e) + \varphi(e')) = \varphi(e') - \varphi(e) = (d^0_n\,\varphi)(e)$.

So $d^0_{n+1}\,\psi$ takes the same value $(d^0_n\,\varphi)(e) = \varphi(e') - \varphi(e)$ on both halves of every subdivided edge, which is precisely $L(d^0_n\,\varphi)$. Hence $L(d^0_n\,\varphi) = d^0_{n+1}\,\psi \in \mathrm{Im}\, d^0_{n+1}$, and L descends to cohomology.

The two maps therefore form a pair on cohomology:

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

The doubled-vertex-value extension $\psi(v) = 2\,\varphi(v)$ is not merely a technical device — it reflects a structural feature of L that is the cochain-level origin of the factor 2 in Theorem 1. §11.2 develops this in detail; for the present section, descent of both maps to cohomology is what matters, and is established.

3.5 Dimension Preservation Under Midpoint Refinement

Midpoint refinement adds one vertex per coarse edge (the midpoint) and replaces each coarse edge with two fine edges. The change in $(|V|, |E|)$ is therefore

$|V(\Lambda_{n+1})| - |V(\Lambda_n)| = |E^{(1)}(\Lambda_n)|$, $|E^{(1)}(\Lambda_{n+1})| - |E^{(1)}(\Lambda_n)| = |E^{(1)}(\Lambda_n)|$.

The cycle rank is preserved:

$\dim H^1(G(\Lambda_{n+1})) = |E^{(1)}(\Lambda_{n+1})| - |V(\Lambda_{n+1})| + 1 = |E^{(1)}(\Lambda_n)| - |V(\Lambda_n)| + 1 = \dim H^1(G(\Lambda_n))$.

This can also be seen from the topological fact that edge subdivision preserves the homotopy type of a graph, so it preserves H^1. So L and Γ^* both map between equidimensional spaces. Their structural relationship is the content of Theorem 1.

---

4. The Three A Priori Possibilities

Before computing, we enumerate the possible outcomes:

(I) Identity (after renormalisation). $\Gamma^* = \alpha \cdot \mathrm{Id}$ for some $\alpha > 0$, where α may be absorbed by renormalisation $\Gamma^* := \Gamma^*/\alpha$. Cohomology classes are exactly refinement-stable up to overall scaling. *VERSF gains a genuine non-trivial continuum sector.*

(II) Strict contraction. All eigenvalues of Γ^* satisfy $|\lambda| < \alpha$ for any natural normalisation α. Cohomology classes decay like scalars did. *The interface ontology fails even at the structured level; the programme requires higher cohomology.*

(III) Partial decomposition. Γ^* has a non-trivial Jordan decomposition with at least one eigenvalue at the natural normalisation α and at least one $|\lambda| < \alpha$. *Some cohomology classes survive, others decay — analogous to the d-cube spectrum of the previous paper, but at cohomology level.*

We will show outcome (I) holds across every admissible substrate tested, with $\alpha = 2$ (the natural unnormalised pullback eigenvalue) — i.e., $\tilde{\Gamma}^* := (1/2)\Gamma^*$ acts as the identity on every H^1 tested.

---

5. Example I — The Single d = 2 Causal Diamond

5.1 Setup

The d = 2 causal diamond $\Lambda^{(2)}_0$ has four events a, b, c, d with covering pairs

$E^{(1)}(\Lambda^{(2)}_0) = \{ a\to b, a\to c, b\to d, c\to d \}$.

So $|V| = 4$, $|E| = 4$, and $\dim H^1(G(\Lambda^{(2)}_0)) = 4 - 4 + 1 = 1$.

5.2 The Plaquette Class

The unique cohomology class is generated by the **plaquette holonomy** — the closed walk around the diamond's square face:

$a \to b \to d \leftarrow c \leftarrow a$

(where "→" means forward traversal of a basis edge and "←" means reverse traversal). In the edge basis (a→b, a→c, b→d, c→d), this is

$\omega_\square = (+1, -1, +1, -1)$.

One verifies that $\omega_\square$ is non-exact: $\omega_\square \notin \mathrm{Im}(d^0)$. The argument is the standard one. If $\omega_\square = d^0\varphi$ for some scalar φ on the four vertices, then the path-integral of $\omega_\square$ along any directed path from a to d would equal $\varphi(d) - \varphi(a)$, independent of which path is taken. But the path $a \to b \to d$ gives integral $\omega_\square(a \to b) + \omega_\square(b \to d) = (+1) + (+1) = +2$, while the path $a \to c \to d$ gives $\omega_\square(a \to c) + \omega_\square(c \to d) = (-1) + (-1) = -2$. These differ, so no consistent φ exists. The plaquette therefore represents a non-trivial cohomology class.

(On a graph with no specified 2-cells, every 1-cochain is closed — d^1 is the zero map — so closedness is vacuous and only non-exactness distinguishes cohomology classes from gradients.)

5.3 Refinement and Lifted Plaquette

Midpoint refinement produces $\Lambda^{(2)}_1$ with eight events:

V(Λ⁽²⁾*1) = { a, b, c, d, m*{ab}, m_{ac}, m_{bd}, m_{cd} },

and eight directed edges (two per coarse edge):

E^(1)(Λ⁽²⁾*1) = { a→m*{ab}, m_{ab}→b, a→m_{ac}, m_{ac}→c, b→m_{bd}, m_{bd}→d, c→m_{cd}, m_{cd}→d }.

Dimension check: |V| − |E| + 1 = 8 − 8 + 1 = 1. ✓

The natural lifted plaquette is the closed walk

a → m_{ab} → b → m_{bd} → d ← m_{cd} ← c ← m_{ac} ← a,

with edge vector in the ordering above:

ω_□' = (+1, +1, −1, −1, +1, +1, −1, −1).

5.4 The Action of Γ*

The pullback Γ* in matrix form, with rows indexed by coarse edges and columns by fine edges, is

```
              a→m_{ab}  m_{ab}→b  a→m_{ac}  m_{ac}→c  b→m_{bd}  m_{bd}→d
c→m_{cd}  m_{cd}→d
a→b:          [  1         1         0         0         0         0         0
0  ]
a→c:          [  0         0         1         1         0         0         0
0  ]
b→d:          [  0         0         0         0         1         1         0
0  ]
c→d:          [  0         0         0         0         0         0         1
1  ]
```

Applied to ω_□':

Γ* ω_□' = (1+1, −1−1, 1+1, −1−1) = (+2, −2, +2, −2) = 2 · ω_□.

So as cohomology classes:

Γ* [ω_□'] = 2 · [ω_□].

The lifted plaquette pulls back to *twice* the coarse plaquette, not the coarse plaquette itself. In the notation of Theorem 1: L sends [ω_□] to [ω_□'], and Γ* sends [ω_□'] back to 2 · [ω_□], so (Γ* ∘ L)[ω_□] = 2 · [ω_□] — the factor-2 instance of Theorem 1 on the diamond.

5.5 What This Means

On the d = 2 diamond, the cohomology class survives refinement: dim H^1 is preserved (1 → 1), the natural fine generator [ω_□'] maps to a non-zero multiple of the natural coarse generator [ω_□], and the multiple is the structurally natural factor 2 (one for each fine edge subdivided from each coarse edge).

This is the first positive refinement-persistence result after the scalar no-go theorems. The plaquette holonomy — the simplest possible substrate-level Wilson loop — survives refinement up to an explicit scaling factor.

---

6. Example II — Two Diamonds Glued at a Vertex

6.1 Setup

Two d = 2 diamonds glued at a vertex d = a′:

V = { a, b, c, d, b′, c′, d′ }, |V| = 7, E^(1) = { a→b, a→c, b→d, c→d, d→b′, d→c′, b′→d′, c′→d′ }, |E| = 8.

So dim H^1 = 8 − 7 + 1 = 2.

6.2 Cycle Basis

Two independent plaquette classes, one per diamond:

ω_1 = (+1, −1, +1, −1, 0, 0, 0, 0) (lower plaquette: a → b → d ← c ← a) ω_2 = (0, 0, 0, 0, +1, −1, +1, −1) (upper plaquette: d → b′ → d′ ← c′ ← d)

Each is closed and non-exact.

6.3 The Action of Γ*

Refinement produces Λ_1 with |V| = 15, |E| = 16, dim H^1 = 2 (preserved). The lifted plaquettes are the natural eight-edge analogues of each four-edge coarse plaquette.

By direct calculation (or by the structural argument of Section 10), each lifted plaquette pulls back to twice its coarse plaquette:

Γ* [ω_1'] = 2 · [ω_1], Γ* [ω_2'] = 2 · [ω_2].

In the L-image basis (L[ω_1], L[ω_2]) = ([ω_1'], [ω_2']) of H^1(Λ_1), Γ* takes the matrix form

$$\Gamma^* = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$$

= 2 · I_2. Equivalently in canonical form: $\Gamma^* \circ L = 2 \cdot \mathrm{Id}$ on H^1.

6.4 Comparison With the Previous Paper

The previous paper showed that on this substrate, scalar gluing modes have eigenvalues $(2 \pm \sqrt{2})/4 = \cos^2(\pi/8), \sin^2(\pi/8)$ — strictly less than 1, hence refinement-decaying. The cohomology analysis here gives a sharply contrasting picture: **scalar gluing modes decay, but cohomological plaquette classes survive refinement** — the relation $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{\{H^1\}}$ of Theorem 1, equivalently $\Gamma^* \circ L = \mathrm{Id}$.

---

7. Example III — Two Diamonds Sharing an Antichain

7.1 Setup

Two diamonds sharing the middle antichain {b, c}:

V = { a_1, a_2, b, c, d_1, d_2 }, |V| = 6, E^(1) = { a_1→b, a_1→c, a_2→b, a_2→c, b→d_1, c→d_1, b→d_2, c→d_2 }, |E| = 8.

So dim $H^1 = 8 - 6 + 1 = 3$.

7.2 Cycle Basis

Using the forward/reverse traversal convention with edges ordered as listed above:

- **The a-loop**: a_1 → b ← a_2 → c ← a_1, with $\omega_a = (+1, -1, -1, +1, 0, 0, 0, 0)$.
- **The d-loop**: b → d_1 ← c → d_2 ← b, with $\omega_d = (0, 0, 0, 0, +1, -1, -1, +1)$.
- **The diagonal**: a_1 → b → d_1 ← c ← a_1, with $\omega_\Delta = (+1, -1, 0, 0, +1, -1, 0, 0)$.

These three classes span $H^1 \cong \mathbb{R}^3$. Independence in H^1 — not just in C^1 — is verified by computing the rank of the augmented matrix $[d^0 \mid \omega_a \mid \omega_d \mid \omega_\Delta]$. The coboundary matrix d^0 has rank 5 on this substrate (one less than |V| = 6, since Λ is connected); the augmented matrix has rank 8 = |E|. The three additional columns therefore add three independent dimensions modulo $\mathrm{Im}(d^0)$, confirming that $\{[\omega_a], [\omega_d], [\omega_\Delta]\}$ is a basis of H^1. The reproducibility code in Appendix A includes this check.

7.3 The Action of Γ*

By direct calculation:

Γ^*(lifted ω_a) = $(+2, -2, -2, +2, 0, 0, 0, 0) = 2 \cdot \omega_a$ Γ^*(lifted ω_d) = $(0, 0, 0, 0, +2, -2, -2, +2) = 2 \cdot \omega_d$ Γ^*(lifted ω_Δ) = $(+2, -2, 0, 0, +2, -2, 0, 0) = 2 \cdot \omega_\Delta$

So in the L-image basis (L[ω_a], L[ω_d], L[ω_Δ]) of $H^1(\Lambda_1)$:

$$\Gamma^* = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}$$

$= 2 \cdot I_3$. Equivalently in canonical form: $\Gamma^* \circ L = 2 \cdot \mathrm{Id}$ on H^1.

7.4 The Sharp Contrast With Antichain Scalars

This is the **most striking** result of the paper. On this exact substrate, the previous paper proved:

- The antisymmetric scalar antichain mode (0, 0, 1, −1, 0, 0) is an exact eigenvector of W with eigenvalue 1/2.
- The symmetric scalar antichain mode (0, 0, 1, 1, 0, 0) decomposes as $\frac{1}{2}\cdot\mathbb{1} + (-\frac{1}{2})\cdot p$ and lands on $(1/2)\cdot\mathbb{1}$ after one step.

Both scalar antichain modes are refinement-trivial. But the *cohomology classes* defined by closed walks crossing the antichain are exactly refinement-stable (up to the structural factor 2). The structural distinction sharpens:

Scalar information localised on antichains decays. Relational transport that crosses antichains persists.

This is precisely the kind of structural separation the interface ontology of the wider VERSF programme requires.

---

8. Example IV — Chains of k Vertex-Glued Diamonds

8.1 Setup

For a chain of k diamonds glued at successive vertices:

$|V| = 3k + 1$, $|E| = 4k$, $\dim H^1 = 4k - (3k + 1) + 1 = k$.

There are k independent plaquette classes $\omega_1, \ldots, \omega_k$, one per diamond.

8.2 The Action of Γ*

By the same argument as Example II (each plaquette lifts to its natural 8-edge counterpart, the pullback sums pairs, factor of 2 per class):

$\Gamma^*[\omega_i'] = 2 \cdot [\omega_i]$ for each $i = 1, \ldots, k$.

So in the L-image basis (L[ω_1], …, L[ω_k]) of $H^1(\Lambda_1)$:

$\Gamma^* = 2 \cdot I_k$. Equivalently in canonical form: $\Gamma^* \circ L = 2 \cdot \mathrm{Id}$ on $H^1(G(\text{chain of } k))$.

8.3 The Sharp Contrast With Near-Marginal Scalars

The previous paper proved that the chain spectrum of W has near-marginal scalar modes:

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

with corresponding mixing time $\tau_{mix}(k, \varepsilon) \approx (16k^2/\pi^2) \cdot \log(1/\varepsilon)$. These are *quasi-stable* — they survive a long time but eventually decay.

The cohomological plaquette classes on the same substrate are **exactly persistent** under Γ^* (eigenvalue 1 after renormalisation). The structural difference is now precise:

- Scalar long-range memory: near-marginal, decay rate $\pi^2/(16k^2)$.
- Cohomological transport: $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{H^1}$ on the chain (equivalently $\Gamma^* \circ L = \mathrm{Id}$) — no decay at any rate.

The chain substrate exhibits both phenomena simultaneously. Their separation is exact at the operator level: W decays everything off-constant; Γ^* preserves the full H^1.

---

9. Example V — Branching Substrate

9.1 Setup

The branching substrate of the previous paper (§7):

$|V| = 9$, $|E| = 11$, $\dim H^1 = 11 - 9 + 1 = 3$.

The three cycle classes correspond to the three independent loop structures generated by the branching topology.

9.2 The Action of Γ^*

The same structural argument (each cycle lifts to its natural refined analogue with twice the edge count; pullback sums pairs; factor of 2 per class) gives, in the L-image basis of $H^1(\Lambda_1)$:

$\Gamma^* = 2 \cdot I_3$. Equivalently in canonical form: $\Gamma^* \circ L = 2 \cdot \mathrm{Id}$ on $H^1(G(\text{branching}))$.

9.3 Robustness to Topology

This confirms that the $\Gamma^* \circ L = 2 \cdot \mathrm{Id}$ result is not an artefact of linear-chain or simple-glue topology. It holds on the smallest non-trivial branching example, and the structural proof of Section 10 will show it holds universally.

---

10. Main Theorem: Cohomological Persistence Under Midpoint Refinement

10.1 Statement

Theorem 1 (Cohomological Persistence Under Midpoint Refinement). Let Λ be any finite connected admissible substrate, let $\Gamma : \Lambda \to \Lambda'$ be midpoint refinement, and let $L : H^1(G(\Lambda)) \to H^1(G(\Lambda'))$ and $\Gamma^* : H^1(G(\Lambda')) \to H^1(G(\Lambda))$ be the lift and cochain pullback respectively (§3.1–3.4). Then:

(a) L is an isomorphism;

(b) $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{H^1(G(\Lambda))}$;

(c) equivalently, $\Gamma^* = 2 \cdot L^{-1}$ as maps $H^1(G(\Lambda')) \to H^1(G(\Lambda))$.

The factor 2 reflects that each coarse edge subdivides into exactly two fine edges under midpoint refinement. For more general refinements with k-fold edge multiplicity per coarse edge, the analogous result yields $\Gamma^* \circ L = k \cdot \mathrm{Id}$.

10.2 Proof

$\Gamma^* \circ L = 2 \cdot \mathrm{Id}$ at the cochain level. This was computed in §3.2: for any $\omega \in C^1(\Lambda)$,

$(\Gamma^* \circ L)(\omega)(e \to e') = L(\omega)(e \to m) + L(\omega)(m \to e') = \omega(e) + \omega(e) = 2 \cdot \omega(e)$.

So $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{C^1(\Lambda)}$ at the cochain level. Since both Γ^* and L descend to cohomology (§3.4), the same relation holds on cohomology classes:

$(\Gamma^* \circ L)[\omega] = [\Gamma^* L(\omega)] = [2\,\omega] = 2 \cdot [\omega]$ in $H^1(G(\Lambda))$.

This establishes (b) and (c).

L is injective on cohomology. If $L[\omega] = 0$ in $H^1(G(\Lambda'))$, then by (b)

$0 = \Gamma^*(L[\omega]) = 2 \cdot [\omega]$ in $H^1(G(\Lambda))$,

so $[\omega] = 0$. Hence $\ker L = 0$.

L is surjective on cohomology. This follows from injectivity plus dimensional preservation. By §3.5,

$\dim H^1(G(\Lambda')) = \dim H^1(G(\Lambda))$,

so L is an injective linear map between finite-dimensional vector spaces of equal dimension — necessarily a bijection. (Equivalently, edge subdivision preserves homotopy type, so it preserves H^1 on the nose, giving surjectivity directly.) This establishes (a).

10.3 Consequence: Canonical Basis Identification

Combining (a) and (b): the lift L provides a canonical identification of $H^1(G(\Lambda))$ with $H^1(G(\Lambda'))$, and in the resulting bases Γ^* takes the explicit matrix form $2 \cdot I$. Concretely, if $(\gamma_1, \ldots, \gamma_d)$ is a cycle basis of $H^1(G(\Lambda))$ and $(\gamma_1', \ldots, \gamma_d') = (L\,\gamma_1, \ldots, L\,\gamma_d)$ is the corresponding lifted basis of $H^1(G(\Lambda'))$, then in these bases

$\Gamma^* = 2 \cdot I_d$.

This is the matrix form of the universal-eigenvalue claim. The form is basis-dependent in the technical sense that any choice of basis on $H^1(G(\Lambda'))$ other than the L-image basis would give a different — though similar — matrix. The structural content of (a)–(c) is basis-free.

10.4 Corollary: No Cohomology Created or Destroyed

By the preceding subsections (§3.5 for dimensional preservation, §10.2 for the L-isomorphism), $\dim H^1(G(\Lambda')) = \dim H^1(G(\Lambda))$, and the lift L identifies the two cohomology spaces canonically. **No cohomology classes are created or destroyed by midpoint refinement** — every coarse class has a unique fine preimage (the lift) and every fine class has a unique coarse image ($\Gamma^*/2$). The map Γ^* itself is a scaled isomorphism rather than an inclusion or quotient.

This is a stronger structural statement than "Γ^* preserves H^1": under midpoint refinement there is no inflation of the cohomology space, no extra fine-scale cycles appearing that the coarse substrate did not already know about. The cohomology is fully determined at any single scale.

---

11. The Factor of 2 — Structural Origin and Renormalisation Choice

11.1 The Structural Origin of the Factor

The factor 2 in $\Gamma^* \circ L = 2 \cdot \mathrm{Id}$ has a single structural source: midpoint refinement subdivides each coarse edge into exactly **two** fine edges, and the cochain pullback sums their values along the subdivided path. The natural lift assigns equal values to both halves, and the sum doubles.

The factor is therefore not a physical statement about decay or growth — it is a combinatorial property of the refinement convention chosen.

For more general refinements with k-fold edge multiplicity per coarse edge, the analogous result yields $\Gamma^* \circ L = k \cdot \mathrm{Id}_{\{H^1(\Lambda)\}}$. The persistence of the cohomology *structure* is independent of k; only the multiplicative scale changes.

11.2 The Factor 2 at the Cochain Level: Why L Is Not a Chain Map

The factor 2 in Theorem 1 has a deeper explanation than "the sum of two equal values is twice one of them." It is the precise quantitative manifestation of L's failure to be a chain map between (C^0_n, C^1_n, d^0_n) and $(C^0_{n+1}, C^1_{n+1}, d^0_{n+1})$.

A strict chain map between two cochain complexes is a pair (L^0, L^1) of maps with $L^1 \circ d^0_n = d^0_{n+1} \circ L^0$. The natural choice of L^0 to pair with our $L^1 = L$ would be the averaging extension:

$L^0_{avg}(\varphi)(v) := \varphi(v)$ for coarse v; $L^0_{avg}(\varphi)(m_e) := \tfrac{1}{2}(\varphi(e) + \varphi(e'))$.

This is "natural" in the sense that it gives midpoints the average of the values at their endpoints. Now compute, on each fine edge:

$(d^0_{n+1}\, L^0_{avg}(\varphi))(e \to m) = L^0_{avg}(\varphi)(m) - L^0_{avg}(\varphi)(e) = \tfrac{1}{2}(\varphi(e) + \varphi(e')) - \varphi(e) = \tfrac{1}{2}(\varphi(e') - \varphi(e))$,

while

$L(d^0_n\, \varphi)(e \to m) = (d^0_n\, \varphi)(e) = \varphi(e') - \varphi(e)$.

The two differ by exactly a factor of 2. So with the averaging extension L^0_{avg}, the chain-map condition fails *by exactly the same factor 2* that appears in Theorem 1:

$L \circ d^0_n = 2 \cdot d^0_{n+1} \circ L^0_{avg}$ on cochains.

The descent argument of §3.4 uses a *different* extension — $\psi(v) = 2\,\varphi(v)$ at coarse vertices, $\psi(m_e) = \varphi(e) + \varphi(e')$ at midpoints, which is exactly $2 \cdot L^0_{avg}$ — and this works precisely because it absorbs the factor 2 by hand. There is no zero-shift natural extension that makes L a strict chain map; the factor 2 is intrinsic.

This is the cochain-level explanation for Theorem 1. The relation $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{\{C^1\}}$ on cochains decomposes as:

- **Γ^* alone** is a strict chain map ($\Gamma^* \circ d^0_{n+1} = d^0_n \circ r$, §3.3). It contributes the identity factor.
- **L alone** is not a strict chain map; with the natural averaging extension it fails by exactly factor 2 (calculation above). It contributes the factor 2.

- The composition $\Gamma^* \circ L$ therefore acts as $2 \cdot \mathrm{Id}$ at the cochain level.

The renormalised lift $\hat{L} := \tfrac{1}{2} L$ would be a strict chain map (with the averaging extension), and would give $\Gamma^* \circ \hat{L} = \mathrm{Id}$ at the cochain level. Equivalently, the renormalised pullback $\hat{\Gamma}^* := \tfrac{1}{2}\Gamma^*$ gives $\hat{\Gamma}^* \circ L = \mathrm{Id}$ at the cochain level. The renormalisation can be placed at either map; the structural fact is that exactly one factor of ½ must be absorbed somewhere.

This asymmetric placement of the renormalisation factor between L and Γ^* is structurally invisible at the level of the current paper — both choices give the same cohomology spectrum and the same canonical isomorphism. But the next paper (on substrate gauge structure) imposes BCB and TPB constraints, of which BCB lives naturally on transport (1-cochains) and TPB on cycles (H^1). Whether the constraints couple cleanly to $\hat{L} = \tfrac{1}{2}L$ or to $\hat{\Gamma}^* = \tfrac{1}{2}\Gamma^*$ will not be obvious until the constraint forms are written down explicitly, and the structural ambiguity flagged here may need resolving at that point. We record it as a hook for the next paper rather than a defect of this one.

11.3 Renormalisation: A Choice With Three Defenses

Define the **renormalised refinement pullback**

$\hat{\Gamma}^* := (1/2) \cdot \Gamma^*$.

Then on every test substrate $\hat{\Gamma}^* \circ L = \mathrm{Id}_{H^1}$, i.e. $\hat{\Gamma}^*$ acts as the identity on $H^1(G(\Lambda'))$ once $H^1(G(\Lambda'))$ is identified with $H^1(G(\Lambda))$ via L. We give three defenses of this renormalisation, in order of strength.

(i) Spectral structure is invariant under overall renormalisation. This is the structurally cleanest argument. The eigenvalues of Γ^* in the L-image basis are all 2; the eigenvalues of $\hat{\Gamma}^*$ are all 1. They differ by an overall scale. Any statement that depends only on the *shape* of the spectrum — whether all eigenvalues are equal, whether the spectrum has refinement-stable and refinement-decaying sectors, whether eigenvalues lie on the unit circle — is invariant under the rescaling. The structural content of Theorem 1 (dim H^1 preserved; spectrum is $\alpha \cdot I_d$ for a uniform constant α; no decomposition into stable and decaying subspaces) is intrinsic.

(ii) Inverse-system regularity. The sequence of substrates $\Lambda_0 \to \Lambda_1 \to \Lambda_2 \to \ldots$ under repeated midpoint refinement, equipped with Γ^*-pullback, forms an inverse system whose inverse limit is the candidate "continuum cohomology" of the programme. The natural fine representative $\omega_n \in C^1(\Lambda_n)$ at refinement depth n of a fixed coarse cohomology class $[\omega_0]$ satisfies $(\Gamma^*)^n \omega_n = 2^n \cdot \omega_0$, so the eigenvalue under iterated coarse-graining grows as 2^n — a bookkeeping divergence with no physical content. With $\hat{\Gamma}^*$ renormalisation, the eigenvalue is 1 at every depth, and the inverse limit is well-defined as a vector space (rather than as a sequence with diverging spectral content).

(iii) Compatibility with a natural metric, if one is chosen. The bare substrate of the previous papers has no intrinsic edge-length structure. However, the most natural *choice* of length-rescaling consistent with midpoint refinement assigns each fine edge half the "length" of the

coarse edge it subdivides. Under this convention, integrating a 1-form along the subdivided path with each fine edge weighted by its length gives ½ · [ω(e → m) + ω(m → e′)] = Γ̃*ω. So Γ̃* is the pullback under integration with the natural midpoint-refinement metric.

We include this third defense because it provides geometric intuition — the continuum limit of midpoint refinement, with edges shrinking like 2^{-n} in length, gives the standard Riemann integral picture — but stress that **it depends on a metric choice that is not deducible from the substrate's defining structure**. The substrate as defined in the previous papers is metric-free; assigning fine edges half the coarse length is a choice compatible with midpoint refinement but not forced by it. The structural content of the paper (§§3–10, including the cochain-level chain-map analysis of §11.2) is invariant under all consistent renormalisation choices; only the explicit numerical scaling depends on this choice.

For the purposes of the next paper, the substantive content of this paper is therefore:

- **dim H^1 is preserved under midpoint refinement** (intrinsic, no choice involved);
- **the lift $L : H^1(\Lambda) \to H^1(\Lambda')$ is an isomorphism** (intrinsic);
- **$\Gamma^* \circ L = k \cdot \mathrm{Id}$ where k is the edge-multiplicity factor of the refinement** (intrinsic; k = 2 for midpoint, but the structure is intrinsic);
- **the renormalised operator $\tilde{\Gamma}^* := (1/k)\Gamma^*$ acts as Id on H^1** (a choice, but a natural and well-justified one).

11.4 Three Honest Framings of the Same Result

The same structural fact admits three equally valid formulations:

- **Unnormalised:** $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{H^1(\Lambda)}$ on every admissible Λ.
- **Renormalised:** $\tilde{\Gamma}^* := (1/2)\,\Gamma^*$ satisfies $\tilde{\Gamma}^* \circ L = \mathrm{Id}_{H^1(\Lambda)}$.
- **Structural:** Midpoint refinement induces a canonical isomorphism $L : H^1(\Lambda) \cong H^1(\Lambda')$, with Γ* recovering the coarse class from the fine class up to the factor 2.

This paper adopts the *structural* framing as primary. The lift L is the carrier of the canonical comparison between H^1 at different refinement scales — the object that the next paper's BCB and TPB constraints will need to respect. The factor 2 (or 1, depending on renormalisation) is a derived numerical feature, not the structural heart of the result.

---

12. Consequence: Scalars Trivialise, Cohomology Persists

We now have, across the three-paper sequence, a structurally complete picture of refinement on admissible substrates.

Sector	Operator	Result	Status
Vertex scalars (bulk)	W on C^0	$\ker(W - I) = \mathbb{R}\cdot\mathbb{1}$ only	Trivialised

Sector	Operator	Result	Status
Vertex scalars (antichain)	W on C^0	Antichain modes decay or collapse	Trivialised
Scalar gradients $\mathrm{Im}(d^0) \subset C^1$	(parametrised by C^0)	Parameter space trivialised	Refinement-trivial
Cohomology $H^1 = C^1/\mathrm{Im}(d^0)$	Γ^* on H^1	$\Gamma^* = \mathrm{Id}$ on every test substrate	**Persistent**

The structural conclusion is now sharp:

Scalar information is refinement-trivial. Cohomological transport is refinement-persistent.

This is the first positive emergence result of the programme. After two papers establishing systematic no-go results for scalar substrate observables, we have a non-trivial sector that survives refinement exactly. The continuum candidate is no longer a scalar field on substrate events. It is a refinement-persistent relational structure: edge transport modulo scalar gradients, with the equivalence forced by scalar trivialisation and the persistence verified by direct computation.

---

13. Interpretation: Why This Looks Like Gauge Structure

The structure of $H^1(G(\Lambda))$ is mathematically identical to the structure of an Abelian gauge potential:

- 1-cochains $\omega \in C^1(\Lambda)$ are the substrate-level analogue of connection 1-forms A.
- The coboundary $d^0\varphi$ is the substrate-level analogue of an exact gauge transformation $d\chi$.
- The equivalence $\omega \sim \omega + d^0\varphi$ is the substrate-level analogue of $A \sim A + d\chi$.
- The cohomology class $[\omega]$ is the substrate-level analogue of the gauge-equivalence class $[A]$.
- The closed-walk integral $\oint_\gamma \omega$ is the substrate-level analogue of the Wilson loop $\oint_\gamma A$.

13.1 What This Section Does Not Claim

Before stating what *is* established, we mark explicitly what is **not**:

- We do not derive the gauge group. The construction here gives ω an additive $\mathbb{R}$-structure (i.e., Abelian and one-dimensional), but whether this is the only group consistent with admissibility is open; the next paper must address this.
- We do not derive gauge dynamics. There is no equation of motion on H^1 in this paper; the BCB and TPB constraint analysis on H^1 is the subject of the next paper.
- We do not establish that the equivalence $\omega \sim \omega + d^0\varphi$ is the *only* physically consistent equivalence — only that it is compatible with the refinement structure and naturally selected by scalar trivialisation.

13.2 What This Section Does Claim

What this paper establishes is more modest but still structurally important. The substrate refinement structure is **compatible with — and naturally selects — exactly the gauge equivalence familiar from U(1) connection theory**. The reasoning:

- The minimal observable sector forced by the scalar refutations (previous paper, §11) is $C^1/\mathrm{Im}(d^0)$, with equivalence $\omega \sim \omega + d^0\varphi$ exactly the substrate analogue of the gauge equivalence.
- The parameter space of equivalences — namely C^0, the space of vertex scalars φ — is the trivialised scalar sector. Different "gauge choices" within an equivalence class differ only by data the refinement flow does not distinguish.
- This sector is refinement-persistent (Theorem 1): $\dim H^1$ is preserved and the natural lift identifies $H^1(\Lambda)$ canonically with $H^1(\Lambda')$.

Combining these: the substrate selects an equivalence relation on edge transport whose mathematical form coincides with the U(1) gauge equivalence, whose parameter space is exactly the refinement-trivial scalar sector, and whose equivalence classes are refinement-stable. This is a strong structural compatibility result.

Whether this compatibility rises to a *derivation* of gauge structure — including the gauge group, the dynamical equations, and the variational principle — depends on the BCB and TPB constraint analysis of the next paper. The current paper provides the carrier; the next paper must demonstrate that imposing admissibility and BCB/TPB on this carrier produces the gauge-theoretic structure observed in physics.

The substantive content of the current paper for the gauge programme is therefore:

Cohomological edge transport modulo scalar gradients is the carrier of the first non-trivial continuum sector in VERSF. It is refinement-persistent, with $H^1(\Lambda)$ canonically isomorphic across all refinement scales via the lift L. The familiar gauge-equivalence structure of U(1) connection theory is *consistent* with — and naturally selected by — the substrate refinement flow on this carrier. Whether substrate dynamics produces actual gauge dynamics is the subject of the next paper.

---

14. Epistemic Status

Proven (within this paper)

- Γ^* commutes with d^0 (§3.3). Descends to H^1 (§3.4).
- $\dim H^1$ is preserved under midpoint refinement (§3.5).
- On every test substrate (single diamond, vertex-glued chain, antichain-glued pair, k-chain, branching), $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{\{H^1(\Lambda)\}}$ holds exactly, where L is the natural lift sending each coarse cycle to its subdivided fine counterpart.

- The eigenvalue 2 is a structural consequence of edge-count doubling (Theorem 1, §10.2), independent of substrate topology.
- The renormalised operator $\tilde{\Gamma}^* = (1/2)\Gamma^*$ acts as Id on every H^1 tested.
- The renormalisation factor is a normalisation choice, not a physical statement (§11.3).

Conditional

- The matrix form $\Gamma^* = 2 \cdot I_d$ displayed in §10.3 holds specifically in the L-image basis. Other bases on $H^1(G(\Lambda'))$ give similar matrices with the same spectrum but different off-diagonal entries. The *basis-free* Theorem 1 ($\Gamma^* \circ L = 2 \cdot Id_{H^1(\Lambda)}$) has no such basis dependence — it is an unconditional structural statement.
- We do not address whether more general (non-midpoint) refinements give the same eigenvalue 2 or different scalings. The midpoint construction is the natural one for the admissible-substrate framework of the previous papers.

Open

- **Higher cohomology.** For substrates with intrinsic 2-cell structure (e.g., 4-dimensional causal diamonds with plaquette 2-cells), the analogous analysis on H^2 is the natural next question. Whether H^2 persists or decays under Γ^* is the subject of the next paper.
- **Fine-scale cycle creation.** Midpoint refinement does not create new cohomology classes (Corollary 10.3), but more general admissible refinements could. Classification of refinement-created cohomology is open.
- **The Maxwell admissibility question.** Whether imposing BCB and TPB constraints on the persistent H^1 sector yields Maxwell-like dynamics is the explicit programme of the subsequent paper.
- **Non-Abelian transport.** The construction here is intrinsically Abelian (it lives on the additive group of edge cochains). Substrate-level non-Abelian transport requires either matrix-valued cochains or a different construction; whether such structure survives refinement is open.

---

15. Conclusion

The scalar sector reached a dead end in the previous two papers: vertex scalars are trivialised by the linearised refinement W, antichain-supported scalars decay or collapse to the constant, and scalar gradients are parametrised entirely by the trivialised scalar sector. No scalar structure beyond the global counting invariant carries non-trivial refinement-stable content.

This paper shows that the *next* minimal candidate — cohomological edge transport modulo scalar gradients — behaves completely differently. The refinement pullback Γ^* commutes with the coboundary, descends to a well-defined operator on $H^1(G(\Lambda))$, and pairs with the canonical natural lift L to satisfy $\Gamma^* \circ L = 2 \cdot Id_{H^1(\Lambda)}$ on every coarse class. The factor 2 is bookkeeping arising from edge-count doubling under midpoint refinement; the substantive content is structural:

The lift $L : H^1(\Lambda) \to H^1(\Lambda')$ is a canonical isomorphism, dim H^1 is preserved exactly, and every cohomology class is uniquely identified across refinement scales via L.

This is the first positive refinement-persistence theorem of the programme. The structural separation is now precise: at every level of the analysis — bulk modes, gluing modes, antichain modes, near-marginal modes — scalar information decays or is trivialised, but cohomological transport persists exactly. The continuum carrier of the programme is therefore not a scalar field on substrate events but the cohomology sector $H^1(G(\Lambda))$ of refinement-persistent edge transport.

The next paper, *Interface-Supported Gauge Structure in VERSF*, can now be written from this foundation. $\mathscr{H}^1$_R is the carrier; $\Gamma^* = \mathrm{Id}$ is its refinement law; the BCB and TPB admissibility constraints can now be imposed on this sector directly; and the substrate-level analogues of gauge transformations, Wilson loops, and Maxwell-like dynamics can be studied on a precisely defined and verifiably persistent structure.

The route from substrate refinement to gauge-like relational structure is now open:

scalar no-go → cohomology forcing → cohomology persistence → admissible gauge transport.

Appendix A — Reproducibility Code

```
import numpy as np

def boundary(omega, edges, n_v):
    """∂ω : C¹ → C⁰, takes edge values to net vertex flux."""
    result = np.zeros(n_v)
    for k, (i, j) in enumerate(edges):
        result[j] += omega[k]
        result[i] -= omega[k]
    return result

def coboundary_matrix(n_v, edges):
    """d⁰ : C⁰ → C¹ as an |E| × |V| matrix."""
    d = np.zeros((len(edges), n_v))
    for k, (i, j) in enumerate(edges):
        d[k, j] = 1
        d[k, i] = -1
    return d

def dim_H1(n_v, edges):
    """Compute dim H¹(G) = |E| − rank(d⁰)."""
    return len(edges) - np.linalg.matrix_rank(coboundary_matrix(n_v, edges))

def midpoint_refine(n_v, edges):
    """Subdivide each edge by adding a midpoint vertex."""
    new_edges = []
    for k, (u, v) in enumerate(edges):
```

```
        m = n_v + k
        new_edges.append((u, m))
        new_edges.append((m, v))
    return n_v + len(edges), new_edges

def build_pullback(n_edges_coarse):
    """Γ* matrix: each coarse edge i sums fine edges 2i and 2i+1."""
    G = np.zeros((n_edges_coarse, 2 * n_edges_coarse))
    for i in range(n_edges_coarse):
        G[i, 2*i] = 1
        G[i, 2*i+1] = 1
    return G

def lift_cycle(omega_coarse):
    """Each coarse edge value v becomes (v, v) on its two refined halves."""
    omega_fine = np.zeros(2 * len(omega_coarse))
    for i, v in enumerate(omega_coarse):
        omega_fine[2*i]     = v
        omega_fine[2*i + 1] = v
    return omega_fine

# ---------------- Example: single diamond ----------------
print("=" * 50)
print("Example I: single d=2 causal diamond")
print("=" * 50)
n_v_0 = 4
edges_diamond = [(0,1), (0,2), (1,3), (2,3)]
print(f"  Λ_0: |V|={n_v_0}, |E|={len(edges_diamond)}, dim H¹ = {dim_H1(n_v_0,
edges_diamond)}")
assert dim_H1(n_v_0, edges_diamond) == 1

n_v_1, edges_1 = midpoint_refine(n_v_0, edges_diamond)
print(f"  Λ_1: |V|={n_v_1}, |E|={len(edges_1)}, dim H¹ = {dim_H1(n_v_1,
edges_1)}")
print(f"  Dimension preserved: {dim_H1(n_v_0, edges_diamond) == dim_H1(n_v_1,
edges_1)}")
assert dim_H1(n_v_0, edges_diamond) == dim_H1(n_v_1, edges_1)

omega_box = np.array([+1, -1, +1, -1], dtype=float)
print(f"  Plaquette closed? {np.allclose(boundary(omega_box, edges_diamond,
n_v_0), 0)}")
assert np.allclose(boundary(omega_box, edges_diamond, n_v_0), 0)

Gamma = build_pullback(len(edges_diamond))
omega_box_lifted = lift_cycle(omega_box)
result = Gamma @ omega_box_lifted
print(f"  Γ*(L(ω_□)) = {result}")
print(f"  2·ω_□      = {2 * omega_box}")
print(f"  Γ* ∘ L = 2·Id? {np.allclose(result, 2 * omega_box)}")
assert np.allclose(result, 2 * omega_box)

# ---------------- Example: antichain glue ----------------
print()
print("=" * 50)
print("Example III: antichain glue")
print("=" * 50)
n_v_anti = 6
```

```
edges_anti = [(0,2),(0,3),(1,2),(1,3),(2,4),(3,4),(2,5),(3,5)]
print(f"  Λ_0: |V|={n_v_anti}, |E|={len(edges_anti)}, dim H¹ =
{dim_H1(n_v_anti, edges_anti)}")
assert dim_H1(n_v_anti, edges_anti) == 3

n_v_r, edges_r = midpoint_refine(n_v_anti, edges_anti)
print(f"  Λ_1: |V|={n_v_r}, |E|={len(edges_r)}, dim H¹ = {dim_H1(n_v_r,
edges_r)}")
assert dim_H1(n_v_r, edges_r) == 3

omega_a    = np.array([+1, -1, -1, +1, 0, 0, 0, 0], dtype=float)
omega_d    = np.array([0, 0, 0, 0, +1, -1, -1, +1], dtype=float)
omega_diag = np.array([+1, -1, 0, 0, +1, -1, 0, 0], dtype=float)

# Verify independence in H¹
d_anti = coboundary_matrix(n_v_anti, edges_anti)
augmented = np.column_stack([d_anti, omega_a, omega_d, omega_diag])
rank_d = np.linalg.matrix_rank(d_anti)
rank_aug = np.linalg.matrix_rank(augmented)
print(f"  {{[ω_a], [ω_d], [ω_Δ]}} independent in H¹? {rank_aug - rank_d ==
3}")
assert rank_aug - rank_d == 3

G = build_pullback(len(edges_anti))
for name, omega in [("ω_a", omega_a), ("ω_d", omega_d), ("ω_Δ", omega_diag)]:
    lifted = lift_cycle(omega)
    pulled = G @ lifted
    print(f"  Γ*(L({name})) = {pulled}  =  2·{name}: {np.allclose(pulled,
2*omega)}")
    assert np.allclose(pulled, 2*omega)

print()
print("All assertions passed.")
```

Running this code reproduces every numerical claim in the paper:

- dim H^1 is preserved at every refinement step;
- the lifted cycle basis is independent in H^1 on each substrate;
- $\Gamma^*(L(\omega)) = 2 \cdot \omega$ exactly on every cycle representative;
- $\Gamma^* \circ L = 2 \cdot \mathrm{Id}_{H^1}$ on every test substrate.