COEFFICIENT-DEPTH BLINDSPOTS IN TOPOLOGICAL MESSAGE PASSING

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning tasks can involve interactions among groups of entities. Graphs represent pairwise relations; simplicial and cellular complexes also represent larger groups. Topological message passing learns through local exchanges between their cells. Homology supplies global summaries that local computations can miss. Yet separate limits of these two channels do not characterize their combined expressivity. For every prime p, integer $s \geq 2$ and prescribed order $k \geq 1$, we construct finite pure two-dimensional full-incidence pairs with identical order-k observations at every finite depth. Their homology agrees over every field in every degree. For $A = \mathbb{Z}/p^{s-1}$, one chain isomorphism matches all permitted finite diagrams inside A-Mod, up to abstract isomorphism. Integral first homology and depth-s modules distinguish different cyclic lengths under the common torsion exponent p^s. Full cubical realization from a height-four lattice preserves the separation, adding a common free summand. Boundary computations and shared-weight sum/attention forwards check three source pairs and one dependent cube realization. An additional source pair extends the neural checks to matching mod-4 and distinct mod-8 information. An order-four simulation covers the standard scalable multicellular network interface. A depth-s homology readout separates each pair; fixed structural order, depth or width cannot uniformly replace this coefficient precision.

1 INTRODUCTION

Graphs represent pairwise relations; simplicial and cellular complexes add higher-dimensional cells, such as filled triangles, for group interactions. Face incidences record which cells are faces of others. Topological message passing updates cell features through face incidences and pools them into input representations (Bodnar et al., 2021b;a; Papillon et al., 2024). Expressivity determines which input differences, including global topology, these models distinguish.

Natural-neighborhood higher-order message passing (HOMP) can miss homological differences (Eitan et al., 2025). The torus-based pair has different Betti numbers, so homology separates it. Higher-order refinements strengthen structural observations and admit logical characterizations (Akbari et al., 2026); relational formulations connect topological message passing to graph-based analysis (Taha et al., 2025).

Persistence features enhance simplicial expressivity (Verma et al., 2024, Proposition 1); RePHINE combines vertex- and edge-level persistence (Immonen et al., 2023); hourglass persistence adds contractions and extends to simplicial and cellular networks (Ji et al., 2026). Homology depends on coefficient choice (Obayashi & Yoshiwaki, 2023; Ballester & Rieck, 2025): fixed precision can hide cyclic lengths in network sheaves (Ghrist & Ding, 2025). These separate limitations leave open whether both channels miss the same pair.

We compare two resources. Structural order k is the cell-tuple length. For a prime p, *coefficient depth* is the least $j \geq 1$ separating a pair's first homology over $\mathbb{Z}/p^j$. Integer boundaries $\mathrm{diag}(8, 8)$ and $\mathrm{diag}(8, 4)$ give first homology groups of exponent 8. With coefficients modulo 2 or 4, these groups are isomorphic; modulo 8, their orders are 64 and 32. Their coefficient depth is 3. The construction must realize this discrepancy while matching structural observations.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

The combined observer uses every nonempty face, its geometric dimension, full containment, and permitted colours and marks. Its coefficient channel supplies abstract all-field homology types and finite *coefficient diagrams*: modules with induced homology/cohomology maps in the module category A-Mod, where $A = \mathbb{Z}/p^{s-1}$ and $s \geq 2$ (Section 2). A *joint blindspot* is a pair with identical complete observations in both channels but different target labels. Arbitrary filtrations require further assumptions; Appendix I compares related torsion and persistence settings.

For each p, $s \geq 2$ and $k \geq 1$ fixed before construction, we produce one pair indistinguishable through both channels. The construction realizes the coefficient discrepancy while preserving structural indistinguishability and a shallow-ring chain isomorphism. Section 3.1 states the four obligations O1–O4. Our contributions are:

- *A joint lower bound for topology-augmented message passing.* One full-incidence pair matches the order-k observations at every finite depth, all-field homology and every permitted finite shallow coefficient diagram (Theorem 4, Proposition 8).
- *Complete geometric realizations of the blindspot.* Full cubical realizations from height-four lattices retain every cell and preserve the separation (Proposition 5, Corollary 6). Height four is the least lattice height for nontrivial first-homology torsion (Proposition 7).
- *A marked filling task with an exact precision threshold.* An always-answering binary decoder of the common observation has exactly $1/2$ accuracy on the uniform opposite-label pair (Section 4.3), even with the prescribed cycle's nonzero shallow homology class (Appendix H.7). A standard depth-s readout separates the family (Section 4.4).

The construction combines affine counting indistinguishability (Atserias et al., 2007), Smith normal form and universal coefficients (Hatcher, 2002), precision locality (Ghrist & Ding, 2025), and cubical subdivision (Athanasiadis, 2012; Krishnan & Rudman, 2025; Kahle et al., 2021). We couple these tools on one fully realized pair, preserving complete face inventories and the prescribed cycle. Table 3 gives the detailed prior-work comparison. Section 5 tests a concrete summary-fusion pipeline. Each pair depends on the prescribed k.

2 Observation Model

To analyze the combined model, we specify the information available to its encoder and decoder. The structural and coefficient channels are defined separately below. For a finite 2-complex X, the universe of $I(X)$ consists of all nonempty faces of dimension $0, 1, 2$. The notation $I(X)$ also denotes this set. Every face x is labelled by its geometric dimension $\mathrm{rank}(x)$. The structure records equality and strict geometric containment Sub between faces, including vertex-to-triangle containment. For simplicial and cubical inputs, Face denotes the same containment. Unmarked inputs carry only these incidence relations and shared rank-dependent colours.

The local refinement propagates cell features through four channels derived from Face: faces, cofaces, and same-dimensional cells sharing a face or coface. Write $B(x, y)$ iff $\mathrm{Face}(y, x)$ and $\mathrm{rank}(y) = \mathrm{rank}(x) - 1$, and let C reverse B. The relation $\mathrm{Down}(x, y)$ holds iff $\mathrm{rank}(x) = \mathrm{rank}(y)$ and some z satisfies $B(x, z) \wedge B(y, z)$. Dually, Up uses a common C-mediator. Both relations permit $y = x$. Every distinct mediator occurrence and ordered pair is retained, including equal values. Atomic types use rank, the stated finite colours, equality and all four relations.

The local and paired profiles below use the combinatorial complex Weisfeiler–Leman (CCWL) recurrences of Akbari et al. (2026, Section 3 and Appendix B). Their whole-input readouts use unaugmented inputs; Appendix D treats broadcast augmentation separately. For initialization, write $\mathrm{col}(x)$ for the shared rank-dependent colour of x, extended by the prescribed marks when present.

Definition 1 (Local profile, $k = 1$)**.** Set $\chi_0(x) = (\mathrm{rank}(x), \mathrm{col}(x))$. Each round hashes the old colour together with four ordered channels: the multiset of B-neighbour colours, the multiset of C-neighbour colours, the multiset of pairs $(\chi_t(y), \chi_t(z))$ over Down-neighbours y with common B-mediator z, and the analogous Up multiset with common C-mediator. The readout compares all cells.

Write $u[\alpha/i]$ for the tuple obtained by replacing entry i of u with α.

Definition 2 (Paired profile, $k \geq 2$)**.** Let $\chi_0(u)$ be the full ordered atomic type of the k-tuple u. For every $(\alpha, \beta) \in I(X)^2$ retain the context consisting of the atomic type $\mathrm{atp}_{k+2}(u, \alpha, \beta)$ and the ordered vector $[(\chi_t(u[\alpha/i]), \chi_t(u[\beta/i]))]_{i=1}^{k}$. The update hashes the old colour together with the whole multiset of these contexts. The readout compares all ordered k-tuples, including tuples with repeats.

Whole-input colour histograms form the decoder's structural observation. A shared partition is stable when further rounds split no colour class. The observation $\mathsf{Obs}_k(X)$ records every finite-round histogram and determines the stable profile. Fix k and an encoder whose layer states are functions of the corresponding refinement colours. Equal histograms and complete side inputs then give equal invariant readouts at every finite depth and width when both inputs use the same parameters.

2.1 Natural-neighborhood higher-order message passing (HOMP)

For natural HOMP, fix finitely many rank predicates R_r: $R_r(x)$ means that cell x has dimension r. Any mark predicates must also be covered by the counting-equivalence premise. Let $F = \mathrm{Face}$ denote strict full face containment, and write $x \preceq y$ for $x = y \vee F(x, y)$. The natural neighborhoods of Eitan et al. (2025) form the following rank-pair families:

$$\begin{aligned} B_{rt}(x, y) &= R_r(x) \wedge R_t(y) \wedge x \preceq y, \\ B^{\mathsf{T}}_{rt}(x, y) &= R_r(x) \wedge R_t(y) \wedge y \preceq x, \\ A_{rt}(x, y) &= R_r(x) \wedge R_r(y) \wedge \exists z \big(R_t(z) \wedge x \preceq z \wedge y \preceq z\big), \\ \mathrm{co}\, A_{rt}(x, y) &= R_r(x) \wedge R_r(y) \wedge \exists z \big(R_t(z) \wedge z \preceq x \wedge z \preceq y\big). \end{aligned}$$

Each neighbouring cell is counted once, however many mediators z witness the adjacency. Definitions 1 and 2 keep mediator multiplicity. Self-neighbours are kept wherever the predicates permit.

Proposition 3 (Natural-neighborhood HOMP coverage)**.** *Let X and Y be finite inputs equivalent in ordinary three-variable counting logic C^3 over this vocabulary. For any fixed finite collection of the displayed neighborhoods, every finite round of rank-aware neighbour-colour refinement has equal whole-input and rank-selected colour multiplicities on X and Y. Consequently HOMP has equal mathematical outputs and invariant readout under these conditions: initial vectors depend only on the stated colours; parameters for each layer, rank and channel are shared between the inputs; neighbour aggregation and the final multiset readout are invariant; and any further decoder inputs agree.*

An explicit neural interface. To express the covered encoders as neural maps, fix a finite list $N_1, \ldots, N_m$ of the displayed neighborhoods. Let $r(x)$ denote geometric rank and $\mu(x)$ the permitted finite mark vector. For input K, the encoder uses T layers, and the decoder also receives the coefficient side input $z(K)$ specified below. The maps take the form:

$$\begin{aligned} h_x^0 &= E_{r(x)}(\mu(x)), \\ m_{\ell j}(x) &= \mathrm{AGG}_{\ell j}\big\{\!\big\{\phi_{\ell j}(h_x^\ell, h_y^\ell) : y \in N_j(x)\big\}\!\big\}, \\ h_x^{\ell+1} &= U_{\ell, r(x)}\big(h_x^\ell, (m_{\ell j}(x))_{j=1}^m\big), \\ f_\theta(K) &= D_\theta\big(R_\theta\big(\big(\{\!\{h_x^T : R_r(x)\}\!\}\big)_{r=0}^2\big), z(K)\big). \end{aligned}$$

The same maps and parameters are used for both compared inputs. Aggregation and pooling are invariant, with a common empty-multiset convention. Finite depth and width are arbitrary. Local softmax attention is also admissible when its scores use only the permitted center and neighbor states (Appendix G).

A concrete summary-fusion model. Take R_θ to concatenate rank-wise mean vectors, with zero for empty ranks. Define the scalar homology summaries:

$$B_F(K) = (\dim_F H_i(K; F))_{i=0}^2, \qquad \boldsymbol{\ell}_j(K) = (\log_p |H_i(K; \mathbb{Z}/p^j)|)_{i=0}^2.$$

Concatenate $z_j(K) = (B_{\mathbb{Q}}(K), B_{\mathbb{F}_p}(K), \boldsymbol{\ell}_j(K))$ before the shared decoder. Boundary matrices compute these summaries but no cellular basis is passed to the model. For $j < s$, both branches agree on the constructed pair. Replacing $\boldsymbol{\ell}_j$ by $\boldsymbol{\ell}_s$ makes the separating readout of Section 4.4

available. This is a concrete information boundary for global-summary fusion; it does not assert coverage of filtration-based or cell-aligned augmentation.

The construction gives C^{k+3} equivalence, hence C^3, including the prescribed marks. Appendix G.1 specifies the scalable multicellular network (SMCN) interface of Eitan et al. (2025) and simulates it by the order-four paired profile. For simultaneous order-k and SMCN coverage, construct at $K = \max\{k, 4\}$ with $\mathrm{tw}(H) > 6(K+3)$.

Observation contract. The structural input contains the specified incidence relations and labels, the six predicates for marked tasks, and the two common roles when the specified broadcast augmentation is used. The coefficient input uses a common encoding of abstract field-homology types and shared finite diagrams from Proposition 8. All modules and short exact sequences lie in A-Mod for $A = \mathbb{Z}/p^{s-1}$; any finite collection is permitted, but not every module simultaneously. The side input $z(K)$ obeys this contract; additional side inputs must agree across the pair. Arbitrary identifiers, orderings, coordinates, signed incidence matrices and cell-basis alignment are excluded. Dimension-skeletal diagrams are covered by Corollary 123; general feature filtrations, cohomology products and additional signed or geometric spectral inputs require further premises. Bounds concern forward computations with shared parameters fixed before the hidden test draw.

Section 3.1 states the compatibility conditions for matching these observations on one pair.

3 Joint blindspots and the separation theorem

3.1 Compatibility obligations

The presentation fixes homology while the face interpretation preserves structural observations. Private cells retain entire ordered relator tuples, so shared coordinates never merge faces from different relation instances. The lattice guards control every Boolean interval of the full cube complex. These requirements are enforced on the same pair:

	Requirement enforced on the same pair	Proof
O1	The two twists have equal complete face counts and support both structural and coefficient comparisons.	A
O2	Arity-three source and arity-two cube encodings preserve all cells, the six marks and whole-input histograms.	A, C–D, F
O3	Equal chain ranks and isomorphic free shallow homology give a chain isomorphism. The prescribed unscaled cycle classes generate free direct summands, allowing a cycle-preserving choice.	H, H.7
O4	Two guard families prevent unwanted Boolean cells; completion adds common free circles and preserves torsion.	B

3.2 The separation theorem

To construct such a pair, fix a prime p and integers $s \geq 2$, $k \geq 1$. Put $q = p^s$, $a = p^{s-1}$ and $r_s = 2(k+3)$, the counting budget. The budget is independent of s. The construction starts from an affine system built on a connected simple cubic graph H with $n = |V(H)|$ and $m = 3n/2$. Require $b = m - n + 1 \geq 1$ and $\mathrm{tw}(H) > 3r_s = 6(k+3)$. The generator and relator counts are $g = 1 + 3qn$ and $M = g + nq^3 + mq^2$; set $h_2 = M - g$.

The first three items give shared observations; the fourth identifies the torsion difference that becomes visible at higher coefficient precision.

Theorem 4 (Coefficient-depth blindspot)**.** *There is a pair (X, Y) of finite pure 2-dimensional complexes such that:*

1. $\mathsf{Obs}_k(X) = \mathsf{Obs}_k(Y)$ *at every finite depth, for the profile of Definition 1 when $k = 1$ and of Definition 2 when $k \geq 2$;*

2. *for every field $\mathbb{F}$ of characteristic p, $H_0 \cong \mathbb{F}$, $H_1 \cong \mathbb{F}^{b+1}$, $H_2 \cong \mathbb{F}^{h_2+b+1}$ for both, and for every field of any other characteristic, including 0, $H_1 = 0$ and $H_2 \cong \mathbb{F}^{h_2}$ for both; $H_0 \cong \mathbb{F}$ and $H_i = 0$ for $i > 2$ in both cases;*

3. *for $A_j = \mathbb{Z}/p^j$ with $1 \le j < s$, both have $H_0 \cong A_j$, $H_1 \cong A_j^{b+1}$, $H_2 \cong A_j^{h_2+b+1}$, and zero above;*

4. *nevertheless $H_1(X;\mathbb{Z}) \cong (\mathbb{Z}/p^s)^{b+1}$ and $H_1(Y;\mathbb{Z}) \cong (\mathbb{Z}/p^s)^b \oplus \mathbb{Z}/p^{s-1}$, and at depth s, as A_s-modules, $H_1(X;A_s) \cong A_s^{b+1}$, $H_2(X;A_s) \cong A_s^{h_2+b+1}$, while $H_1(Y;A_s) \cong A_s^b \oplus \mathbb{Z}/p^{s-1}$ and $H_2(Y;A_s) \cong A_s^{h_2+b} \oplus \mathbb{Z}/p^{s-1}$.*

Both integral torsion subgroups have total exponent p^s; the orders $|\operatorname{Tor} H_1|$ differ by exactly one factor of p. Moreover, for any commutative ring A, $H_1(X;A) \cong (A/p^sA)^{b+1}$ and $H_2(X;A) \cong A^{h_2} \oplus A[p^s]^{b+1}$, whereas $H_1(Y;A) \cong (A/p^sA)^b \oplus A/p^{s-1}A$ and $H_2(Y;A) \cong A^{h_2} \oplus A[p^s]^b \oplus A[p^{s-1}]$, with $A[d] = \{x : dx = 0\}$, $H_0 \cong A$, and $H_i = 0$ for $i > 2$.

Items 2–4 give coefficient depth exactly s. Items 1–3 meet O1–O2. The fixed chain isomorphism gives O3 (Proposition 8); the guarded completion preserves the separating cycle for O4 (Section 4.1).

Spanning-tree elimination and Smith normal form are performed over the integers. Tensoring the resulting chain complex then gives the displayed coefficient homology types. Proposition 8 extends shallow agreement to coherent diagrams inside A-Mod.

The source pair is $X = S(E^0)$, $Y = S(E^a)$, where the affine systems differ by a twist $\epsilon \in \{0, a\}$ at one distinguished base vertex. Each source encodes all ordered coloured equations over $\mathbb{Z}/q$. The relation for $x_1 + \cdots + x_j = c$ is $z_{x_1} + \cdots + z_{x_j} - ct$; together with qt, qz_x, these present the abelian group $P(E^\epsilon)$. After the integral shear $z_{(v,e,i)} = y_{(v,e,i)} + i\,t$, unit and tree elimination leave b chord generators with relations qz and the distinguished generator t with relations $qt, \epsilon t$. The resulting group is:

$$P(E^\epsilon) \cong (\mathbb{Z}/q)^b \oplus \mathbb{Z}/\gcd(q, \epsilon).$$

The two twists give the groups in Theorem 4, without division in $\mathbb{Z}/q$. The positive word $x_1 \cdots x_j\, t^{(-c) \bmod q}$ realizes each relation geometrically. Its private vertices keep repeated letters distinct where needed; equal face counts and the logical interpretation hold for this same realization (Appendix A).

Let f_i count source i-cells. Their exact census in Appendix A gives at most $1+5g+(24q+49)M = O(nq^4)$ cells. Explicit cubic bases yield $O((K+3)^2p^{4s})$ cells for the source and full-cube witnesses, with $K = k$ or $\max\{k, 4\}$ for simultaneous SMCN coverage (Proposition 127).

4 Construction and Proof Overview

4.1 Height-four realization

The lattice lift must preserve the source's separating torsion. For a finite lattice L, $\mathcal{Q}(L)$ has one cell per Boolean interval of L, an interval isomorphic to a finite subset lattice. Lattice height counts strict chain steps; cell rank is geometric dimension. For connected pure two-dimensional X, the two-guard lift $L_2(X)$ has height four and adds only a free homology summand (Appendix B).

Proposition 5 (Full height-four realization)**.** *$\mathcal{Q}(L_2(X)) \simeq |X| \vee \bigvee_R S^1$ with $R = f_0+3f_1+f_2-2$. Every lattice element remains a cubical 0-cell, and the cell counts are $N = c_0 = f_0+4f_1+f_2+2$, $c_1 = f_0+8f_1+4f_2$, $c_2 = 3f_2$. Only the corner intervals $[v, t_T]$ are squares, and no Boolean interval of rank ≥ 3 occurs; the full Boolean-interval classification is in Appendix B, and the exact* Face *table in Appendix C. The corner-square core Q is homeomorphic to X. The extra vertices number $2+3f_1$ and the extra edges $f_0+6f_1+f_2$; contracting the two explicit rooted trees leaves exactly R free circles.*

Guard intervals. The guards control which lattice intervals become cubical cells. For a source edge $\eta = \{u, v\}$, Appendix B identifies three intermediates in $[\hat{0}, E_\eta]$: the endpoints and lower guard λ_η. A Boolean square interval has exactly two intermediate elements, so this interval cannot form a square. If d source triangles contain η, $[E_\eta, \hat{1}]$ contains those triangles and the two upper guards

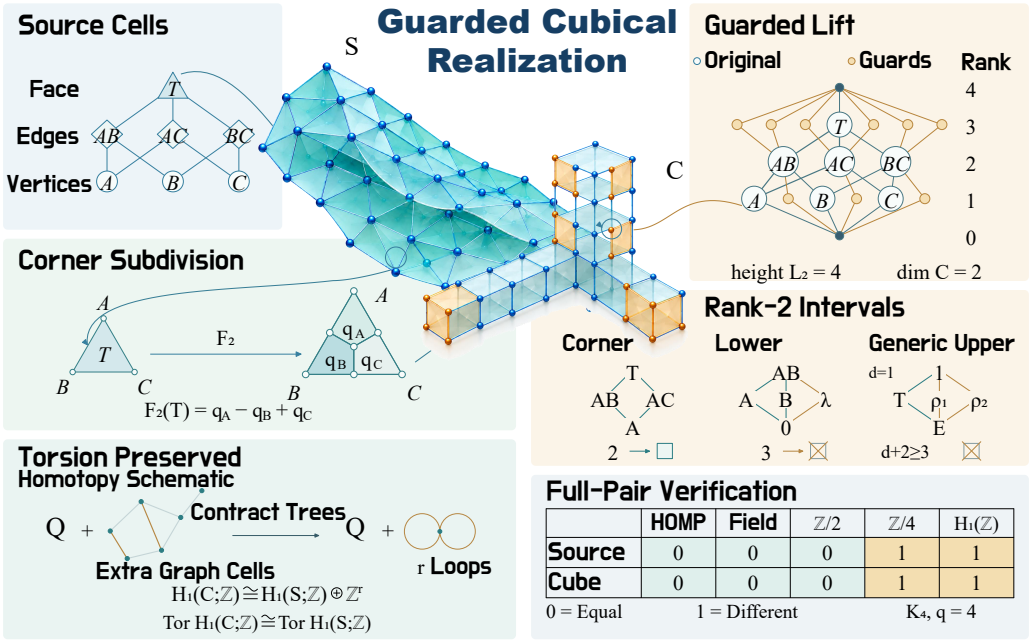

Figure 1: **Guarded cubical realization and finite verification.** A source triangle T subdivides into three corner squares in the corner-square core $Q \subset C = \mathcal{Q}(L_2(S))$, where C is the full complex. Only corner lattice-rank-two intervals form squares. The upper sketch shows the generic $d = 1$ case, with d counting incident triangles. The selected edge's full interval is not drawn. Construction colours and lattice ranks annotate the lift. The figure's oriented bases satisfy $q_A = Q_0$, $q_B = -Q_1$, $q_C = Q_2$ in Eq. (20) of Appendix F. Contracting the attached trees leaves $R = f_0(S) + 3f_1(S) + f_2(S) - 2$ free circles, labelled r in the figure, using all source face counts $f_i(S)$ and preserving first-homology torsion. The matrix compares one $K_4, q = 4$ pair in its source and dependent cubical realizations. An entry 0 means equal observations and 1 means different observations. Entries are not accuracies. Natural HOMP starts from geometric cell ranks only and uses all 36 channels. Field denotes all fields. Large bodies and free-loop views are schematic. Local incidences and matrix values come from exact calculations.

ρ^1_η, ρ^2_η. Purity gives $d \geq 1$, so the upper interval also has at least three intermediates. Appendix B excludes all other square types and higher Boolean cells. Thus only $[v, t_T]$ for $v \in T$, with its two edge intermediates, is a corner square.

Full cubes retain every guard and incident edge, have maximum dimension 2 and may be non-pure. Figure 1 summarizes the realization and finite comparison. For any commutative ring A, $H_1(\mathcal{Q}(L_2(X)); A) \cong A^R \oplus H_1(X; A)$, while H_i is unchanged for $i \geq 2$. Equal source face vectors give the same R for Y.

Corollary 6 (Joint separation on full height-four lattice inputs)**.** *The full inputs $\widehat{X} = \mathcal{Q}(L_2(X))$ and $\widehat{Y} = \mathcal{Q}(L_2(Y))$ satisfy $\mathsf{Obs}_k(\widehat{X}) = \mathsf{Obs}_k(\widehat{Y})$, have isomorphic homology in every degree over every field and over each $\mathbb{Z}/p^j$ for $1 \leq j < s$, and have integral first homology:*

$$H_1(\widehat{X}; \mathbb{Z}) \cong \mathbb{Z}^R \oplus (\mathbb{Z}/p^s)^{b+1},$$
$$H_1(\widehat{Y}; \mathbb{Z}) \cong \mathbb{Z}^R \oplus (\mathbb{Z}/p^s)^{b} \oplus \mathbb{Z}/p^{s-1}.$$

Their depth-s modules differ as well. The torsion exponents remain equal to p^s.

Proposition 7 (Sharp height)**.** *If a finite lattice L has height at most 3, then $H_1(\mathcal{Q}(L); \mathbb{Z})$ is free abelian. For every nontrivial finite abelian group G, some finite lattice L of height 4 has* Tor $H_1(\mathcal{Q}(L); \mathbb{Z}) \cong G$*, and 4 is the least such height (Appendices B.7–B.8).*

Roadmap. Logical encodings transfer counting equivalence to source and cubical complexes, then to whole-input readouts. Affine systems are $\mathsf{C}^{6(k+3)}$-equivalent when $\mathrm{tw}(H) > 6(k+3)$ (Atserias

et al., 2007). The source encoding costs $3r_s$ variables at $r_s = 2(k+3)$ (Appendix A); the full-cube encoding uses $2(k+3)$, yielding C^{k+3} equivalence (Appendix C). Appendix D gives whole-input histogram equality, including repeated tuples and rank/equality-pattern selection. The one-way transfers are $\mathrm{TC}^3 \to \mathsf{C}^4$ and $\mathrm{TC}^{k+2} \to \mathsf{C}^{k+3}$; pair-counting variables have distinct names but may repeat values. Section 4.2 gives coherent shallow-diagram agreement on the same pair.

4.2 Coherent shallow coefficient information

Matching homology modules separately is insufficient when the coefficient channel also sees maps between them. A fixed chain isomorphism over the shallow coefficient ring matches the entire diagram, on both source and full-cube pairs.

Proposition 8 (Coherent shallow coefficient diagrams)**.** *Let $A = \mathbb{Z}/p^{s-1}$, and let (X, Y) be either the source pair of Theorem 4 or the full height-four cubes $(\widehat{X}, \widehat{Y})$ of Corollary 6. Write ρ for the common free first-homology rank ($\rho = 0$ for the sources, $\rho = R$ for the cubes), h for the common free second-homology rank, and b for the base cycle rank, so that, with coefficients in A, both sides have $H_0 \cong A$, $H_1 \cong A^{\rho+b+1}$ and $H_2 \cong A^{h+b+1}$.* (i) *The ordinary unreduced cellular A-chain complexes of the two sides are isomorphic as A-chain complexes.* (ii) *One fixed pair of mutually inverse chain isomorphisms induces, for every A-module M and every A-linear map, natural isomorphisms $H_i(X; M) \cong H_i(Y; M)$ and $H^i(X; M) \cong H^i(Y; M)$; and for every short exact sequence of A-modules the associated connecting homomorphisms vanish on both sides.* (iii) *Every finite diagram assembled from A-modules, A-linear coefficient maps, short exact sequences in A-Mod, and the induced homology and cohomology arrows has the same abstract isomorphism type on X and on Y.*

Equal chain ranks and free A-homology ranks give identical split standard complexes, hence a fixed chain isomorphism (Appendix H). It identifies coefficient diagrams naturally, including nonflat modules and nonsplit sequences. The prescribed broadcast-star augmentation and dimension-skeletal persistence preserve the pointed conclusion (Proposition 122, Corollary 123). Appendix H.11 extends coefficient coherence to mixed groups below the same p-primary threshold.

The A-Mod restriction excludes coefficient sequences revealing higher precision. At $s = 2$, where $A = \mathbb{F}_p$, the first Bockstein connecting homomorphism separates the pair using $0 \to \mathbb{F}_p \to \mathbb{Z}/p^2 \to \mathbb{F}_p \to 0$. This sequence lies outside A-Mod: multiplication by p is nonzero on its middle term. Precision locality of connecting information comes from prior work (Ghrist & Ding, 2025, Lemma 5.9). Section 4.3 extends agreement to the prescribed nonzero shallow homology class and its naturally induced maps.

4.3 Marked filling and certificates

We express the separation as a binary filling task: determine whether a prescribed cycle bounds an integral 2-chain. Recall $q = p^s$ and $a = p^{s-1}$. In Appendix F, γ_t is the oriented three-edge circle in the auxiliary presentation graph, through vertices o, A_t and B_t. The marked query $z = a\gamma_t$ fills over $\mathbb{Z}$ and $\mathbb{Z}/q$ only on the twisted side. At every shallower depth and in characteristic p, the query is literally zero. In other fields, the chain $(a/q)D_{t^q}$, obtained from the power disk D_{t^q} of Appendix F, fills the query. The chain-level facts give a certificate-bearing arithmetic witness on the marked pair.

Appendix F proves that the marked source pair is C^{2r_m}-equivalent and the marked full-cube pair is C^{r_m}-equivalent when $r_m \geq 2$ and $\mathrm{tw}(H) > 6r_m$. Source marks name o, A_t, B_t and their three consecutive edges through genuine zero-arity tags. The cube marks name six 0-cells on the subdivided circle, alternately corresponding to vertex and edge elements. The claims cover these prescribed marks, with arbitrary marking and unmarked definability outside scope. At $r_m = k + 3$, the treewidth threshold is unchanged.

The filling labels come with certificates defined on the actual cells. Mutually inverse maps identify integer presentation solutions with actual fillings for cycles in the presentation graph. Explicit modular cocycles obstruct negative instances, and subdivision preserves the certificates in full cubes. Signed coordinates are reserved for verification and excluded from the encoding. The decoder bound concerns the common marked observation and equal side inputs. When the two inputs are sampled with equal probability, accuracy is exactly $1/2$ for an always-answering decoder. With abstention or a required valid certificate, success is at most $1/2$.

Table 1: Fixed marked tasks at $q = 4$. Cell counts are per input. Stable rounds refer to marked natural-HOMP refinement. The gap is the largest pooled-vector difference across the 64 sum/attention settings per pair.

Input	Cells	Stable round	Max. pooled gap
K_4 source	41,063	10	2.85×10^{-14}
$K_{3,3}$ source	61,543	11	2.23×10^{-15}
Prism source	61,543	14	1.10×10^{-14}
Full K_4 cube	376,614	13	9.19×10^{-15}

The unscaled class as coefficient side input. We strengthen the coefficient input by supplying the unscaled cycle's nonzero shallow class, with its inclusion and evaluation maps. The query $a\gamma_t$ remains zero at every shallower depth on both sides. For six-marked sources and full two-guard cubes, the chain isomorphism in Appendix H.7 matches the prescribed unscaled t-cycles. Within A-Mod for $A = \mathbb{Z}/p^{s-1}$, the permitted finite diagrams remain abstractly isomorphic after including this class, its two split maps and natural coefficient maps. Thus the joint blindspot persists under Section 2's contract. Section 5 checks the finite marked observations and filling certificates.

4.4 A SEPARATING READOUT

After establishing the prediction limit, we give a separating readout with access to depth-s coefficients. Assign labels 0 and 1 to the lifted pair $\widehat{X}$ and $\widehat{Y}$. Appendices F and E extend the Section 4.3 bound to this pair. When the two inputs are sampled with equal probability, the equal joint observation limits accuracy to chance.

The standard depth-s homology readout (Hatcher, 2002) lies outside the permitted model profile (Appendices D and E) and is untrained. Let $\ell_s(K) = \log_p |H_1(K; \mathbb{Z}/p^s)|$. Then $\ell_s(\widehat{X}) = s(R + b + 1)$ and $\ell_s(\widehat{Y}) = s(R + b) + s - 1$. For the fixed p, s constructed family, predicting label 1 iff $\ell_s \not\equiv 0 \pmod{s}$ is uniformly correct.

For fixed p and s, every prescribed finite order k admits a pair with equal permitted observations at all finite depths and widths. Each pair depends on k: for k at least the number of cells, equal round-zero histograms in Definition 2 force an incidence isomorphism. Thus, no fixed structural order removes all joint blindspots in this family at shallow coefficient precision. Proposition 127 gives an explicit sufficient witness-size bound.

5 COMPUTATIONAL CHECKS

Constructed inputs. With $p = s = 2$ and $q = 4$, we evaluate the three source pairs and dependent K_4 cube pair in Table 1. Each input retains all cells and six marks. Base graphs have at most six vertices and treewidth at most five, below the general theorem premise. Signed-boundary reduction, trace replay and measured Smith factors yield the coefficient summaries. Independent integer fillings and mod-4 cocycles certify labels; labels and twists never enter features.

Fixed encoder, different coefficient information. Sum HOMP and admissible local attention use all 36 natural channels and rank-wise mean pooling. Widths $8, 16$, seeds $0, 1, 2, 3$ and depths $0, 1, 2, 4$ give 256 paired settings across four realizations. A shared random affine head combines pooled features with zeros, field summaries, or field summaries augmented by $\boldsymbol{\ell}_1$ or $\boldsymbol{\ell}_2$. No weights are trained. Count-weighted equitable quotients implement the same full-cell recurrences.

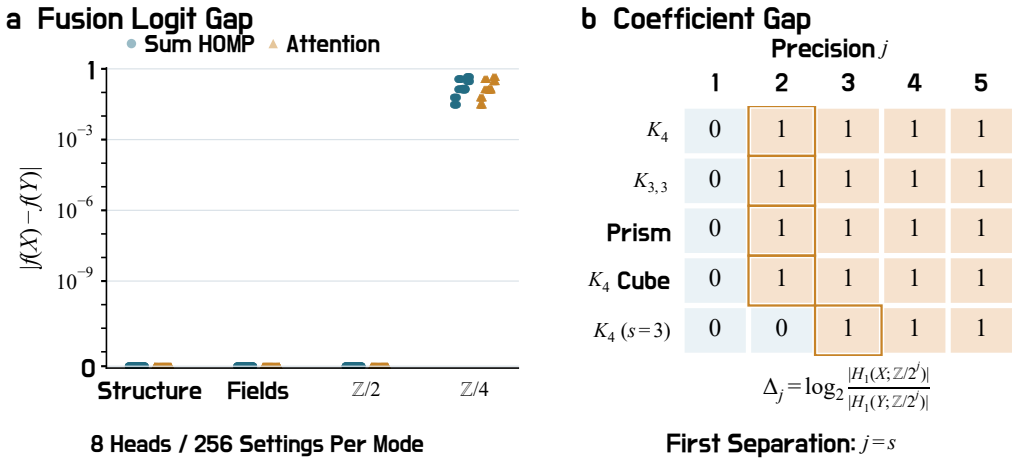

Figure 2: **Finite coefficient-precision checks.** (a) Paired-logit gaps for the four $s = 2$ realizations: 256 settings per mode share eight heads. The axis is logarithmic outside a linear neighborhood of zero; horizontal offsets separate overlapping records without changing values. (b) Measured coefficient-order gaps for those pairs and the additional $s = 3$ source pair. Boxes mark each first separating precision. The cube depends on the K_4 source; cells and settings are not independent task samples.

Results. All pooled vectors and the first three fusion modes (Figure 2a) match within the specified tolerance (10^{-9} absolute, 10^{-10} relative); the largest shallow-fusion logit gap is 2.85×10^{-14}. The ℓ_2 mode gives distinct logits in all 256 settings, with gaps 0.0302–0.4488. Pooled vectors match, and each pair's summaries at its coefficient depth differ by one fixed vector, so the gap depends only on the head up to roundoff. The fixed arithmetic readout at coefficient depth two in Section 4.4 correctly labels all eight physical inputs.

Same-head controls isolate the summary contribution: removing the neural branch changes paired gaps by at most 7.28×10^{-12}. Implementation checks, replay and twelve diagonal-chain controls appear in Appendix I.3.

A non-field shallow ring. At $p = 2, s = 3$, an additional K_4 source pair has 391,815 cells per input and equal six-mark refinement stable at round 12. The same encoders and eight heads give 64 paired settings: structure, fields, mod-2 and mod-4 modes match, while mod-8 separates all 64 with the same eight gaps up to roundoff. The exact readout matches both filling labels. Mod-4 cocycles pair to one with each unscaled marked cycle (Appendix I.3).

Interpretation. Across both tested coefficient thresholds, the fixed encoders share shallow observations and the coefficient branch supplies the separating information.

6 Conclusion

For every prime p, depth $s \geq 2$ and structural order fixed before construction, our full-incidence pairs match permitted message-passing, all-field homology and finite shallow-diagram observations, also after the height-four lattice lift. A standard depth-s homology readout separates each pair (Section 4.4). Boundary and partition certificates verify the tested finite instances; neural evaluations check the covered interface.

Coefficient precision is an expressivity resource beyond structural capacity. The extensions cover specified broadcast nodes, skeletal diagrams and standard SMCN, with explicit witness-size bounds. Rounding tolerates scalar-summary errors strictly below $1/2$ (Remark 125). Measuring the prevalence of these gaps and their training consequences in natural complexes is an empirical direction.

AI USE STATEMENT

The authors originated and directed the core research ideas. AI tools assisted literature research, refinement of mathematical formulations and claims, proof development and writing, experimental design, synthetic-data generation, implementation and analysis, and manuscript preparation, including translation and figures. The authors are responsible for verifying the mathematical claims and computational results and for the final manuscript.

REPRODUCIBILITY STATEMENT

Section 2 specifies the observation model and neural interface, and Section 5 describes the physical inputs, encoder configurations, coefficient summaries and numerical tolerances. Appendices A–H provide complete proofs and construction details, including the signed filling certificates in Appendix F. Appendix I specifies the forward recurrences, quotient implementation, sensitivity controls and replay protocol. The supplementary material provides the code, input inventories and raw records needed to reproduce the reported computations.

REFERENCES

Ralph Abboud, İsmail İlkan Ceylan, Martin Grohe, and Thomas Lukasiewicz. The surprising power of graph neural networks with random node initialization. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 2112–2118, 2021. doi: 10.24963/ijcai.2021/291.

Amirreza Akbari, Amauri H. Souza, and Vikas Garg. The Logical Expressiveness of Topological Neural Networks. In *International Conference on Learning Representations (ICLR)*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/22650d2b38807c88a7226eb1fed31e43-Paper-Conference.pdf.

Christos A. Athanasiadis. Cubical subdivisions and local h-vectors. *Annals of Combinatorics*, 16 (3):421–448, 2012. doi: 10.1007/s00026-012-0140-y. URL https://arxiv.org/pdf/1007.3154v3. Example 7.12 locators refer to arXiv:1007.3154v3, 30 January 2011.

Albert Atserias, Andrei Bulatov, and Anuj Dawar. Affine Systems of Equations and Counting Infinitary Logic. In *Automata, Languages and Programming*, volume 4596 of *Lecture Notes in Computer Science*, pp. 558–570. Springer, 2007. doi: 10.1007/978-3-540-73420-8_49. URL https://www.cl.cam.ac.uk/~ad260/papers/abd-icalp07.pdf. ICALP 2007. Theorem locators refer to the 12-page author version.

Muhammet Balcilar, Pierre Héroux, Benoit Gaüzère, Pascal Vasseur, Sébastien Adam, and Paul Honeine. Breaking the limits of message passing graph neural networks. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 599–608. PMLR, 2021. URL https://proceedings.mlr.press/v139/balcilar21a.html.

Rubén Ballester and Bastian Rieck. On the Expressivity of Persistent Homology in Graph Learning. In *Proceedings of the Third Learning on Graphs Conference*, volume 269 of *Proceedings of Machine Learning Research*, pp. 42:1–42:31. PMLR, 2025. URL https://proceedings.mlr.press/v269/ballester25a.html. LoG 2024; final proceedings published in 2025.

Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yu Guang Wang, Pietro Liò, Guido Montúfar, and Michael Bronstein. Weisfeiler and Lehman go cellular: CW networks. In *Advances in Neural Information Processing Systems*, volume 34, pp. 2625–2640. Curran Associates, Inc., 2021a. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/157792e4abb490f99dbd738483e0d2d4-Paper.pdf.

Cristian Bodnar, Fabrizio Frasca, Yu Guang Wang, Nina Otter, Guido Montúfar, Pietro Liò, and Michael M. Bronstein. Weisfeiler and Lehman go topological: Message passing simplicial networks. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 1026–1037. PMLR, 2021b. URL https://proceedings.mlr.press/v139/bodnar21a.html.

Cristian Bodnar, Francesco Di Giovanni, Benjamin Paul Chamberlain, Pietro Liò, and Michael M. Bronstein. Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs. In *Advances in Neural Information Processing Systems*, volume 35. Curran Associates, Inc., 2022. URL https://papers.nips.cc/paper_files/paper/2022/hash/75c45fca2aa416ada062b26cc4fb7641-Abstract-Conference.html.

Anuj Dawar, Erich Grädel, and Moritz Lichter. Limitations of the Invertible-Map Equivalences. *Journal of Logic and Computation*, 33(5):961–969, 2023. doi: 10.1093/logcom/exac058. URL https://arxiv.org/pdf/2109.07218v2. Theorem locators use author version 2 of 26 September 2022.

Stefania Ebli, Michaël Defferrard, and Gard Spreemann. Simplicial neural networks. *arXiv preprint arXiv:2010.03633*, 2020. URL https://arxiv.org/abs/2010.03633.

Yam Eitan, Yoav Gelberg, Guy Bar-Shalom, Fabrizio Frasca, Michael Bronstein, and Haggai Maron. Topological Blindspots: Understanding and Extending Topological Deep Learning Through the Lens of Expressivity. In *International Conference on Learning Representations (ICLR)*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/ca3b1f24fc0238edf5ed1ad226b9d655-Paper-Conference.pdf.

Robert Ghrist and Cassie Ding. Precision-Graded Cohomology and Arithmetic Persistence for Network Sheaves. arXiv:2511.00677v1, 2025. URL https://arxiv.org/html/2511.00677v1. Version 1, 1 November 2025; preprint.

Lorenzo Giusti, Claudio Battiloro, Lucia Testa, Paolo Di Lorenzo, Stefania Sardellitti, and Sergio Barbarossa. Cell Attention Networks. In *International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2023. doi: 10.1109/IJCNN54540.2023.10191530.

Martin Grohe and Wied Pakusa. Descriptive complexity of linear equation systems and applications to propositional proof complexity. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pp. 1–12, 2017. doi: 10.1109/LICS.2017.8005081. URL https://www.logic.rwth-aachen.de/pub/pakusa/GrohePak17.pdf.

Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. URL https://pi.math.cornell.edu/~hatcher/AT/AT.pdf.

Max Horn, Edward De Brouwer, Michael Moor, Yves Moreau, Bastian Rieck, and Karsten Borgwardt. Topological Graph Neural Networks. In *International Conference on Learning Representations (ICLR)*, 2022. URL https://openreview.net/forum?id=oxxUMeFwEHd.

Johanna Immonen, Amauri H. Souza, and Vikas Garg. Going Beyond Persistent Homology Using Persistent Homology. In *Advances in Neural Information Processing Systems*, volume 36, 2023. URL https://papers.nips.cc/paper_files/paper/2023/hash/c78f81a878a72566422f37279bca0fd0-Abstract-Conference.html.

Mattie Ji, Indradyumna Roy, and Vikas Garg. Contraction and Hourglass Persistence for Learning on Graphs, Simplices, and Cells. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2604.17548.

Matthew Kahle, Elliot Paquette, and Érika Roldán. Topology of random 2-dimensional cubical complexes. *Forum of Mathematics, Sigma*, 9:e76, 2021. doi: 10.1017/fms.2021.64. URL https://www.cambridge.org/core/journals/forum-of-mathematics-sigma/article/topology-of-random-2dimensional-cubical-complexes/0C13EA097E0A9E98B4A66F0CF713A9BC.

Donghan Kim. Persistent Homology is Complementary to the Weisfeiler–Leman Test. 53rd KAIST Math Graduate Student Seminar, Department of Mathematical Sciences, KAIST, December 2025. URL https://mathsci.kaist.ac.kr/gradseminar/notice/2025-12-04/. Seminar announcement; talk on 4 December 2025; accessed 15 September 2026.

Sanjeevi Krishnan and Emily Rudman. A convenient category of cubes. arXiv:2503.13663v1, 2025. URL https://arxiv.org/html/2503.13663v1. Version 1, 17 March 2025.

Ippei Obayashi and Michio Yoshiwaki. Field Choice Problem in Persistent Homology. *Discrete & Computational Geometry*, 70:645–670, 2023. doi: 10.1007/s00454-023-00544-7. URL https://link.springer.com/article/10.1007/s00454-023-00544-7.

Mathilde Papillon, Sophia Sanborn, Mustafa Hajij, and Nina Miolane. Architectures of topological deep learning: A survey of message-passing topological neural networks. arXiv:2304.10031v3, 2024. URL https://arxiv.org/abs/2304.10031v3. Version 3, 21 February 2024; first submitted in 2023; preprint.

T. Mitchell Roddenberry, Nicholas Glaze, and Santiago Segarra. Principled simplicial neural networks for trajectory prediction. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9020–9029. PMLR, 2021. URL https://proceedings.mlr.press/v139/roddenberry21a.html.

The Stacks project authors. The Stacks project. https://stacks.math.columbia.edu, 2026. Tags 00NX and 0117; accessed 9 September 2026.

Diaaeldin Taha, James Chapman, Marzieh Eidi, Karel Devriendt, and Guido Montúfar. Demystifying Topological Message-Passing with Relational Structures: A Case Study on Oversquashing in Simplicial Message-Passing. In *International Conference on Learning Representations (ICLR)*, pp. 53412–53461, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/8433bb4f7477bf8202614ce1ae8b1169-Abstract-Conference.html. Claim locators use arXiv:2506.06582v1, Sections 1 and 2.

Yogesh Verma, Amauri H. Souza, and Vikas Garg. Topological Neural Networks go Persistent, Equivariant, and Continuous. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 49388–49407. PMLR, 2024. URL https://proceedings.mlr.press/v235/verma24a.html.

Maria Walch. Torsion in persistent homology and neural networks. arXiv:2506.03049v3, 2025. URL https://arxiv.org/abs/2506.03049v3. Version 3, revised 19 September 2025.

Maosheng Yang, Elvin Isufi, and Geert Leus. Simplicial convolutional neural networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8847–8851. IEEE, 2022. doi: 10.1109/ICASSP43922.2022.9746017.

APPENDIX OVERVIEW

A **Source construction.** Affine systems, complete face inventories and coefficient homology.
B **Lattice realization.** Boolean intervals, preserved homology and the minimum height.
C **Counting transfer.** Face encodings and counting-logic variable budgets.
D **Model bounds.** Local and paired observations, invariant readouts and parameter transfer.
E **Decision and readout.** The equal-observation bound and an exact separating rule.
F **Marked filling.** Six query marks, filling certificates and cocycle transport.
G **Neural interfaces.** Natural-neighborhood HOMP, local attention and order-four SMCN simulation.
H **Coherent coefficients.** Chain isomorphisms, prescribed cycles, broadcast and skeletal extensions, mixed coefficients and witness sizes.
I **Computational methods.** Related work, fusion-model equations and finite validation.

A AFFINE REALIZATIONS AND COEFFICIENT DEPTH

A.1 AFFINE INPUT AND CONSTRUCTION

Fix a prime p and an integer $s \geq 2$, and put $q = p^s$, $a = p^{s-1}$, $\epsilon \in \{0, a\}$. The letter ϵ always denotes the varying twist, so the two compared systems are E^0, E^a and $E^\epsilon, P(E^\epsilon), \gcd(q, \epsilon)$ are unambiguous; the indexed vertex symbols $a_{\rho,i}$ are distinct from the scalar a. Let H be a connected simple cubic graph with a distinguished vertex u, and

$$n = |V(H)|, \qquad m = |E(H)| = \tfrac{3n}{2}, \qquad b = m - n + 1 = \tfrac{n}{2} + 1 \ \geq 1.$$

Definition 9 (The affine source systems)**.** Let E^ϵ be the relational affine system over $\mathbb{Z}/q$ attached to (H, u, ϵ) in the ordered-tuple form of (Atserias et al., 2007, Definition 1), with universe

$$U = \{\, x_{v,e,i} \ : \ v \in V(H),\ e \in E(H) \text{ incident with } v,\ i \in \mathbb{Z}/q \,\}, \qquad |U| = 3qn,$$

since H has $2m = 3n$ vertex–edge incidences. The vocabulary consists of the fixed symbols R_c^j, $1 \leq j \leq 3$, $c \in \mathbb{Z}/q$, where R_c^j contains the ordered j-tuples constrained by $x_1 + \cdots + x_j = c$ in $\mathbb{Z}/q$. Local port enumerations and an order of the two ends of each edge are fixed solely to specify tuple slots and add no order or index predicate to the universe. The tuples of E^ϵ are: for every vertex v with incident edges e_1, e_2, e_3 and every $(i_1, i_2, i_3) \in (\mathbb{Z}/q)^3$, the ordered triple $(x_{v,e_1,i_1}, x_{v,e_2,i_2}, x_{v,e_3,i_3})$ of colour $c \equiv i_1 + i_2 + i_3 + \epsilon \cdot \mathbf{1}[v = u] \pmod q$; and for every edge $e = \{v, w\}$ and every $(i, j) \in (\mathbb{Z}/q)^2$, the ordered pair $(x_{v,e,i}, x_{w,e,j})$ of colour $c \equiv i + j \pmod q$. No construction index, orientation predicate or ordering of the universe belongs to this vocabulary; the only symbols are the R_c^j.

The construction uses the following affine-source indistinguishability theorem.

Theorem 10 (Affine-source indistinguishability; (Atserias et al., 2007, Definition 1 and Lemmas 2–3))**.** *For every $\ell \geq 1$, if $\mathrm{tw}(H) > \ell$ then*

$$E^0 \ \equiv_{C^\ell} \ E^\epsilon \qquad \text{for } \epsilon \in \{0, a\},$$

and this holds even for the infinitary counting logic with ℓ variables. Moreover, connected cubic base graphs H of arbitrarily large treewidth exist.

Remark 11. The cited definition and lemma numbers refer to the 2007 author version. The cited result establishes affine-source indistinguishability; we prove the realization and exact variable budgets below. See Grohe & Pakusa (2017) for counting-logic background.

We realize the affine presentation, encode every face with at most three source coordinates, and compute its coefficient modules.

A.2 THE ABELIAN PRESENTATION AND ITS EXACT VALUE

Definition 12 (The presentation attached to a system)**.** For a finite system E with universe U, use the following presentation:

$$P(E) = \Big\langle\, t,\ (z_x)_{x \in U} \ \Big|\ qt,\ (qz_x)_{x \in U},\ z_{x_1} + \cdots + z_{x_j} - c\,t \ \text{ for every } (x_1, \ldots, x_j) \in R_c^j \,\Big\rangle_{\mathrm{ab}}.$$

Repeated coordinates inside a tuple are allowed and the relator is read with multiplicity, e.g. $(x, x, y) \in R_c^3$ contributes $2z_x + z_y - ct$. Distinct coloured tuples remain distinct relators even on the same set of variables; a duplicate record of one and the same coloured tuple is not a second relation of E and is not instantiated twice.

Remark 13. $P(E)$ is an abelian presentation. Lemma 27 identifies it with first homology through $H_1(S(E); \mathbb{Z}) \cong P(E)$.

All manipulations below are integral row and generator transformations: no relation is discarded before it has been shown redundant, and there is no division in $\mathbb{Z}/q$.

Lemma 14 (Integral shear)**.** *The substitution $z_{v,e,i} = y_{v,e,i} + i\,t$ is an integral unimodular change of generators, and, with colours c read as residues in $\{0, \ldots, q-1\}$, it turns the presentation of $P(E^\epsilon)$, modulo integral multiples of the present relation qt, into the relators*

$$\sum_{k=1}^{3} y_{v,e_k,i_k} - \epsilon \cdot \mathbf{1}[v = u]\, t \quad \textit{(vertex tuples)}, \qquad y_{v,e,i} + y_{w,e,j} \quad \textit{(edge tuples)},$$

together with qt and the relators $qy_{v,e,i}$.

Proof. The substitution is triangular with unit diagonal. A vertex relator acquires the t-coefficient $i_1+i_2+i_3-c \equiv -\epsilon\mathbf{1}[v=u] \pmod q$, an edge relator $i+j-c \equiv 0$, and a power relator becomes $qy_{v,e,i} + i\,qt$. Subtracting integral multiples of the retained row qt gives the stated rows. □

Lemma 15 (Unit edge elimination and the incidence presentation)**.** *Let B be the signed incidence matrix of H for an auxiliary orientation of its edges and e_u its standard column indexed by u. Then*

$$P(E^\epsilon) \;\cong\; \big\langle\, t,\ (z_e)_{e\in E(H)} \;\big|\; qt,\ (qz_e),\ Bz - \epsilon\, e_u t \,\big\rangle_{\mathrm{ab}}.$$

The orientation serves only to write this integer algebra; it is not part of the source vocabulary and is not observed anywhere in Subsection A.4.

Proof. For $e=\{v,w\}$, subtract its $(0,0)$ row from the $(i,0)$ and $(0,j)$ rows. This gives $y_{v,e,i}-y_{v,e,0}$ and $y_{w,e,j}-y_{w,e,0}$, and every edge row is their sum with the $(0,0)$ row:

$$y_{v,e,i}+y_{w,e,j} = (y_{v,e,i}-y_{v,e,0}) + (y_{w,e,j}-y_{w,e,0}) + (y_{v,e,0}+y_{w,e,0}).$$

Unit elimination therefore leaves one generator $z_e = y_{v,e,0} = -y_{w,e,0}$. All vertex tuples at v then give $(Bz)_v - \epsilon\mathbf{1}[v=u]t$, and all power rows become copies of qz_e. □

Lemma 16 (Spanning-tree elimination and preservation of the q-torsion block)**.** *In the presentation of Lemma 15 the $n-1$ generators indexed by the edges of a spanning tree of H can be eliminated by integral unimodular operations, leaving exactly $b=m-n+1$ chord generators subject only to the relations qz, together with t subject to qt and ϵt. No q-torsion condition is lost.*

Proof. Replace the u-row by the sum of the vertex rows, obtaining $-\epsilon t$ because each column of B sums to zero. The reduced incidence matrix on any spanning tree has determinant ±1: expand at a leaf other than u and induct. Thus its inverse is integral, and the remaining $n-1$ vertex rows eliminate the tree generators in terms of the b chords.

These changes retain the torsion relations. Under a unimodular generator change V, the block qI becomes qV and a row change by V^{-1} restores qI. Eliminating a generator by a unit row turns its power row into a combination of the retained power rows. Hence only qz remains on each chord, while t has the two rows $qt, \epsilon t$, equivalent by the integer Euclidean algorithm to $\gcd(q,\epsilon)t$. □

Proposition 17 (Exact value of the presentation)**.** *For $\epsilon \in \{0,a\}$,*

$$P(E^\epsilon) \;\cong\; (\mathbb{Z}/q)^b \;\oplus\; \mathbb{Z}/\gcd(q,\epsilon), \qquad \gcd(q,0)=q.$$

Proof. Lemmas 14–16 leave b uncoupled q-torsion generators and the factor generated by t. □

Corollary 18 (The two integral groups)**.** *With $q=p^s$ and $a=p^{s-1}$,*

$$P(E^0) \cong (\mathbb{Z}/p^s)^{b+1}, \qquad P(E^a) \cong (\mathbb{Z}/p^s)^b \oplus \mathbb{Z}/p^{s-1}.$$

Granting the realization of Subsection A.3, these are the integral first homology groups

$$H_1(S(E^0);\mathbb{Z}) \cong (\mathbb{Z}/p^s)^{b+1}, \qquad H_1(S(E^a);\mathbb{Z}) \cong (\mathbb{Z}/p^s)^b \oplus \mathbb{Z}/p^{s-1},$$

used in Subsections A.6 and A.7. Neither group has a free summand. Their orders differ by the factor p, but both groups have exponent p^s: the distinguishing information lies in the multiset of cyclic factors, not in the exponent of the whole group.

Proof. Substitute $\gcd(p^s,0)=p^s$ and $\gcd(p^s,p^{s-1})=p^{s-1}$ into Proposition 17. Since $b\ge1$, both groups contain a factor of order p^s. □

A.3 EXPLICIT SIMPLICIAL REALIZATION

Here E is an arbitrary finite system in the signature of Definition 9, with universe U; put $G = \{t\} \sqcup U$ and $g = |G| = |U| + 1$.

Definition 19 (Base graph)**.** Let B_G have vertex set $\{o\} \cup \{A_z, B_z : z \in G\}$ and, for each $z \in G$, the three edges oA_z, A_zB_z, B_zo. The triangle on $\{o, A_z, B_z\}$ is *not* declared. We write z also for the positive closed walk o, A_z, B_z, o. Distinct circles meet exactly in o.

Definition 20 (Relator instances and their words)**.** The relator instances are: (i) for every $z \in G$, including $z = t$, one distinct *power relator* with word z^q; (ii) for every coloured source tuple $\rho = (j, c, x_1, \dots, x_j)$ with $(x_1, \dots, x_j) \in R_c^j$, one *equation relator* with the following positive word:

$$x_1 x_2 \cdots x_j \, t^d, \qquad d = (-c) \bmod q \in \{0, \dots, q-1\}.$$

Each letter is traversed as the three-edge circle o, A_z, B_z, o, so a word with ℓ letters expands to the closed walk $y_0, y_1, \dots, y_N = y_0$ with $N = 3\ell \geq 3$.

Lemma 21 (The positive words present the same abelian group)**.** *The abelianized word of the equation relator for ρ equals the relator $z_{x_1} + \cdots + z_{x_j} - c\,t$ of Definition 12 plus an integer multiple of qt, and the abelianized word of a power relator equals qz. Since the power relator for t is present, the abelian group presented by the words of Definition 20 is $P(E)$.*

Proof. Since $d \equiv -c \pmod q$, abelianization changes the equation row only by a multiple of the present power row qt. □

Lemma 22 (Nondegeneracy of the walk)**.** *For every relator instance ρ the expanded walk satisfies $N \geq 3$, $y_N = y_0$ and $y_i \neq y_{i+1}$ for all i (indices mod N), regardless of repeated letters, repeated source coordinates, or repeated visits to o.*

Proof. Each word has at least one letter. Its expansion traverses o, A_z, B_z, o for each letter, with junctions $B_z, o, A_{z'}$; every adjacent pair is distinct, even for repeated letters. □

Definition 23 (The complex $S(E)$)**.** For each relator instance ρ with walk length $N = N_\rho$, adjoin private vertices $a_{\rho,0}, \dots, a_{\rho,N-1}$, indices read cyclically modulo N, and one apex h_ρ. These are pairwise distinct, distinct from the vertices of B_G, and distinct from the private vertices of every other instance. Fixing ρ and writing $a_i = a_{\rho,i}$, declare exactly the triangles

$$A_i = \{a_i, a_{i+1}, y_{i+1}\}, \qquad B_i = \{a_i, y_i, y_{i+1}\}, \qquad C_i = \{h_\rho, a_i, a_{i+1}\} \qquad (i \in \mathbb{Z}/N),$$

and let $S(E)$ be the abstract simplicial complex generated by B_G together with the downward closure of all these triangles, over all instances ρ.

Lemma 24 (Dimension)**.** *Every declared triangle has three distinct vertices, and $S(E)$ has dimension exactly 2.*

Proof. Use $N \geq 3$, distinct private vertices and $y_i \neq y_{i+1}$. No higher simplex is declared, and the power relator for t supplies a triangle. □

A.3.1 ALL FACES AND THE ABSENCE OF UNINTENDED IDENTIFICATIONS

Lemma 25 (Complete face inventory)**.** *Fix an instance ρ and write $a_i = a_{\rho,i}$. The edges of $S(E)$ that are not edges of B_G and arise from ρ are exactly*

$$d_i = \{a_i, a_{i+1}\}, \quad v_i = \{a_i, y_i\}, \quad z_i = \{a_i, y_{i+1}\}, \quad k_i = \{h_\rho, a_i\} \qquad (i \in \mathbb{Z}/N),$$

all remaining edges of A_i, B_i, C_i being edges $\{y_i, y_{i+1}\}$ of B_G. Within one instance the four edge families and the three triangle families are pairwise disjoint and each is indexed injectively by i; across instances, private faces are disjoint. Consequently the list

$$\{o\} \cup \{A_z, B_z\}, \quad \textit{the three circle edges per } z, \quad \{a_{\rho,i}\},\ \{h_\rho\},\ d_i, v_i, z_i, k_i,\ A_i, B_i, C_i$$

represents every nonempty face of $S(E)$ exactly once, provided the faces of B_G are represented globally, i.e. once and not once per instance.

Proof. The edge boundaries are as follows:

$$\partial A_i = \{d_i, z_i, v_{i+1}\}, \quad \partial B_i = \{v_i, z_i, \{y_i, y_{i+1}\}\}, \quad \partial C_i = \{k_i, k_{i+1}, d_i\}.$$

The cyclic pairs $\{i, i+1\}$ are distinct for $N \geq 3$, and $v_i \neq z_i$ because adjacent walk values differ. Private-vertex patterns distinguish d, k, the mixed edges v, z, and the old edges. They also distinguish the three triangle families; d_i determines the index in A, C, and a_i in B. Every private face retains its instance's private vertices, so different instances cannot collide. Downward closure adds exactly the listed faces. □

Lemma 26 (Repeated coordinates and duplicate records)**.** *Repeated source coordinates inside a tuple identify only the corresponding old generator vertices and circle edges; they never identify private vertices of a relator, and never make two adjacent walk values equal. Two distinct relation symbols, or two distinct ordered tuples of the same symbol, always yield distinct relator instances with disjoint private faces, even if their expanded words coincide letterwise. A duplicate record of one and the same coloured tuple is not a second relation and is instantiated only once.*

Proof. Private vertices carry the entire instance index $\rho = (j, c, x_1, \ldots, x_j)$ and their position. Coinciding coordinates therefore affect only old vertices. Lemma 22 handles adjacent walk values, and relations are sets of tuples. □

A.3.2 The realized attachment is the intended mapping cone

Lemma 27 (Mapping cylinder, cone, and H_1)**.** *Fix an instance ρ with walk $y_0, \ldots, y_N = y_0$. Triangulate the cylinder over the N-cycle with bottom vertices a_i and top vertices b_i using triangles $\{a_i, a_{i+1}, b_{i+1}\}$, $\{a_i, b_i, b_{i+1}\}$ $(i \in \mathbb{Z}/N)$. Attach the top edgewise to B_G by $b_i \mapsto y_i$. Cone the bottom cycle to the new apex h_ρ. Then:*

1. *the quotient triangles are exactly A_i, B_i, C_i of Definition 23;*
2. *every simplex of the cylinder-with-cone embeds in the quotient, and any two image simplices meet in a common face; hence the quotient is homeomorphic to the geometric realization of $S(E)$ restricted to B_G and this instance;*
3. *globally, $S(E)$ is obtained from B_G by attaching one 2-cell along the edge word of each relator instance, and consequently $H_1(S(E); \mathbb{Z}) \cong P(E)$.*

Proof. The quotient sends the two strip triangles to A_i, B_i and the cone triangle to C_i. They remain nondegenerate. To check the quotient triangulation, pairs sharing private vertices have only the following intersections: adjacent A triangles share an a-vertex; A_i meets B_i, B_{i+1} in z_i, v_{i+1}; cone triangles meet along a spoke or at the apex; and cone–strip intersections lie in the bottom cycle. These statements include cyclic wraparound and $N = 3$. Pairs without a common private vertex intersect only through their top faces in the graph B_G. Thus each image simplex embeds and intersections are common faces, with exactly the declared identifications.

The cylinder capped at its bottom is a disk attached along the given top walk, even when that walk revisits old vertices or edges. Private faces of different instances are disjoint, so these are precisely the specified relator attachments. The standard presentation-space and abelianization results (Hatcher, 2002, Proposition 1.26 and Theorem 2A.1), together with Lemma 21, give $H_1(S(E); \mathbb{Z}) \cong P(E)$. □

A.3.3 Connectedness, purity and census

Lemma 28 (Connectedness and purity)**.** *$S(E)$ is connected and pure of dimension 2.*

Proof. Every private vertex connects to B_G through v_i or a spoke k_i, giving connectedness. Every private face lies in a declared triangle. Each old circle edge occurs in its generator's power relator and hence in a B_i triangle; every old vertex lies on such an edge. □

Lemma 29 (Census). *Let M be the number of relator instances and $\Lambda = \sum_\rho N_\rho$. Then*

$$f_0 = 1 + 2g + \Lambda + M, \qquad f_1 = 3g + 4\Lambda, \qquad f_2 = 3\Lambda, \qquad \chi = f_0 - f_1 + f_2 = 1 - g + M.$$

For the ordered-source convention of Definition 9, $g = |U| + 1 = 1 + 3qn$ and $M = g + nq^3 + mq^2$, every power word has $N_\rho = 3q$, and every equation word has $N_\rho = 3(j + d) \le 3(q + 2)$ since $j \le 3$ and $d \le q - 1$. Consequently

$$\#\{\textit{nonempty faces}\} = f_0 + f_1 + f_2 = 1 + 5g + M + 8\Lambda \;\le\; 1 + 5g + (24q + 49)M = O_q(n),$$

using $m = \frac{3n}{2}$.

Proof. The inventory in Lemma 25 contributes $1 + 2g, 3g$ old vertices and edges, and $N_\rho + 1, 4N_\rho, 3N_\rho$ faces per instance. Summing gives the formulas. There are g power, nq^3 vertex and mq^2 edge instances, and $\Lambda \le 3(q + 2)M$. □

Remark 30 (Face-count bound). Lemma 29 is a face-count upper bound for fixed q, with a finite constant that depends on q but not on r or on the size of the source.

Lemma 31 (The paired inputs have equal parameters). *The systems E^0 and E^a have the same universe and the same number of ordered tuples of each arity and each colour, hence the same g, M and Λ. Therefore $S(E^0)$ and $S(E^a)$ have the same f-vector and the same Euler characteristic $\chi = 1 - g + M$.*

Proof. For each vertex and colour c, any two indices uniquely determine the third, giving q^2 triples of that colour for either twist; each edge similarly gives q pairs per colour. Arity-1 relations are empty and power instances agree. Equation lengths depend only on arity and colour, so g, M, Λ and hence the face vectors agree. □

A.4 The full incidence structure is a bounded tagged interpretation

Definition 32 (The observed structure). For a complex S of dimension at most 2, let $I(S)$ be the structure whose universe is the set of *all* nonempty faces of S — vertices, edges and triangles alike — with the unary dimension predicates $\mathrm{Dim}_0, \mathrm{Dim}_1, \mathrm{Dim}_2$ and the binary strict containment relation Sub, where $\mathrm{Sub}(F, G)$ holds iff $F \subsetneq G$. In particular vertex-in-triangle containment is observed. No further predicate is present: the internal tags below are indices of the translation, not observable colours, and no predicate names individual source elements, construction roles or algebraic orientations; no global order is supplied.

We exhibit $I(S(E))$ as a quantifier-free interpretation in E with tuple arity at most 3. For each fixed q the following is a finite tag schema; expanding its stated finite index ranges produces a literal finite table.

Table 2: Tag schema for the nonempty faces of $S(E)$.

Face family	Tuple arity	Source index and domain	Internal tags
old vertices	0 and 1	$()$ for o, A_t, B_t; $x \in U$ for A_x, B_x	distinct base tag for o and distinct generator-position tags for the A/B roles
old edges	0 and 1	$()$ for the three t-circle edges; $x \in U$ for the three x-circle edges	edge position $0, 1, 2$ on the circle
t^q instance	0	$()$, with the true (always satisfied) domain formula	one position-independent tag for h_ρ; for each fixed $i \in \{0, \dots, 3q - 1\}$, tags for $a_{\rho,i}$, the four edge families and the three triangle families

Face family	Tuple arity	Source index and domain	Internal tags
x^q instance	1	$x \in U$, with domain formula "$x = x$"	one position-independent apex tag and the same position-indexed inventory for $i \in \{0, \dots, 3q-1\}$
$\rho = (j, c, x_1, \dots, x_j)$ instance	$j \in \{1, 2, 3\}$	$(x_1, \dots, x_j) \in R^j_c$, one tag family per pair (j, c)	one position-independent apex tag and the same position-indexed inventory, with $i \in \{0, \dots, N-1\}$ and $N = 3\big(j + ((-c) \bmod q)\big)$ fixed by (j, c)

Lemma 33 (Faithfulness of the schema)**.** *Every tag arity lies in* $\{0, 1, 2, 3\}$*. The disjoint union of the tagged domains of Table 2 is in bijection with the set of all nonempty faces of* $S(E)$*, without any quotient.*

Proof. The table assigns arity 0, 1 or $j \le 3$ according to the instance index. Lemmas 25 and 26 give the bijection because private tags retain the full ordered tuple. □

Remark 34 (Entire tuples are retained)**.** Every private face retains the entire ordered relator tuple, including unused endpoint coordinates; projecting any away could merge distinct instances. The apex tag is position-independent, and other private tags record role and position. These tags are translation indices, not observed colours.

Lemma 35 (Quantifier-free atoms)**.** *Every listed tag comes with a list of one, two or three* vertex descriptors*, each consisting of a tag and coordinate projections of the face's source tuple. Descriptors of old vertices are* o*,* A_t/B_t*, or* A_{x_l}/B_{x_l} *with* l *determined by the fixed word position; descriptors of private vertices consist of their fixed role, their fixed internal position and the full relator tuple. Then, on legal representations:*

1. *equality of two represented faces is false whenever the tags differ, is coordinatewise equality of the source tuples when the tags agree, and is true when the tag has arity* 0*;*
2. Dim_i *is the fixed dimension of the tag;*
3. *strict containment is given by the quantifier-free formula*

$$\mathrm{Sub}(F, G) \iff \dim F < \dim G \wedge \bigwedge_{v \in \mathrm{Vert}(F)} \bigvee_{w \in \mathrm{Vert}(G)} \mathrm{EqVertex}(v, w),$$

where EqVertex *is the finite tag test together with the corresponding coordinate equalities.*

All conjunctions and disjunctions are finite expansions over lists of at most three descriptors and are not quantifiers over a source domain; a binary face atom uses at most six supplied source symbols. All dimension pairs not listed give the value false, and no construction-role predicate is added to $I(S(E))$*.*

Proof. Equality follows from the bijection, and dimension is fixed by the tag. In an abstract simplicial complex strict containment is precisely inclusion of vertex sets with increasing dimension. Expanding the lists of at most three descriptors uses only the two arguments' coordinates, hence at most six symbols and no auxiliary witnesses. □

A.5 Ordinary counting translation with cost $3r$

Ordinary counting quantifiers count *elements*, not tuples, so the tuple counts required by Lemma 35 must be produced rather than assumed.

Lemma 36 (Ordinary counting of ordered tuples). *Let $d \in \{0,1,2,3\}$, let $\theta(u_1,\dots,u_d,\bar{z})$ be a formula and let $n \geq 0$ be a fixed natural number. Define $\mathsf{C}^d_{\geq n}[\theta]$ by:*

- $\mathsf{C}^d_{\geq 0}[\theta] = \mathrm{true}$ *for every* $d \in \{0,1,2,3\}$*;*
- *for* $d = 0$*:* $\mathsf{C}^0_{\geq 1}[\theta] = \theta$ *and* $\mathsf{C}^0_{\geq n}[\theta] = \mathrm{false}$ *for* $n \geq 2$*;*
- *for* $d = 1$ *and* $n \geq 1$*:* $\mathsf{C}^1_{\geq n}[\theta] = \exists^{\geq n} u_1\, \theta$*, the ordinary unary counting quantifier;*
- *for* $d \geq 2$ *and* $n \geq 1$*:*

$$\mathsf{C}^d_{\geq n}[\theta] = \bigvee_{\substack{m_1,\dots,m_n \in \{0,\dots,n\} \\ m_1+\cdots+m_n=n}} \bigwedge_{j\,:\,m_j>0} \exists^{\geq m_j} u_1\ \mathsf{C}^{d-1}_{\geq j}[\theta],$$

where the inner count binds $u_2,\dots,u_d$ and leaves u_1 free as a parameter.

Then, on every structure and every assignment to $\bar{z}$, $\mathsf{C}^d_{\geq n}[\theta]$ holds iff at least n ordered d-tuples $(u_1,\dots,u_d)$ satisfy θ.

Proof. For finite structures this is Lemma 76 with designated block $(u_1,\dots,u_d)$ and $d \leq 3$; the displayed names are distinct and disjoint from the remaining free parameters. Its zero-arity convention counts the unique empty tuple, and its zero-threshold convention applies also to an empty universe. The block may contain names not free in θ.

The same finite-threshold argument applies to an arbitrary structure. A tuple count of at least n has a finite set of n witnesses. Conversely, a successful allocation has only finitely many positive requirements, each with finitely many first-coordinate and fibre witnesses. Restricting the satisfying relation to these witnesses reduces either direction to the finite layer identity in that lemma. The formula reuses the designated names, permits repeated coordinate values and introduces no auxiliary name. □

Proposition 37 (The $3r$ translation). *Fix q and a pool of at most $r \geq 1$ target variable symbols. For each C^r formula φ and each fixed tag assignment $\boldsymbol{\tau}$ to its free variables, there is a C^{3r} source formula $\varphi^*_{\boldsymbol{\tau}}$ whose free variables are contained in the active coordinates of those free-variable representations, such that for every finite system E and every legal assignment of faces represented with these prescribed tags and coordinates $\bar{u}$,*

$$I(S(E)), \eta_{\boldsymbol{\tau},\bar{u}} \models \varphi \quad \Longleftrightarrow \quad E \models \varphi^*_{\boldsymbol{\tau}}(\bar{u}),$$

where $\eta_{\boldsymbol{\tau},\bar{u}}$ assigns each free variable the face represented by its prescribed tag and coordinates. For a sentence the free-tag assignment is empty, giving a single C^{3r} sentence φ^*. Thus the sentence-translation cost is $c_{p,s}(r) = 3r$, independent of source size.

Proof. Reserve three disjoint source symbols for each target symbol x_i, using only the first $d(\sigma)$ coordinates when its tag is σ. No representation is padded. Translate atoms by Lemma 35; for Boolean connectives restrict the tag assignment to each operand's free variables.

For $\exists^{\geq n} x_i\, \psi$, keep all other free tags fixed and put $\theta_{\boldsymbol{\tau},\sigma} = D_\sigma \wedge \psi^*_{\boldsymbol{\tau}[x_i \mapsto \sigma]}$, restricting the extended assignment to $\mathrm{FV}(\psi)$. Use the following translation:

$$\bigvee_{\substack{(n_\sigma)\in\{0,\dots,n\}^{\mathcal{T}} \\ \sum_\sigma n_\sigma = n}} \bigwedge_{\sigma\in\mathcal{T}} \mathsf{C}^{d(\sigma)}_{\geq n_\sigma}[\theta_{\boldsymbol{\tau},\sigma}].$$

The tagged representation is bijective, so its witness count is the sum of these tag counts. The allocation formula is therefore correct by Lemma 36. Arity-0 tags contribute 0 or 1, including for an empty source universe.

This induction preserves truth on legal representations and keeps free source symbols within the active coordinates of free target variables. Quantifying x_i binds only its reserved block, and nested rebinding shadows the same block; other blocks are disjoint, so no capture occurs. Lower-arity

tags leave unused coordinates absent. All Boolean branches reuse the same pool of $3r$ symbols. If needed, true and false are local closed formulas over one reserved symbol. The empty tag assignment for sentences yields the asserted C^{3r} translation, with no formula-length or quantifier-rank bound. □

Corollary 38 (Indistinguishable realizations with equal f-vectors)**.** *Let $r \geq 2$ and choose, by Theorem 10, a connected cubic H with $\mathrm{tw}(H) > 3r$. Then $I(S(E^0)) \equiv_{C^r} I(S(E^a))$, and the two complexes have equal f-vectors.*

Proof. Apply Theorem 10 at $\ell = 3r$ and pull back sentences using Proposition 37. Equal face vectors follow from Lemma 31. □

A.6 Integral splitting and coefficients in an arbitrary ring

Throughout this subsection and the next, A denotes a commutative coefficient ring; the letters A_i, A_z of Subsection A.3 do not occur here. Set $S = S(E^0)$ and $T = S(E^a) = S(E^{p^{s-1}})$.

Lemma 39 (Integral shape of the chain complex)**.** *S and T are connected, two-dimensional, have the integral first homology groups of Corollary 18, and have equal Euler characteristic. Consequently $H_0(\,\cdot\,;\mathbb{Z}) = \mathbb{Z}$, $H_1(\,\cdot\,;\mathbb{Z})$ is finite, $H_2(\,\cdot\,;\mathbb{Z})$ is free of the common rank $h_2 = \chi - 1 = M - g$, and $H_i = 0$ for $i > 2$. For each $X \in \{S, T\}$, integral invertible basis changes split its full cellular chain complex as a direct sum of a contractible complex $\mathcal{U}_X$ and the residual summands*

$$\mathbb{Z} \text{ in degree } 0, \qquad \left(\mathbb{Z} \xrightarrow{d} \mathbb{Z}\right) \text{ in degrees } 2, 1 \text{ for each nonunit Smith factor } d, \qquad \mathbb{Z}^{h_2} \text{ in degree } 2,$$

where $\mathcal{U}_X$ is a finite direct sum of unit complexes $\mathbb{Z} \xrightarrow{1} \mathbb{Z}$ in degrees $(1,0)$ or $(2,1)$. The residual summand has no unmatched degree-1 generator.

Proof. The preceding results give connectedness, finite H_1, equal χ, and dimension two. Split off the spanning-tree unit blocks of d_1, leaving one degree-0 copy of $\mathbb{Z}$ and the free cycle group $\ker d_1$ in degree 1. Integral Smith basis changes for $C_2 \to \ker d_1$ then give unit blocks and the displayed nonunit blocks; this is the standard free-chain decomposition (Hatcher, 2002, §3A, pp. 262–264). Finite H_1 excludes an unmatched degree-1 generator. The remaining degree-2 generators form H_2, free of rank $\chi - 1 = M - g$. All changes are integral and invertible. □

Proposition 40 (Homology over an arbitrary commutative ring)**.** *Let A be any commutative ring and let $X \in \{S, T\}$ have nonunit Smith factors d as in Lemma 39. Then*

$$H_0(X;A) = A, \qquad H_1(X;A) = \bigoplus_d A/dA,$$
$$H_2(X;A) = A^{h_2} \oplus \bigoplus_d A[d], \quad H_i(X;A) = 0 \quad (i > 2),$$

where $A[d] = \{x \in A : dx = 0\}$.

Proof. Integral invertible basis changes remain invertible after base change to A. Unit blocks become contractible $A \xrightarrow{1} A$, while each $A \xrightarrow{d} A$ has cokernel A/dA and kernel $A[d]$. The isolated degree-0 and degree-2 copies give the remaining summands. This is an A-module computation and needs no flatness assumption. □

A.7 Every field and the exact coefficient-depth boundary

By Corollary 18, the nonunit Smith factors of S are p^s with multiplicity $b+1$, and those of T are p^s with multiplicity b together with one factor p^{s-1}.

Proposition 41 (All fields)**.** *Let K be any field, of any characteristic, including arbitrary extensions.*

- *If $\mathrm{char}\, K = p$, then for both $X = S$ and $X = T$,*

$$H_0(X;K) = K, \qquad H_1(X;K) = K^{b+1},$$
$$H_2(X;K) = K^{h_2+b+1}, \quad H_i(X;K) = 0 \quad (i > 2).$$

- *If* $\operatorname{char} K \neq p$*, in particular if* $\operatorname{char} K = 0$*, then for both* $X = S$ *and* $X = T$*,*

$$H_0(X;K) = K, \qquad H_1(X;K) = 0,$$
$$H_2(X;K) = K^{h_2}, \quad H_i(X;K) = 0 \quad (i > 2).$$

In particular no field detects the difference between S *and* T *in any degree.*

Proof. In Proposition 40, every nonunit factor acts as zero when $\operatorname{char} K = p$ and is invertible otherwise; both sides have $b + 1$ factors. □

Proposition 42 (The prime-power tower and the first distinguishing depth)**.** *Let* $A_j = \mathbb{Z}/p^j$ *for a positive integer* j*.*

- *For every* j *with* $1 \leq j < s$*, both* $X = S$ *and* $X = T$ *satisfy the following identities:*

$$H_0(X;A_j) = A_j, \qquad H_1(X;A_j) = A_j^{b+1},$$
$$H_2(X;A_j) = A_j^{h_2+b+1}, \quad H_i(X;A_j) = 0 \quad (i > 2).$$

- *For* $j = s$*,*

$$H_1(S;A_s) = A_s^{b+1}, \qquad H_2(S;A_s) = A_s^{h_2+b+1},$$
$$H_1(T;A_s) = A_s^{b} \oplus \mathbb{Z}/p^{s-1}, \qquad H_2(T;A_s) = A_s^{h_2+b} \oplus \mathbb{Z}/p^{s-1}.$$

These modules have different orders, so both H_1 *and* H_2 *separate* S *from* T *at depth* s*.*

Hence s *is the first distinguishing depth in the tower* $(A_j)_{j\geq 1}$ *of positive-integer prime-power coefficients.*

Proof. For $e \geq 1$, both $A_j/p^e A_j$ and $A_j[p^e]$ are cyclic A_j-modules of order $p^{\min(e,j)}$. Substitute $e = s, s-1$ into Proposition 40. If $j < s$, both contributions are A_j; at $j = s$ the exceptional factor on T has order p^{s-1} instead of p^s. Thus the first difference occurs at depth s in both degrees. □

Remark 43 (Coefficient information)**.** These are module isomorphism types, and Appendix H supplies compatible shallow coefficient diagrams.

The parameter substitution and full model consequences of this source construction are given in Appendix D.

B Height-Four Boolean-Interval Lifts and the Structural Threshold

This appendix realizes arbitrary finite abelian first-homology torsion by full Boolean-interval cube complexes and proves the sharp height-four threshold. The main text applies this construction to its matched pair. Geometric conventions follow Appendix A. Section B.4 lists the encoding inventory; Appendix C proves the bounded-tuple encoding and gives the complete cell and atomic face tables.

B.1 Source complexes and the two-guard lattice $L_2(S)$

Definition 44 (Source complex)**.** A *source complex* is a finite connected pure abstract simplicial 2-complex S with vertex set V, edge set E and triangle set F, and f-vector $(v, e, f) = (f_0(S), f_1(S), f_2(S))$. Purity of dimension 2 means that every maximal face is a triangle; equivalently every vertex and every edge lies in some triangle. Individual vertices are written u, individual edges η or ε, individual triangles T; the symbols v, e, f always denote the three face counts and never an individual face. For $\eta \in E$ let $d_\eta = \#\{T \in F : \eta \subset T\}$, and for $u \in V$ let $d(u)$ be the number of incident edges and $f(u)$ the number of incident triangles. No manifold hypothesis and no upper bound on d_η is imposed.

Definition 45 (The two-guard lift)**.** Let S be a source complex. The poset $L_2(S)$ has underlying set

$$\{\hat{0}\} \sqcup \big(V \sqcup \{\lambda_\eta\}_{\eta\in E}\big) \sqcup \{E_\eta\}_{\eta\in E} \sqcup \big(\{t_T\}_{T\in F} \sqcup \{\rho^1_\eta, \rho^2_\eta\}_{\eta\in E}\big) \sqcup \{\hat{1}\},$$

the five blocks carrying ranks $0, 1, 2, 3, 4$ respectively. The λ_η are the *lower guards* and ρ^1_η, ρ^2_η the *upper guards* of η; E_η is the *edge element* of η and t_T the *triangle element* of T. The five rank symbols are pairwise disjoint. The order is the transitive closure of exactly the following cover families, and of no others:

(C1) $\hat{0} \lessdot u$ for every $u \in V$;

(C2) $\hat{0} \lessdot \lambda_\eta$ for every $\eta \in E$;

(C3) $u \lessdot E_\eta$ whenever $u \in \eta$;

(C4) $\lambda_\eta \lessdot E_\eta$ for every $\eta \in E$;

(C5) $E_\eta \lessdot t_T$ whenever $\eta \subset T$;

(C6) $E_\eta \lessdot \rho^j_\eta$ for $j = 1, 2$;

(C7) $t_T \lessdot \hat{1}$ for every $T \in F$;

(C8) $\rho^j_\eta \lessdot \hat{1}$ for $j = 1, 2$.

All further comparabilities are transitive consequences of (C1)–(C8).

Counting the families gives the following totals:

$$\begin{aligned} |L_2(S)| &= v + 4e + f + 2, \\ \#\{\text{covers}\} &= \underbrace{v}_{\text{C1}} + \underbrace{e}_{\text{C2}} + \underbrace{2e}_{\text{C3}} + \underbrace{e}_{\text{C4}} + \underbrace{3f}_{\text{C5}} + \underbrace{2e}_{\text{C6}} + \underbrace{f}_{\text{C7}} + \underbrace{2e}_{\text{C8}} = v + 8e + 4f. \end{aligned}$$

Writing $\eta = \{u_1, u_2\}$, the principal order ideals are as follows:

$$\begin{aligned} &\downarrow u = \{\hat{0}, u\}, \qquad \downarrow \lambda_\eta = \{\hat{0}, \lambda_\eta\}, \qquad \downarrow E_\eta = \{\hat{0}, u_1, u_2, \lambda_\eta, E_\eta\}, \\ &\downarrow \rho^j_\eta = (\downarrow E_\eta) \cup \{\rho^j_\eta\}, \quad \downarrow t_T = \{\hat{0}\} \cup \bigcup_{\eta\subset T} (\downarrow E_\eta) \cup \{t_T\}, \quad \downarrow \hat{1} = L_2(S). \end{aligned}$$

B.2 Meets, joins and exact height 4

Lemma 46 (All meets exist)**.** *Any two elements of $L_2(S)$ have a greatest lower bound, given case by case as follows.*

1. *Equality, and any pair involving $\hat{0}$ or $\hat{1}$: $x \wedge x = x$, $x \wedge \hat{0} = \hat{0}$, $x \wedge \hat{1} = x$.*
2. *A rank-1 element x meets a comparable element in x and an incomparable element in $\hat{0}$.*
3. *Two distinct edge elements E_η, E_ε meet in their common vertex if $\eta \cap \varepsilon$ is a vertex, and in $\hat{0}$ otherwise.*
4. *An edge element E_η and a rank-3 element z are either comparable, in which case the meet is E_η, or meet in the unique common vertex of η with the support of z, or in $\hat{0}$.*
5. *Two distinct triangle elements $t_T, t_{T'}$ meet in $\hat{0}$, in a vertex, or in an edge element, according as $T \cap T'$ is empty, a vertex, or an edge.*
6. *A triangle element t_T and an upper guard ρ^j_η obey the same rule with ρ^j_η treated as supported on η: the meet is E_η if $\eta \subset T$, the common vertex if $|\eta \cap T| = 1$, and $\hat{0}$ if $\eta \cap T = \varnothing$.*
7. *Two distinct upper guards of the same edge meet in E_η; upper guards on distinct edges meet in their common vertex if there is one, and in $\hat{0}$ otherwise.*

Moreover λ_η lies below a rank-3 element exactly when E_η does.

Proof. Intersect the displayed principal ideals. For a rank-1 element the intersection is its two-element ideal or $\{\hat{0}\}$. Distinct edge supports intersect in at most one vertex; an edge not contained in a triangle also intersects it in at most one vertex. Distinct triangles intersect in an empty face, a vertex or an edge, giving ideals with maxima $\hat{0}, u, E_\eta$ respectively. An upper guard has the ideal of its edge plus the guard itself, so guard pairs reduce to the same cases, except that two guards of one edge meet at E_η. These exhaust ranks $(1,*),(2,2),(2,3),(3,3)$; ranks $0, 4$ and equality are immediate. The lower-guard assertion follows because its only upward cover is E_η. □

Proposition 47 (Lattice and height)**.** *$L_2(S)$ is a finite lattice of height exactly 4.*

Proof. A finite meet-semilattice with top has joins: $x \vee y$ is the meet of their nonempty finite set of common upper bounds. All covers increase the displayed rank by one, and any flag $u \in \eta \subset T$ gives $\hat{0} < u < E_\eta < t_T < \hat{1}$, so the height is exactly four. □

B.3 Exhaustion of the Boolean intervals

Proposition 48 (Rank-two intervals)**.** *The Boolean intervals of rank 2 in $L_2(S)$ are exactly the* corner intervals $[u, t_T]$ *with $u \in T$. There are $3f$ of them.*

Proof. Since the rank function is compatible with covers, a rank-2 interval has endpoint ranks $(0,2)$, $(1,3)$ or $(2,4)$, and is Boolean iff it has exactly two intermediate elements, all incomparable. The table exhausts the comparable endpoint pairs.

interval	condition	intermediate elements	count
$[\hat{0}, E_\eta]$	$\eta = \{u_1, u_2\} \in E$	$u_1,\ u_2,\ \lambda_\eta$	3
$[u, t_T]$	$u \in T$	E_η for the two η with $u \in \eta \subset T$	2
$[\lambda_\eta, t_T]$	$\eta \subset T$	E_η	1
$[\lambda_\eta, \rho_\eta^j]$	$j \in \{1,2\}$	E_η	1
$[u, \rho_\eta^j]$	$u \in \eta$, $j \in \{1,2\}$	E_η	1
$[E_\eta, \hat{1}]$	$\eta \in E$	t_T for $\eta \subset T$; ρ_η^1, ρ_η^2	$d_\eta + 2 \geq 3$

Only $[\hat{0}, E_\eta]$ has endpoint ranks $(0,2)$ and only $[E_\eta, \hat{1}]$ has $(2,4)$; purity gives $d_\eta \geq 1$, so the last row has at least three intermediates. In the second row a triangle contains exactly two edges through a given vertex, so $[u, t_T]$ has two atoms and two coatoms and is B_2. All other rows have one or at least three intermediates. □

Proposition 49 (No Boolean interval of rank ≥ 3)**.** *$L_2(S)$ contains no interval isomorphic to B_3, and hence none isomorphic to B_k for any $k \geq 3$.*

Proof. A rank-3 interval has endpoint ranks $(0,3)$ or $(1,4)$, and B_3 has exactly three atoms and three coatoms. The four possible types are:

interval	atoms / coatoms of the interval	count
$[\hat{0}, t_T]$	atoms: three vertices of T and three lower guards λ_η, $\eta \subset T$	6
$[\hat{0}, \rho_\eta^j]$	coatoms: E_η only	1
$[\lambda_\eta, \hat{1}]$	atoms: E_η only	1
$[u, \hat{1}]$	coatoms: t_T for $T \ni u$, and ρ_η^j for $\eta \ni u$, $j = 1,2$	$f(u) + 2d(u) \geq 5$

Purity gives $f(u) \geq 1$ and, since a triangle through u carries two edges through u, also $d(u) \geq 2$; hence $f(u) + 2d(u) \geq 5$. None of the four types is B_3, and they exhaust the rank-3 endpoint types. Finally, any interval isomorphic to B_k with $k \geq 4$ contains a rank-3 subinterval isomorphic to B_3, just excluded. □

B.4 Cells and strict faces of the cube complex

Definition 50 (Full Boolean-interval cube complex)**.** For a finite lattice L, the complex $\mathcal{Q}(L)$ has one cell for *every* interval $[x, y]$ of L that is isomorphic to a Boolean lattice B_k; this cell is a k-cube. In particular every singleton $[x, x]$ is a 0-cell and every cover $x \lessdot y$ is a 1-cell. The combinatorial face relation is subinterval inclusion, $[x', y'] \subseteq [x, y]$ with $x \le x' \le y' \le y$, and $\mathrm{Face}(\sigma, \tau)$ denotes the induced *strict* geometric face relation between cells of different dimension in the realisation. The lattice rank of an element is unrelated to the dimension of the cell it indexes: a 0-cell $[x, x]$ may carry lattice rank $2, 3$ or 4. The complex need not be pure; for instance $[\hat{0}, \lambda_\eta]$ is a maximal 1-cell of $\mathcal{Q}(L_2(S))$. We never replace $\mathcal{Q}(L)$ by its corner subcomplex.

By the element and cover counts of Definition 45 and Propositions 48–49, the full complex has $v + 4e + f + 2$ vertices, $v + 8e + 4f$ edges and $3f$ corner squares. The corner square $[u, t_T]$ has vertices u, the two E_η with $u \in \eta \subset T$, and t_T, with the four intervening cover edges. These incidences use at most two source coordinates per cell. The single complete table of 19 bijective tags, domains and strict face formulas is in Subsections C.2–C.3; no cell is omitted.

B.5 Homotopy type and cell counts

Definition 51. Let $Q \subseteq \mathcal{Q}(L_2(S))$ be the subcomplex spanned by the $3f$ corner squares, i.e. by the cells $\mathsf{v}, \mathsf{E}, \mathsf{t}, \mathsf{vE}, \mathsf{Et}, \mathsf{q}$ in the tag notation of Definition 64 (Appendix C).

Lemma 52 (The corner subcomplex realises S)**.** *Q has $v + e + f$ vertices, $2e + 3f$ edges and $3f$ squares, and $|Q|$ is homeomorphic to $|S|$.*

Proof. Map the cells indexed by a vertex, edge and triangle to that vertex, its edge midpoint and triangle barycentre. The square $[u, t_T]$ maps to the quadrilateral with corners u, the two edge midpoints through u, and the barycentre of T. Thus Q is the standard cubical subdivision of S (Athanasiadis, 2012, Example 7.12). Its cells count the nonempty faces and the vertex–edge, edge–triangle and vertex–triangle incidences, giving the stated census. □

Theorem 53 (Homotopy type of the two-guard lift)**.** *Let S be a source complex and $r = v + 3e + f - 2$. Then*

$$\mathcal{Q}(L_2(S)) \;\simeq\; |S| \;\vee\; \bigvee_{r} S^1, \qquad H_1\big(\mathcal{Q}(L_2(S)); \mathbb{Z}\big) \;\cong\; H_1(S; \mathbb{Z}) \oplus \mathbb{Z}^r,$$

and the cell counts are $(c_0, c_1, c_2) = (v + 4e + f + 2,\ v + 8e + 4f,\ 3f)$.

Proof. The cells outside Q are $2 + 3e$ vertices ($\hat{0}, \hat{1}$, and the guards) and $v + 6e + f$ edges, with no higher cell. Fix $u_0 \in V$ and $T_0 \in F$. In the tag notation of Definition 64 in Appendix C, the stars

$$A = \{\mathsf{bv}(u_0)\} \cup \{\mathsf{bl}(\eta) : \eta \in E\}, \qquad B = \{\mathsf{th}(T_0)\} \cup \{\mathsf{r}_j\mathsf{h}(\eta) : \eta \in E,\ j = 1, 2\}$$

are disjoint trees, each meeting Q in its root, and absorb every outside vertex. Successively contracting them gives a homotopy equivalence fixing Q (Hatcher, 2002, Proposition 0.17). The remaining $(v + 6e + f) - (2 + 3e) = v + 3e + f - 2$ edges have both endpoints in connected Q, so each adds a free circle. Lemma 52 gives the homotopy and homology statements. The interval classification gives the cell counts. Contraction is used only in the proof; every original cell remains in the input. □

Remark 54. The free summand $\mathbb{Z}^r$ is part of $H_1(\mathcal{Q}(L_2(S)); \mathbb{Z})$, which in general differs from the torsion group being realised; the preserved torsion satisfies $\mathrm{Tor}\, H_1(\mathcal{Q}(L_2(S)); \mathbb{Z}) \cong \mathrm{Tor}\, H_1(S; \mathbb{Z})$.

B.6 The one-upper-guard variant $L_1(S)$

Proposition 55. *Assume in addition that $d_\eta \ge 2$ for every $\eta \in E$. Let $L_1(S)$ be defined as in Definition 45 but with a single upper guard ρ_η per edge, families (C6) and (C8) becoming $E_\eta \lessdot \rho_\eta \lessdot \hat{1}$. Then $L_1(S)$ is a lattice of height 4, its only rank-2 Boolean intervals are the $3f$ corner intervals $[u, t_T]$, it has no B_3, and*

$$|L_1(S)| = v + 3e + f + 2, \quad c_1 = v + 6e + 4f, \quad c_2 = 3f, \quad \mathcal{Q}(L_1(S)) \simeq |S| \vee \bigvee_{v+2e+f-2} S^1.$$

Proof. The meet proof is unchanged with one upper guard. In the rank-two table only $[E_\eta, \hat{1}]$ changes, now having $d_\eta + 1 \geq 3$ intermediates. For rank three, the first three exclusions are unchanged; $[u, \hat{1}]$ has $f(u) + d(u) \geq 4$ coatoms. Indeed $d(u) = 2$ would force the unique possible triangle through u, contradicting $d_\eta \geq 2$ on either incident edge.

Deleting one guard and its two covers per edge gives the stated census. Outside Q remain $2 + 2e$ vertices and $v + 4e + f$ edges. The same two rooted-star contractions as in Theorem 53, with one upper guard, leave $v + 2e + f - 2$ circles. □

B.7 Explicit cyclic inputs and arbitrary finite abelian torsion

Definition 56 (The cyclic source S_m)**.** Fix an integer $m \geq 2$ and put $N = 3m$. Take distinct vertices $a_0, \ldots, a_{N-1}, b_0, b_1, b_2$ and c, with a-indices read modulo N and b-indices modulo 3. The complex S_m has exactly the triangles, for $i = 0, \ldots, N-1$,

$$\{a_i, a_{i+1}, b_{i+1 \bmod 3}\}, \qquad \{a_i, b_{i \bmod 3}, b_{i+1 \bmod 3}\}, \qquad \{c, a_i, a_{i+1}\},$$

together with all their faces.

These $3N = 9m$ triples are distinct nondegenerate abstract triangles. The first two families triangulate the mapping cylinder of the simplicial degree-m map $C_{3m} \to C_3$, $a_i \mapsto b_{i \bmod 3}$, and the third cones off the source circle.

Lemma 57 (Strip and cone intersections)**.** *Let R_i be the quadrilateral strip on $\{a_i, a_{i+1}, b_j, b_{j+1}\}$ with $j = i \bmod 3$. Then adjacent strips R_i, R_{i+1}, including the wraparound pair R_{N-1}, R_0, meet exactly in the vertical edge $a_{i+1}b_{j+1}$; nonadjacent strips with equal residues modulo 3 meet exactly in their common image edge $b_j b_{j+1}$; nonadjacent strips with distinct residues meet exactly in the single common vertex of those two image edges. No two nonadjacent strips share an a-vertex. The cone triangles $\{c, a_i, a_{i+1}\}$ meet the union of the strips exactly in the source circle, and meet each other in c, or, for adjacent indices, in their common spoke.*

Proof. Distinct cyclic pairs $\{i, i+1\}$ share an a-vertex exactly for adjacent strips, including wraparound; their common b-vertex then gives the vertical edge. Nonadjacent strips meet only in their image edges in C_3, which coincide for equal residues and otherwise share one vertex. Cone–strip intersections lie in the source cycle, and cone triangles share only the apex or an adjacent spoke. Thus the declared diagonals give exactly the mapping-cylinder and cone identifications. □

Proposition 58 (Homology of S_m)**.** *$H_1(S_m; \mathbb{Z}) \cong \mathbb{Z}/m$ and $H_2(S_m; \mathbb{Z}) = 0$. Moreover S_m is connected and pure of dimension 2, with*

$$(v, e, f) = (3m + 4,\ 12m + 3,\ 9m),$$

each image edge $b_j b_{j+1}$ has triangle degree m and every other edge has degree 2.

Proof. By Lemma 57, this is the mapping cone of a degree-m circle map, with reduced CW chain model $\mathbb{Z} \xrightarrow{m} \mathbb{Z}$. Hence $H_1 = \mathbb{Z}/m$ and $H_2 = 0$; this is not the full triangulated chain matrix. There are $N + 4$ vertices and $4N + 3$ edges: N each of the source-cycle, vertical, diagonal and spoke edges, plus three image edges. Each image edge lies in m strip triangles; every other edge lies in two declared triangles. Connectedness and purity follow from this inventory. □

Theorem 59 (Cyclic realisation at height four)**.** *For every $m \geq 2$ the one-upper-guard lattice $L_m := L_1(S_m)$ has height 4 and*

$$|L_m| = 48m + 15, \quad (c_0, c_1, c_2) = (48m + 15,\ 111m + 22,\ 27m),$$

$$H_1\big(\mathcal{Q}(L_m); \mathbb{Z}\big) \cong \mathbb{Z}^{36m+8} \oplus \mathbb{Z}/m, \qquad H_2 = 0.$$

In particular $m = 3$ gives a 159-element height-four lattice with 355 covers, 81 squares and $H_1 \cong \mathbb{Z}^{116} \oplus \mathbb{Z}/3$.

Proof. Apply Proposition 55 to the census and homology of Proposition 58; its edge-degree hypothesis holds since $m \geq 2$. Substitution gives all displayed formulas. □

Theorem 60 (Arbitrary finite abelian torsion)**.** *Let A be a nontrivial finite abelian group with a cyclic decomposition $A \cong \bigoplus_{i=1}^{k} \mathbb{Z}/m_i$, $m_i \geq 2$. Then there is a finite lattice L of height 4 with*

$$\operatorname{Tor} H_1\big(\mathcal{Q}(L);\mathbb{Z}\big) \cong A, \qquad |L| = 48\sum_i m_i + 12k + 3,$$

and $H_1(\mathcal{Q}(L);\mathbb{Z}) \cong \mathbb{Z}^{36\sum_i m_i + 9k - 1} \oplus A$. The trivial group is realised at height 4 by the five-element chain.

Proof. Wedge the S_{m_i} at a single vertex, keeping all other vertices disjoint. The result is connected and pure, with unchanged edge degrees at least two, $H_1 \cong A$, and

$$v = 3\sum_i m_i + 3k + 1, \qquad e = 12\sum_i m_i + 3k, \qquad f = 9\sum_i m_i.$$

Apply Proposition 55. For the trivial group, the five-element chain has height four and its full Boolean-interval complex is an interval. □

B.8 Torsion-freeness at height at most three

Theorem 61. *Let L be a finite lattice of height at most 3. Then $H_1(\mathcal{Q}(L);\mathbb{Z})$ is free abelian.*

Proof. For height at most 1, the cube complex is a point or a single edge, and the conclusion is immediate. We may therefore assume that L has height 2 or 3. Orient all cover edges upward. If x is proper and is neither an atom nor a coatom, there are y, z with $\hat{0} < y < x < z < \hat{1}$, a chain of four strict steps, contradicting the height bound. So every proper element is an atom or a coatom. If m is both, then $\hat{0} \lessdot m$ and $m \lessdot \hat{1}$, so no other proper element is comparable with m.

Hence every cover of L is of one of the forms $\hat{0} \lessdot a$, $c \lessdot \hat{1}$, or $a \lessdot c$ with a an atom, c a coatom and $a \neq c$; write the corresponding 1-cells p_a, q_c, r_{ac}. These form a basis of C_1. Change basis by replacing each r_{ac} by

$$t_{ac} = p_a + r_{ac} + q_c,$$

and, for each element m that is both an atom and a coatom, replacing q_m by $w_m = p_m + q_m$. Each replacement changes one basis vector by a combination of others, so it is unimodular; the pivots are distinct because r_{ac} occurs only in t_{ac}, and because an atom-coatom m carries no cell $r_{m\cdot}$ or $r_{\cdot m}$ by the incomparability just proved.

Every rank-2 Boolean interval touches $\hat{0}$ or $\hat{1}$: an interval $[x, y]$ with $\hat{0} < x$ and $y < \hat{1}$ would give a chain $\hat{0} < x < z < y < \hat{1}$. The three cases have the following boundaries in the new basis:

$$\partial[\hat{0}, c] = p_a + r_{ac} - p_b - r_{bc} = t_{ac} - t_{bc}, \qquad \partial[a, \hat{1}] = r_{ac} + q_c - r_{ad} - q_d = t_{ac} - t_{ad},$$

$$\partial[\hat{0}, \hat{1}] = p_m + q_m - p_n - q_n = w_m - w_n.$$

Thus every column of the square boundary matrix is a difference of two distinct basis vectors, so in the new basis it is the incidence matrix of a graph, possibly with additional zero rows. Its image is the set of integer vectors whose coordinates sum to zero on each connected component of that graph, which is a saturated subgroup of C_1; hence $C_1/\operatorname{im}\partial_2$ is free abelian. Since $H_1 = \ker\partial_1/\operatorname{im}\partial_2$ embeds in $C_1/\operatorname{im}\partial_2$, and every subgroup of a free abelian group is free abelian, H_1 is free abelian. Cells of dimension ≥ 3, which may exist when $L \cong B_3$, do not enter the computation of H_1. □

Corollary 62. *For every nontrivial finite abelian group A, the minimum height of a finite lattice L with $\operatorname{Tor} H_1(\mathcal{Q}(L);\mathbb{Z}) \cong A$ is exactly 4.*

Proof. Theorem 61 excludes heights ≤ 3, and Theorem 60 realises A at height 4. □

C COUNTING TRANSFER FOR THE FULL GUARDED LIFT

C.1 PRECISE STATEMENT

Let S, T be finite connected pure simplicial 2-complexes. The source $I(S)$ contains all nonempty faces with dimension and strict containment, including vertex–triangle containment, and no names, order or orientation. The target $J(S)$ contains all cells of $\mathcal{Q}(L_2(S))$, geometric dimension, strict Face, and the stipulated finite construction roles. For $J^+(S)$ add β and one edge to every old vertex, with two new roles and no new higher cells.

Here C^r uses at most r symbols and ordinary unary counting at fixed integer thresholds; formulas are finite with unrestricted length and quantifier rank.

Theorem 63. *For every $k \geq 2$ and every C^k sentence φ over the signature of J there is a uniform source C^{2k} sentence φ^* such that $J(S) \models \varphi$ if and only if $I(S) \models \varphi^*$; the same holds for the signature of J^+. Consequently*

$$I(S) \equiv_{C^{2k}} I(T) \quad \Longrightarrow \quad J(S) \equiv_{C^k} J(T), \qquad J^+(S) \equiv_{C^k} J^+(T).$$

Uniformity means dependence only on the formula and fixed role table, not on source size or a chosen source element.

C.2 REPETITION-FREE FINITE TAG REPRESENTATION

Here $V(u)$, $E(\eta)$ and $F(\Delta)$ abbreviate the unary dimension predicates Dim_0, Dim_1 and Dim_2 of $I(S)$ (Appendix A), and $U(S)$ is the universe of $I(S)$. Use the following incidence predicates:

$$D_{VE}(u,\eta) = V(u) \wedge E(\eta) \wedge \mathrm{Sub}(u,\eta), \qquad D_{EF}(\eta,\Delta) = E(\eta) \wedge F(\Delta) \wedge \mathrm{Sub}(\eta,\Delta),$$
$$D_{VF}(u,\Delta) = V(u) \wedge F(\Delta) \wedge \mathrm{Sub}(u,\Delta).$$

Definition 64. A *tag* is a symbol τ with an arity $a_\tau \in \{0,1,2\}$ and a source domain formula D_τ with a_τ free variables; $D_\tau(S) \subseteq U(S)^{a_\tau}$ denotes its set of solutions in $I(S)$, with the convention $U^0 = \{()\}$. The *tagged carrier* of S is the disjoint union $\bigsqcup_{\tau\in\mathcal{T}} \{\tau\} \times D_\tau(S)$, and ι denotes the map sending $(\tau, \bar{a})$ to the cell listed in the last column of the tables below.

Tags are finite indices of the translation, not source constants or additional observable colours. The following tables give every cell from the classification in Appendix B.

0-cell tag	source coordinates	domain	corresponding lattice element
b	$()$	true	bottom
h	$()$	true	top
v	(u)	$V(u)$	u
l	(η)	$E(\eta)$	λ_η
E	(η)	$E(\eta)$	E_η
t	(Δ)	$F(\Delta)$	t_Δ
r_1, r_2	(η)	$E(\eta)$	ρ^1_η, ρ^2_η

This table has 8 tags in total. Bottom and top each carry the unique empty tuple and are distinguished by their tags alone; no source point has been selected to represent them.

1-cell tag	source coordinates	domain	the two endpoints
bv	(u)	$V(u)$	$\mathsf{b}()$, $\mathsf{v}(u)$
bl	(η)	$E(\eta)$	$\mathsf{b}()$, $\mathsf{l}(\eta)$
vE	(u,η)	D_{VE}	$\mathsf{v}(u)$, $\mathsf{E}(\eta)$
lE	(η)	$E(\eta)$	$\mathsf{l}(\eta)$, $\mathsf{E}(\eta)$
Et	(η,Δ)	D_{EF}	$\mathsf{E}(\eta)$, $\mathsf{t}(\Delta)$
Er_1, Er_2	(η)	$E(\eta)$	$\mathsf{E}(\eta)$, $\mathsf{r}_1(\eta)$ resp. $\mathsf{r}_2(\eta)$
th	(Δ)	$F(\Delta)$	$\mathsf{t}(\Delta)$, $\mathsf{h}()$
$\mathsf{r}_1\mathsf{h}$, $\mathsf{r}_2\mathsf{h}$	(η)	$E(\eta)$	$\mathsf{r}_1(\eta)$ resp. $\mathsf{r}_2(\eta)$, $\mathsf{h}()$

This table has 10 tags in total. The unique 2-cell tag is q, of arity 2 with coordinates (u,Δ) and domain D_{VF}, representing the corner square $[u, t_\Delta]$. Thus $\mathcal{T}$ consists of 19 tags altogether.

Lemma 65. *The map ι is a bijection from the tagged carrier of S onto the universe of $J(S)$, and it matches cell dimension and construction role with the tag.*

Proof. The interval classification exhausts the rows. Each row's cell recovers its source coordinates: the relevant source face, the endpoints of a cover, or the endpoints (u, t_Δ) of a square. Distinct tags name different cell families. Thus the representation is bijective, with dimensions and stipulated roles fixed by each tag. □

The role of each coordinate is fixed by its position, so no orientation of edges or triangles and no choice of representative is required anywhere. Even when the lattice rank of a 0-cell is 2, 3 or 4, its cell dimension is still 0; rank and dimension are different data.

C.3 Complete translation of the atomic formulas

Fix a tag for each variable of an atom. On legal representations, equality between distinct tags is identically false; equality within one tag is coordinatewise equality of the source variables, and equality of empty coordinate tuples is identically true. Each unary predicate of the target signature — the dimension predicates and the stipulated construction roles — is a union of tag classes, so under a fixed tag it translates to the domain formula D_τ if the tag lies in the class, and to falsity otherwise; proof tags are not required to be observable roles of the target. For a 0-cell and a 1-cell, Face holds if and only if the former equals one of the two endpoints listed in the 1-cell table; the endpoints are merely tags together with coordinate projections, and no source function and no auxiliary variable is introduced.

Fix $\mathsf{q}(u, \Delta)$. The remaining strict face relations are exactly those in the following table, where the displayed parameters are understood to satisfy their respective domain formulas.

lower-face tag	condition for being a face of $\mathsf{q}(u, \Delta)$
$\mathsf{v}(w)$	$w = u$
$\mathsf{E}(\eta)$	$\mathrm{Sub}(u, \eta) \wedge \mathrm{Sub}(\eta, \Delta)$
$\mathsf{t}(\Delta')$	$\Delta' = \Delta$
$\mathsf{vE}(w, \eta)$	$w = u \wedge \mathrm{Sub}(\eta, \Delta)$
$\mathsf{Et}(\eta, \Delta')$	$\Delta' = \Delta \wedge \mathrm{Sub}(u, \eta)$

Lemma 66. *Under any fixed assignment of tags to its variables, every atomic formula over the signature of J is equivalent, along ι, to a source formula whose free variables lie among the coordinates of its own arguments; at most four source symbols occur.*

Proof. The square table lists its four vertices and four edges. All unlisted, equal-dimensional and reversed face combinations are false. Equality and unary predicates are tag tests, so each atom uses only its arguments' coordinates, at most four symbols. □

C.4 The broadcast augmentation

Definition 67. Add the arity-0 tag z, with domain true, representing β. For each of the eight old 0-cell tags τ add a tag $\mathsf{g}\tau$ with the same arity, the same coordinates and the same domain as τ, representing the new edge from β to the cell τ. This adds $1 + 8 = 9$ tags, so $\mathcal{T}$ has 28 tags for J^+; the maximal arity is still 2.

Here z is one new 0-cell role and all eight tags $\mathsf{g}\tau$ satisfy one and the same new broadcast-edge role: the eight internal representations do not add eight observable colours to the target. The cells $\mathsf{gb}()$ and $\mathsf{gh}()$ are distinct new edges; the length of the empty tuple is a feature of the representation and not the cell dimension.

Lemma 68. *Lemmas 65 and 66 hold verbatim for J^+ with the 28 tags of Definition 67.*

Proof. The nine new tags bijectively name β and the broadcast edges. The only new faces of $\mathsf{g}\tau(\bar{a})$ are $\mathsf{z}()$ and $\tau(\bar{a})$, tested by tag and coordinate equality. All broadcast edges have the same role, and no square acquires a new face. □

C.5 Elimination lemma for pair counting

Definition 69. Let $\theta(p, q, \bar{z})$ be a counting formula in which p and q are distinct symbols not occurring in $\bar{z}$, and let $n \geq 1$. Let $\mathcal{A}_n$ be the set of all n-tuples of nonnegative integers $m = (m_1, \ldots, m_n)$ with every entry at most n and with sum n. Put

$$\mathrm{Pair}^{p,q}_{\geq n}[\theta] = \bigvee_{m \in \mathcal{A}_n} \bigwedge_{j\,:\,m_j > 0} \exists^{\geq m_j} p\, \exists^{\geq j} q\, \theta(p, q, \bar{z}),$$

and let $\mathrm{Pair}^{p,q}_{\geq 0}[\theta]$ be true.

Lemma 70. *Fix a structure and an interpretation of the parameters $\bar{z}$. Then $\mathrm{Pair}^{p,q}_{\geq n}[\theta]$ holds if and only if there are at least n ordered pairs (a, b) with $\theta(a, b, \bar{z})$. The formula is finite and uses only p, q and the symbols already occurring in θ.*

Proof. This is exactly Lemma 36 at $d = 2$, with $(u_1, u_2) = (p, q)$, including its zero-threshold case. It uses two source symbols. □

The integers n, j and m_j are fixed thresholds; only the $d \leq 2$ cases of the counting lemma are used.

C.6 Cross-tag counting and the induction over all formulas

For the target symbols $x_1, \ldots, x_k$ fix pairwise disjoint blocks (p_i, q_i) of source symbols, $2k$ symbols in all. For a formula ψ, a *tag assignment* α assigns a tag to each $x_i \in \mathrm{FV}(\psi)$. The active coordinate set is defined as follows:

$$A_\alpha(\psi) = \bigcup_{x_i \in \mathrm{FV}(\psi)} \big(\{p_i : a_{\alpha(i)} \geq 1\} \cup \{q_i : a_{\alpha(i)} = 2\}\big).$$

An assignment $\bar{a}$ to these source coordinates is *legal* if each active block satisfies $D_{\alpha(i)}$; for an arity-0 tag the block is the empty tuple. The notation $\iota(\alpha, \bar{a})$ denotes the induced assignment to the free target variables. A sentence has the unique empty tag assignment.

Proposition 71. *For every C^k formula ψ over the signature of J, and every tag assignment α on $\mathrm{FV}(\psi)$, there is a source formula ψ^α such that:*

(i) all symbols occurring in ψ^α, free or bound, belong to $\{p_1, q_1, \ldots, p_k, q_k\}$;

(ii) $\mathrm{FV}(\psi^\alpha) \subseteq A_\alpha(\psi)$;

(iii) for every S and every $\bar{a}$ legal for α, $J(S) \models \psi[\iota(\alpha, \bar{a})]$ if and only if $I(S) \models \psi^\alpha[\bar{a}]$.

The same holds for J^+ with the 28 tags. The translation is indexed by α, and (ii) permits a proper subset of the active free coordinates.

Proof. Induct on ψ, using the indicated active coordinates without padding. Atoms follow from Lemmas 66 and 68; Boolean connectives restrict α to each operand's free variables.

For $\exists^{\geq n} x_i\, \psi$, put $\theta_\tau = D_\tau \wedge \psi^{\alpha_\tau}$, where $\alpha_\tau = \alpha[i \mapsto \tau]$ is restricted to $\mathrm{FV}(\psi)$. Define the following auxiliary formulas:

$$K_{\tau,m} = \begin{cases} \mathrm{true}, & m = 0, \\ \theta_\tau, & a_\tau = 0, m = 1, \\ \mathrm{false}, & a_\tau = 0, m \geq 2, \\ \exists^{\geq m} p_i\, \theta_\tau, & a_\tau = 1, m \geq 1, \\ \mathrm{Pair}^{p_i,q_i}_{\geq m}[\theta_\tau], & a_\tau = 2, m \geq 1. \end{cases}$$

These are ordinary source formulas with domain guards as conjuncts. Each tests whether tag τ contributes at least m witnesses; arity zero contributes 0 or 1 through the unique empty tuple. Use the following formula to translate the quantifier:

$$\bigvee_{\substack{(n_\tau) \in \{0,\ldots,n\}^{\mathcal{T}} \\ \sum_\tau n_\tau = n}} \bigwedge_{\tau \in \mathcal{T}} K_{\tau, n_\tau}.$$

Bijectivity makes the total witness count $\sum_\tau N_\tau$. This allocation exists exactly when the sum is at least n, proving truth preservation, including $n = 0$ and $x_i \notin \mathrm{FV}(\psi)$.

Only the block (p_i, q_i) is rebound. Nested rebinding shadows that same block, other blocks are disjoint, and lower-arity tags introduce no unused free coordinate. Thus all symbols stay within the $2k$ pool and free symbols form the claimed subset. True and false, if not primitive, can be local closed formulas over p_1. □

Proof of Theorem 63. Apply Proposition 71 with the empty tag assignment. The resulting sentence uses at most $2k$ symbols and preserves truth, hence transfers C^{2k} equivalence to C^k equivalence. The fixed tables make this uniform for both J and J^+. □

Remark 72. The factor 2 gives the sufficient variable budget for this forward translation.

D Model upper bounds and coefficient-depth transfer

D.1 Precise neighbourhoods and semantic counting lemmas

Definition 73 (Input signature)**.** Fix finite $d, \ell \in \mathbb{N}$. A finite ranked incidence structure K has a finite universe X, unary rank predicates $D_0, \dots, D_d$ partitioning X, unary bit-colour predicates $P_1, \dots, P_\ell$, equality, and one binary relation Face; we write $\mathrm{rk}(x) = a$ for $D_a(x)$. In geometric instances $\mathrm{Face}(y, x)$ means that y is a strict geometric face of x, all cells and all strict full faces are present, and the rank of a cell is its geometric cell dimension, never a lattice-order rank; the logical statements below are proved for arbitrary finite relational data over this signature. The rank and colour alphabets are fixed and shared between any two compared inputs. There are no individual names, no source orders, no orientations, and no spectral or filtration attributes. The universe may be empty.

Definition 74 (Derived neighbourhoods)**.** For $x \in X$ put

$$B(x) = \{y : \mathrm{Face}(y, x),\ \mathrm{rk}(y) = \mathrm{rk}(x) - 1\}, \qquad C(x) = \{y : \mathrm{Face}(x, y),\ \mathrm{rk}(y) = \mathrm{rk}(x) + 1\},$$

$$\mathrm{Down}(x) = \{y : \mathrm{rk}(y) = \mathrm{rk}(x),\ B(x) \cap B(y) \neq \emptyset\},$$
$$\mathrm{Up}(x) = \{y : \mathrm{rk}(y) = \mathrm{rk}(x),\ C(x) \cap C(y) \neq \emptyset\}.$$

For a neighbourhood operator N we write $E_N(x, y)$ for "$y \in N(x)$". There is no side condition $y \neq x$: a cell is its own Down- and Up-neighbour whenever the corresponding intersection is nonempty. Every distinct mediator occurrence is counted separately wherever mediators are counted below.

These operators are definable with very few names. With finite rank disjunctions set $B(x, y) := \mathrm{Face}(y, x) \wedge \bigvee_{a=1}^{d}(D_a(x) \wedge D_{a-1}(y))$, $C(x, y) := B(y, x)$, $\mathrm{EqRank}(x, y) := \bigvee_{a=0}^{d}(D_a(x) \wedge D_a(y))$, and

$$L(x, y) := \mathrm{EqRank}(x, y) \wedge \exists w\,(B(x, w) \wedge B(y, w)),$$
$$U(x, y) := \mathrm{EqRank}(x, y) \wedge \exists w\,(C(x, w) \wedge C(y, w)).$$

Thus L defines E_{Down} and U defines E_{Up}, with the convention "y is a neighbour of the centre x". The formulas B, C, EqRank are quantifier free; L, U use a single additional bound name w, and both may hold with $x = y$. If $d = 0$ the disjunction defining B is empty, hence false, so there are no adjacent-rank boundary neighbours; in geometric rank-0 instances there are no strict faces, while arbitrary relational data may still have Face pairs. No long-range relation, orientation, or unmentioned exclusion enters.

Definition 75 (The two logics)**.** For $m \geq 2$, C^m denotes first-order logic with equality, the predicates of Definition 73 and unary counting quantifiers $\exists^{\geq n} x\, \varphi$ for fixed natural thresholds n, using at most m distinct variable symbols, with arbitrary finite formula length and quantifier rank; see Grohe & Pakusa (2017). The language TC^m is the semantic pair-count language of Akbari et al. (2026) over the derived neighbourhoods: its atoms are $D_a(x_i)$, $P_c(x_i)$, $x_i = x_j$ and $E_N(x_i, x_j)$ for $N \in \{B, C, \mathrm{Down}, \mathrm{Up}\}$; it is closed under Boolean operations; and for distinct indices $i \neq j$ and $N \geq 0$ it has the pair quantifier $Q^{\geq N}_{(x_i, x_j)}\psi$, true under an assignment iff $|\{(\alpha, \beta) \in X^2 : K \models \psi[x_i \mapsto \alpha, x_j \mapsto \beta]\}| \geq N$. The two bound symbols must be distinct

symbols, but the counted pairs range over all of X^2, including pairs with $\alpha = \beta$ and pairs sharing a value with another free coordinate. Unary counting is *not* a primitive of TC^m. The constants $\top, \bot$ are width-preserving abbreviations; for $m \geq 2$ the closed formula $Q^{\geq 0}_{(x_1,x_2)}(x_1 = x_1)$ is such a $\top$.

Lemma 76 (Ordered tuple counting)**.** *For $t \geq 0$ fix a designated ordered block of distinct counted names $(u_1, \ldots, u_t)$, disjoint from the remaining free parameters $\bar{z}$ of θ; the block need not occur free in θ and is part of the operator although suppressed in the notation. Put $\mathrm{Count}^t_{\geq 0} := \top$ for every arity, $\mathrm{Count}^0_{\geq 1}[\theta] := \theta$, $\mathrm{Count}^0_{\geq n} := \bot$ for $n \geq 2$, $\mathrm{Count}^1_{\geq n}[\theta] := \exists^{\geq n} u_1\, \theta$, and for $t \geq 2$, $n \geq 1$,*

$$\mathrm{Count}^t_{\geq n}[\theta] \;:=\; \bigvee_{\substack{(a_1,\ldots,a_n)\in\{0,\ldots,n\}^n \\ a_1+\cdots+a_n=n}} \;\bigwedge_{j\,:\,a_j>0} \exists^{\geq a_j} u_1\, \mathrm{Count}^{t-1}_{\geq j}[\theta],$$

where the designated block of the recursive occurrence is $(u_2, \ldots, u_t)$ and u_1 is held fixed. Then $\mathrm{Count}^t_{\geq n}[\theta]$ holds exactly when at least n ordered tuples $(\alpha_1, \ldots, \alpha_t) \in X^t$ satisfy θ. For fixed t, n this is a single finite formula using only the designated names and the symbols of θ, which Definition 75 permits.

Proof. Induct on t. The cases $t = 0, 1$ follow from the unique empty tuple and ordinary unary counting; threshold zero is always true. For $t \geq 2$ and $n \geq 1$, fix the parameters and let c_α count the ordered $(t-1)$-tuples satisfying θ with $u_1 = \alpha$. By induction the inner count defines the level set of size $A_j = |\{\alpha : c_\alpha \geq j\}|$. A branch succeeds exactly when an allocation $a_j \leq A_j$, $\sum_{j=1}^n a_j = n$, exists, equivalently when $\sum_{j=1}^n A_j \geq n$. The layer identity

$$\sum_{j=1}^{n} A_j = \sum_{\alpha} \min(c_\alpha, n)$$

makes this equivalent to $\sum_\alpha c_\alpha \geq n$, the required tuple count. The level sets need not be disjoint, and coordinate values may repeat. Empty universes give no positive-arity witnesses. Names absent from the free variables of θ are still counted over the designated block. Each branch uses only that block and the symbols already in θ; its internal binders retain their scopes, and the surrounding counting binders affect only free occurrences of the designated names. Fixed t, n give a finite formula. □

We use $\mathrm{Pair}^{u,v}_{\geq n}[\theta] := \mathrm{Count}^2_{\geq n}[\theta]$ for the designated block (u, v), so $\mathrm{Pair}^{u,v}_{\geq 0}[\theta] = \top$.

Lemma 77 (Exact guarded multiset)**.** *Let $\mathcal{L}$ be a logic closed under Boolean combinations and equipped, for a fixed block $\bar{w}$ of counted names and every threshold n, with a construct $\mathsf{C}_{\geq n}\bar{w}\,[\vartheta]$ expressing that at least n assignments of $\bar{w}$ satisfy ϑ. Let $\Gamma(\bar{x}, \bar{w}) \in \mathcal{L}$ be a guard and let $(\Phi_a)_{a\in S}$ be finitely many formulas of $\mathcal{L}$ that are pairwise contradictory, so that each assignment of $\bar{w}$ has at most one* class *$a \in S$. For a finite multiset M with support S and multiplicities n_a, use the following formula:*

$$\mathrm{Ex}_{\Gamma,M} \;:=\; \bigwedge_{a\in S}\Big(\mathsf{C}_{\geq n_a}\bar{w}[\Gamma \wedge \Phi_a] \;\wedge\; \neg\,\mathsf{C}_{\geq n_a+1}\bar{w}[\Gamma \wedge \Phi_a]\Big) \;\wedge\; \neg\,\mathsf{C}_{\geq 1}\bar{w}\Big[\Gamma \wedge \neg\bigvee_{a\in S}\Phi_a\Big].$$

Then $\mathrm{Ex}_{\Gamma,M}$ holds at $\bar{x}$ if and only if the multiset of classes of the assignments of $\bar{w}$ satisfying Γ is exactly M. It uses no name outside those of Γ, of the Φ_a and of $\bar{w}$.

Proof. The Φ_a are pairwise contradictory, so the guarded assignments are partitioned into the listed classes together with the unlisted remainder. The two-sided threshold pair for a says exactly that the class-a part has cardinality n_a, and the final conjunct says that the remainder is empty; conversely both hold when the multiset is M. If $S = \emptyset$ the empty conjunction is true and the final conjunct states that no assignment satisfies Γ. □

D.2 From the pair-count language to ordinary counting logic

Theorem 78 (Width transfer for the pair-count language)**.** *Let $m \geq 2$. For every TC^m formula φ there is a uniformly constructed ordinary C^{m+1} formula φ^* over the* Face *signature such that for every finite input K and every assignment of the free variables of φ, $K \models \varphi$ iff $K \models \varphi^*$. The*

free variables of φ^* are contained in those of φ; no further free coordinate is introduced and no numerical oracle is used.

Proof. Reserve the original names $x_1, \dots, x_m$ and one fresh name w, and translate by structural induction. Rank, colour and equality atoms translate to themselves. The translations $E_B \mapsto B$ and $E_C \mapsto C$ are quantifier free. The translations $E_{\mathrm{Down}} \mapsto L$, $E_{\mathrm{Up}} \mapsto U$ each bind w within the atom they expand. Every occurrence of w is therefore locally bound; it is neither an original bound name nor an assignment coordinate. Boolean steps are recursive. Set $(Q^{\geq N}_{(x_i,x_j)}\psi)^* := \mathrm{Pair}^{x_i,x_j}_{\geq N}[\psi^*]$. The condition $i \neq j$ supplies distinct symbols for the two counting blocks. The threshold $N = 0$ yields $\top$, in agreement with the semantics; this may delete syntactic occurrences of variables, whence containment rather than equality of free-variable sets, but never creates a free variable. Nested pair quantifiers may rebind names already in use: the inner blocks $\exists^{\geq a} x_i\, \exists^{\geq j} x_j$ shadow the outer values exactly inside the scope in which the pair-quantifier semantics also shadows them, and the helper scopes are strictly local to single atoms, so nothing is captured. No symbol outside $\{x_1, \dots, x_m, w\}$ is introduced, so $\varphi^* \in C^{m+1}$. On an empty universe every $\exists^{\geq a}$ with $a \geq 1$ is false, thresholds 0 are true, and the equivalence holds vacuously. □

D.3 Exact-width preservation of the anchor augmentation

Definition 79 (Anchor augmentation)**.** $\mathrm{Aug}(K)$ contains all cells of K, one new rank-0 cell β, and one new rank-1 cell e_v for each old rank-0 cell v. The new strict-face pairs are exactly (β, e_v) and (v, e_v); all old relations are unchanged. Old bit colours are preserved and the vocabulary is extended by fixed fresh tags, shared between compared inputs, marking β and the star edges. There are no new rank-2 cells, no other new relations and no per-cell unique colours. For cubical inputs this is the usual geometric star augmentation.

In the tag representation, $\mathrm{Aug}(K)$ is the broadcast augmentation of Definition 67 in Appendix C. Represent $\mathrm{Aug}(K)$ inside K by three disjoint tags, in the free-tag-context style of Proposition 71: $\mathsf{O}(u)$ with domain $\top$ for the old cell u; $\mathsf{A}(\,)$, of arity zero and domain $\top$, for the anchor β; and $\mathsf{S}(u)$ with domain $D_0(u)$ for the star edge e_u.

Lemma 80 (Faithful tagged representation)**.** *This assignment is a bijection onto the universe of* $\mathrm{Aug}(K)$*, and all* $\mathrm{Aug}(K)$ *relations are defined on legal representations by*

$$\begin{aligned} \mathrm{Face}(\mathsf{O}(u), \mathsf{O}(v)) &\iff \mathrm{Face}(u, v), \\ \mathrm{Face}(\mathsf{A}(\,), \mathsf{S}(v)) &\iff \top, \\ \mathrm{Face}(\mathsf{O}(u), \mathsf{S}(v)) &\iff u = v, \end{aligned}$$

with all remaining tag combinations false. Equality of two objects with the same tag is coordinate-wise equality, equality of two A*-objects is true, and objects with different tags are never equal. Rank is the old rank on* O*, is* 0 *on* A *and is* 1 *on* S*; old colour predicates are inherited on* O*, take fixed constant values on the new tags, and the fresh marks are tag tests.*

Proof. The A tag has arity zero and denotes exactly one object, a genuine empty tuple, so β is represented once even when $X = \emptyset$; S, restricted to D_0, bijects with the old rank-0 cells, and O with the old cells. Distinct tags denote disjoint classes, so the map is a bijection with no duplicated representation. The new strict-face pairs of Definition 79 are exactly (β, e_v), giving the second clause, and (v, e_v) with $\mathrm{rk}(v) = 0$, giving the third; there the domain condition on the S coordinate already forces $D_0(v)$, so the equality test correctly restricts the old endpoint to rank 0. For ranks, if $d = 0$ the augmented alphabet must include D_1, false on all old objects; in general one uses ranks up to $\max(d, 1)$. Every displayed formula uses only the coordinates supplied by the tags. □

Theorem 81 (Augmentation preserves counting width)**.** *Let $m \geq 2$. If $K \equiv_{C^m} L$ then* $\mathrm{Aug}(K) \equiv_{C^m} \mathrm{Aug}(L)$*, in their fixed matching extended colour vocabularies.*

Proof. Reserve one source name u_i for each target symbol x_i, $1 \leq i \leq m$. Translate a target formula φ relative to a *tag context* τ. It assigns a tag in $\{\mathsf{O}, \mathsf{A}, \mathsf{S}\}$ to each variable of $\mathrm{FV}(\varphi)$; it need not assign all m symbols. Atoms translate by Lemma 80, using the coordinate u_i exactly when $\tau(x_i) \in \{\mathsf{O}, \mathsf{S}\}$ and no coordinate when $\tau(x_i) = \mathsf{A}$; Boolean steps restrict the context to the

free variables of each immediate subformula. For $\exists^{\geq n} x_i\, \psi$ under τ, the context domain need not contain x_i. Form τ_σ by extending τ with $x_i \mapsto \sigma$ and restricting to $\mathrm{FV}(\psi)$. This also handles a quantified name unused in ψ. Put $\theta_\sigma := \psi^{\tau_\sigma}$. By disjointness of the tags the witness count splits as $n_{\mathsf{O}} + n_{\mathsf{A}} + n_{\mathsf{S}}$, with n_{O} counting old cells satisfying θ_{O}, n_{S} rank-0 old cells satisfying θ_{S}, and $n_{\mathsf{A}} \in \{0, 1\}$ according to whether θ_{A} holds; hence

$$\big(\exists^{\geq n} x_i\, \psi\big)^{\tau} \;:=\; \bigvee_{\substack{n_{\mathsf{O}}, n_{\mathsf{S}} \in \{0,\ldots,n\},\ n_{\mathsf{A}} \in \{0,1\} \\ n_{\mathsf{O}} + n_{\mathsf{A}} + n_{\mathsf{S}} = n}} \Big(\exists^{\geq n_{\mathsf{O}}} u_i\, \theta_{\mathsf{O}} \;\wedge\; \big[\theta_{\mathsf{A}}\big]_{n_{\mathsf{A}}} \;\wedge\; \exists^{\geq n_{\mathsf{S}}} u_i\, \big(D_0(u_i) \wedge \theta_{\mathsf{S}}\big) \Big),$$

where $[\theta_{\mathsf{A}}]_1 := \theta_{\mathsf{A}}$, $[\theta_{\mathsf{A}}]_0 := \top$, and a threshold 0 on a unary block is replaced by $\top$. If the actual counts are $N_{\mathsf{O}}, N_{\mathsf{A}}, N_{\mathsf{S}}$ with $N_{\mathsf{A}} \leq 1$, their sum is at least n exactly when an allocation bounded by these three capacities exists, and the conjunction enforces exactly those bounds, so the translation is exact. Only u_i is ever bound for x_i, and rebinding x_i inside ψ rebinds u_i only; the zero-arity branch keeps no unused free coordinate, so at most m source names occur and the translation is in C^m. A target sentence has empty tag context, so every C^m sentence about $\mathrm{Aug}(K)$ pulls back to a C^m sentence about K, whose truth value is shared with L, and the same pullback computes its value on $\mathrm{Aug}(L)$. If K is empty the translation still produces exactly the one-element structure $\{\beta\}$, without choosing a dummy source element. □

D.4 Characteristic formulas for the paired refinement

Definition 82 (Paired recurrence, $k \geq 2$)**.** For $k \geq 2$ we use Definition 2 verbatim. Coordinate values and α, β may repeat; the coding is injective and shared across compared inputs, so colours are shared recursive records rather than unrelated per-input integers; two inputs are compared by the multiset of $\chi^t_k(\bar{u})$ over $\bar{u} \in X^k$.

Proposition 83 (Characteristic formulas)**.** *Let $k \geq 2$. For every finite round t and every colour γ realized in either of two compared inputs there is a TC^{k+2} formula Φ^t_γ, with free variables among the k tuple names $x_1, \ldots, x_k$, such that $\chi^t_k(\bar{u}) = \gamma$ iff $\bar{u}$ satisfies Φ^t_γ.*

Proof. At stage 0 the atomic type is a finite quantifier-free conjunction on the k names. The alphabets are fixed and finite. Assume the claim at round t. A round-$(t+1)$ record is $\Gamma = (\gamma, M)$ with M the context multiset of Definition 82; a context $P \in \mathrm{supp}\, M$ consists of an extended atomic type on $k+2$ positions together with round-t colours $\gamma_{i,\alpha}, \gamma_{i,\beta}$ for each i. With a, b the two spare names of the fixed pool $\{x_1, \ldots, x_k, a, b\}$ set

$$\Psi_P(\bar{x}, a, b) \;:=\; \mathrm{atp}^P_{k+2}(\bar{x}, a, b) \;\wedge\; \bigwedge_{i=1}^{k} \Big(\Phi^t_{\gamma_{i,\alpha}}\big(\bar{x}[a/i]\big) \;\wedge\; \Phi^t_{\gamma_{i,\beta}}\big(\bar{x}[b/i]\big) \Big),$$

where each shifted formula uses a capture-safe bijection of the whole variable pool. It sends the designated block of k tuple names to $x_1, \ldots, x_{i-1}, a, x_{i+1}, \ldots, x_k$ (respectively with b) and acts on bound occurrences as well. Thus x_i and the other spare name may occur bound, but not free. Internal rebinding of the overwritten x_i is harmless: x_i is absent from the free arguments, and its scope does not reach neighboring conjuncts. Distinct contexts have contradictory Ψ_P: extended atomic types are mutually exclusive and each shifted round-t colour is pinned down by exactly one characteristic formula, so every $(\alpha, \beta) \in X^2$ realizes exactly one context. Now apply Lemma 77 with counted block $\bar{w} = (a, b)$, construct $\mathsf{C}_{\geq n} = Q^{\geq n}_{(a,b)}$, guard $\Gamma \equiv \top$ (all ordered pairs, diagonal included) and classes $(\Psi_P)_{P \in \mathrm{supp}\, M}$ with multiplicities n_P, and conjoin the old colour: $\Phi^{t+1}_{(\gamma,M)}(\bar{x}) := \Phi^t_\gamma(\bar{x}) \wedge \mathrm{Ex}_{\top,M}$, including the guarded exclusion of unlisted contexts. This defines exactly the tuples of old colour γ and context multiset M and lies in TC^{k+2}. □

By Theorem 78 with $m = k + 2 \geq 4$, each Φ^t_γ has a truth-preserving ordinary C^{k+3} translation over the Face signature, with free variables contained in the same k tuple names.

D.5 Global multiplicities, stability and selected readouts

Theorem 84 (Histogram transfer, $k \geq 2$)**.** *Let $k \geq 2$ and let $K \equiv_{C^{k+3}} L$ be finite inputs over the shared signature. Then for every finite round t the multisets $\{\{\chi^t_k(\bar{u}) : \bar{u} \in X^k_K\}\}$ and $\{\{\chi^t_k(\bar{u}) : \bar{u} \in X^k_L\}\}$ are equal.*

Proof. Fix t and a colour γ realized in either input. By Proposition 83 and Theorem 78 there is an ordinary C^{k+3} formula Φ with free variables among $x_1, \ldots, x_k$ defining exactly the tuples of colour γ. Apply Lemma 76 with tuple length k, taking the k tuple names as the designated counted block; the at most three remaining symbols of Φ, namely the two spare pair names and the adjacency helper, occur only bound inside Φ and are untouched by the counting blocks, which reuse $x_1, \ldots, x_k$. Hence every sentence $\mathrm{Count}^k_{\geq n}[\Phi]$ lies in C^{k+3} and has the same truth value on K and L. Since both universes are finite, comparison of all thresholds determines the exact number of ordered k-tuples of colour γ. Only finitely many colours are realized. □

Remark 85 (Stability)**.** The refinement partitions of the finite tuple sets induced by χ^t_k under the shared canonical coding are non-decreasing in t and stabilize after finitely many rounds. Theorem 84 holds at every finite round, hence at any common stable stage. Stability means stabilization of these shared partitions, not literal equality of nested hash strings.

Remark 86 (Selected readouts)**.** Every colour retains its initial atomic type through its old-colour component, so any fixed rank pattern or coordinate-equality pattern selects a union of colour classes. Equality of the complete colour multiplicities therefore implies equality of the selected multiplicities, including rank-zero pooling and pooling over diagonal tuples. Nothing in the converse direction follows from a single selected match.

D.6 The rank-aware local recurrence and augmented profiles

Definition 87 (Local recurrence, $k = 1$)**.** For $k = 1$ we use Definition 1 with the CCWL initialization $\chi^0_1(x) := (\mathrm{rk}(x), \mathrm{color}(x))$ (Akbari et al., 2026, Appendix B, p. 14). Forgetting rank gives the colour-only coarsening. The global signature is the multiset over all cells; diagonal neighbours $y = x$ are included and every mediator occurrence is counted.

Proposition 88 (Local characteristic formulas)**.** *For every finite round t and every realized cell colour c there is a TC^3 formula $\Theta^t_c(x)$ with free variable x defining exactly the cells of colour c at round t.*

Proof. Use the three names x, y, z with x the free centre. At round 0 take the conjunction of the rank atom and the bit-colour literals of c. Assume the claim at round t. Relocate Θ^t_c to y or z by a capture-safe bijection of the whole three-name pool, including bound occurrences. A relocated subformula may locally rebind x: x is absent from its free arguments, and the binding does not extend to the outside guards. Each of the four channels is a guarded pair channel on the block (y, z). A boundary neighbour of colour c contributes exactly one ordered pair, the diagonal one, under the guard $y = z \wedge B(x, y)$ with class formula $\Theta^t_c(y)$; the coboundary channel uses C in place of B. The lower mediated channel has guard $\mathrm{EqRank}(x, y) \wedge B(x, z) \wedge B(y, z)$ and classes $\Theta^t_c(y) \wedge \Theta^t_d(z)$ indexed by pairs (c, d). This guard is equivalent to $L(x, y) \wedge B(x, z) \wedge B(y, z)$, because z itself witnesses the existential quantifier in L, and it uses only the three names. The upper channel uses EqRank, C, C in the same way. Apply Lemma 77 once per channel, with counted block (y, z), construct $Q^{\geq n}_{(y,z)}$, that channel's guard and its listed classes; the exclusion of unlisted types is thereby guarded to its own channel and places no constraint on arbitrary non-neighbour pairs. Conjoining the four resulting formulas with the previous centre-colour formula $\Theta^t_{c_0}(x)$ yields a TC^3 formula defining exactly the next colour. □

Theorem 89 (Local histogram transfer)**.** *If $K \equiv_{C^4} L$ then for every finite round t the multisets of χ^t_1 over all cells of K and of L coincide. The same holds for the colour-only initialization.*

Proof. By Proposition 88 and Theorem 78 with $m = 3$, each realized colour has an ordinary C^4 formula in the names x, y, z, w with x free. Unary counting $\exists^{\geq n} x$ reuses x, so every threshold sentence stays in C^4 and has the same value on K and L; comparing all thresholds gives equal multiplicities for each of the finitely many colours. Forgetting the rank at initialization and running the same four updates is a coarsening: an induction on layers produces a common map sending each rank-aware colour to the corresponding colour-only one, so equality of the richer histograms implies equality of the coarser ones. □

Corollary 90 (Same bounds after augmentation)**.** *Let $k \geq 2$ and $K \equiv_{C^{k+3}} L$. Then the conclusion of Theorem 84 holds verbatim for the paired recurrence run on $\mathrm{Aug}(K)$ and $\mathrm{Aug}(L)$. Likewise,*

if $K \equiv_{C^4} L$*, the conclusion of Theorem 89 holds for the local recurrence run on* $\mathrm{Aug}(K)$ *and* $\mathrm{Aug}(L)$*.*

Proof. Apply Theorem 81 with $m = k+3$, respectively $m = 4$, to obtain $\mathrm{Aug}(K) \equiv_{C^{k+3}} \mathrm{Aug}(L)$, respectively $\mathrm{Aug}(K) \equiv_{C^4} \mathrm{Aug}(L)$, in their matched extended colour vocabularies; then apply Theorem 84, respectively Theorem 89, to the augmented inputs, whose derived neighbourhoods are computed from the augmented Face relation exactly as in Definition 74. □

D.7 One-directional factorization of the permitted computations

Definition 91 (Permitted computations)**.** A permitted computation of order $k \geq 2$ initializes the embedding of each tuple by a common function of atp_k and, at each finite layer, computes a common function of its previous embedding and a permutation-invariant function of the same context multiset as in Definition 82, with previous colours replaced by previous embeddings and the atomic context retained. A permitted computation of order 1 initializes by any common function of rank and colour, including forgetting the rank, and each layer computes a common function of the previous embedding and the four message channels of Definition 87, with optional common message maps applied to the central, neighbour and mediator embeddings. Embedding dimensions, real parameters, invariant functions and the finite number of layers are arbitrary. The readout is a common permutation-invariant function of all final tuple, respectively cell, embeddings, possibly after selection by a fixed rank or equality pattern. The only admissible additional argument is global side information that has separately been proved equal for the two compared inputs and is supplied uniformly.

Theorem 92 (Colour factorization and equal outputs)**.** *Fix a permitted computation as in Definition 91, with the same maps and parameters on both inputs. For every finite layer* t *there is a function* F_t*, common to both inputs, such that the layer-*t *embedding of any tuple, respectively cell, equals* F_t *applied to its discrete colour* χ^t*. Consequently, if* $K \equiv_{C^{k+3}} L$ *for order* $k \geq 2$*, or* $K \equiv_{C^4} L$ *for order* 1*, the computation returns the same output on* K *and* L*, and likewise on* $\mathrm{Aug}(K)$ *and* $\mathrm{Aug}(L)$*, at every finite depth and for every allowed common parameter choice.*

Proof. Induction on layers. At layer 0 the embedding is a common function of the atomic type, respectively of rank and colour, and χ^0 is exactly that data, or in the colour-only case a coarsening of it through which the initialization factors; this gives F_0. Assume the claim at layer t and let two objects have the same colour χ^{t+1}. By injectivity of the shared coding their previous colours agree and their entire context multisets agree, with multiplicities and channel labels; applying F_t to every shifted, neighbour and mediator colour turns these equal discrete multisets into equal multisets of numerical messages, with the same atomic context. The common message, invariant aggregation and update maps therefore receive equal arguments; hence the layer-$(t+1)$ embeddings agree, defining F_{t+1} on colour classes, with no injectivity of any neural map required. For the output, Theorem 84, respectively Theorem 89, gives equality of the complete final colour multisets, and applying F_T elementwise gives equality of the final embedding multisets, also after selection by a fixed rank or equality pattern, by Remark 86; a common permutation-invariant readout therefore returns the same value. Global side information separately proved equal enters as the same fixed argument and does not disturb the induction. The augmented case is Corollary 90 together with the same argument on the augmented inputs. □

D.8 Faithful height-four inputs and the parameter substitution

Write $I(S), I(T)$ for the source incidence structures of Appendices A and C, and set $X_S := \mathcal{Q}(L_2(S))$, $X_T := \mathcal{Q}(L_2(T))$, with full cubical incidence structures $J(S), J(T)$ carrying the dimensions and stipulated construction-role predicates of Appendix C; $\mathrm{Aug}(X_S)$ is the geometric star augmentation, with structure $\mathrm{Aug}(J(S))$, and likewise for T. The symbol Sub denotes binary strict containment. Fix a profile order $k \geq 1$, a prime p and $s \geq 2$, and put $r := 2(k+3)$.

By Appendix A, whose base construction follows Atserias et al. (2007), there are connected pure two-dimensional abstract simplicial complexes S, T with identical f-vectors and $I(S) \equiv_{C^r} I(T)$, whose affine interpretation costs $3r$ and uses a connected cubic base H with $\mathrm{tw}(H) > 3r = 6(k+3)$. Write $b = |E(H)| - |V(H)| + 1 \geq 1$. For each triple (k, p, s), choose its pair after fixing k; the tuple

order is fixed independently of the input. Applying Theorem 63 of Appendix C with $h := k+3 \geq 4$, so that $2h = r$, gives $J(S) \equiv_{C^{k+3}} J(T)$, both for the unaugmented full cube structures and, independently, for their anchor augmentations. Each structure contains all cube cells, the full strict geometric face relation, the cube dimensions, and the fixed construction-role predicates, which may be encoded as shared finite colour bits or forgotten.

Lemma 93 (Faithful cell-set representation)**.** *Take the ground vertices of the abstract cell complex to be all elements of the height-four lattice $L_2(S)$, each a 0-cell of X_S. Represent each 0-cell by its singleton, each cover edge by its two endpoints, and each corner square $[v, t_\Delta]$ by $\{v, E_{\eta_1}, E_{\eta_2}, t_\Delta\}$. Then distinct cells have distinct vertex sets, and strict inclusion of vertex sets coincides exactly with the full strict geometric face relation.*

Proof. Apply Proposition 108 to the Boolean intervals classified in Appendix B. They are exactly the singletons, cover edges and corner squares listed here, so its vertex-set identification preserves every cell and the full strict face relation, with geometric dimensions as ranks. □

Every positive-dimensional cell has a descending boundary path to a 0-cell. Adjoining a new ground vertex β together with the sets $\{\beta, v\}$ for all old 0-cells v produces exactly the anchor augmentation of Definition 79, with no duplicated and no higher new cells, and with no individual cell names and no arbitrary extra attributes.

D.9 The full homology statement, including the free summands

Let $R := f_0(S) + 3f_1(S) + f_2(S) - 2$, the same number for S and T because their f-vectors agree; by Appendix B the cell counts and the added free-circle counts of the two full cube complexes coincide, and $X_S \simeq |S| \vee (\bigvee_R \mathbb{S}^1)$.

Theorem 94 (Coefficient-depth separation of the pair)**.** *With the notation above,*

$$H_1(X_S;\mathbb{Z}) \cong \mathbb{Z}^R \oplus (\mathbb{Z}/p^s)^{b+1}, \qquad H_1(X_T;\mathbb{Z}) \cong \mathbb{Z}^R \oplus (\mathbb{Z}/p^s)^b \oplus \mathbb{Z}/p^{s-1}.$$

For the anchor augmentations, with $N := |L_2(S)| = |L_2(T)|$,

$$H_1(\mathrm{Aug}(X_S);\mathbb{Z}) \cong \mathbb{Z}^{R+N-1} \oplus (\mathbb{Z}/p^s)^{b+1},$$
$$H_1(\mathrm{Aug}(X_T);\mathbb{Z}) \cong \mathbb{Z}^{R+N-1} \oplus (\mathbb{Z}/p^s)^b \oplus \mathbb{Z}/p^{s-1}.$$

In both the original and the augmented case, the homology modules with coefficients in any field $\mathbb{F}$ and with coefficients in $\mathbb{Z}/p^j$ for every $1 \leq j < s$ are isomorphic in every degree; at $j = s$ the first homology modules are not isomorphic.

Proof. The displayed integral groups are the wedge decomposition of Appendix B combined with the source groups of Appendix A, using that homology of a wedge of well-pointed spaces splits (Hatcher, 2002); the free rank R is the same for both since it depends only on the common f-vector. For the anchor, the full cube complex is connected with N old 0-cells, and the augmentation adds one vertex β, N new edges and no 2-cells; contracting one new edge is a homotopy equivalence absorbing β, after which the remaining $N-1$ new edges have both endpoints in the connected old complex and each contributes one free circle, so the augmented space is the old one wedged with $N-1$ circles. Higher homology is unchanged. Over any field $\mathbb{F}$ and over $\mathbb{Z}/p^j$ with $1 \leq j < s$ the source modules agree in every degree by Appendix A, and the free contributions produce the same modules $\mathbb{F}^R$, respectively $(\mathbb{Z}/p^j)^R$, and likewise with $R+N-1$; the wedge decomposition therefore gives isomorphic modules in every degree over each coefficient ring. At $j = s$ the integral H_0 is free, so the degree-one universal coefficient isomorphism is $H_1(-;\mathbb{Z}/p^s) \cong H_1(-;\mathbb{Z}) \otimes \mathbb{Z}/p^s$, and every displayed torsion cyclic factor has order dividing p^s, hence survives tensoring. The common free summand agrees, while the torsion subgroups $(\mathbb{Z}/p^s)^{b+1}$ and $(\mathbb{Z}/p^s)^b \oplus \mathbb{Z}/p^{s-1}$ have orders $p^{s(b+1)}$ and p^{sb+s-1}, differing by a factor p. Both have exponent p^s, so the separation lies in the individual cyclic-factor orders, not in the total exponent, and the degree-one modules with $\mathbb{Z}/p^s$ coefficients are non-isomorphic. □

Remark 95 (Observation profile)**.** Theorem 92 bounds the distinguishing power of the two profiles in Definition 91, with their stated signatures and recurrences. The homology comparisons identify individual module isomorphism types. Appendix H extends the agreement to the permitted abstract coefficient diagrams, including the specified star augmentation and skeletal observations. These diagrams follow the abstract-input contract of Section 2.

E JOINT OBSERVATION CONSEQUENCE

Definition 96 (Constructed family)**.** Fix a prime p, an integer $s \geq 2$ and a permitted tuple order k (local profile for $k = 1$, paired profile for $k \geq 2$). Let (X, Y) be either the matched source pair of connected pure two-dimensional simplicial complexes of Appendix D, for which the common free first-homology rank is $R = 0$, or its matched full height-four lattice cube pair, for which $R \geq 0$ is the common free rank; a common anchor shifts R identically in both members. Let $b = |E(H)| - |V(H)| + 1 \geq 1$ be the common base cycle rank. By Theorems 89 ($k = 1$), 84 ($k \geq 2$) and 92 (neural realization) the permitted model outputs of X and Y agree at every finite depth at order bound k (R1), while $H_i(X; \mathbb{F}) \cong H_i(Y; \mathbb{F})$ for every field $\mathbb{F}$ and degree i (R2), $H_i(X; \mathbb{Z}/p^j) \cong H_i(Y; \mathbb{Z}/p^j)$ as modules for every i and every $1 \leq j < s$ (R3), and

$$H_1(X; \mathbb{Z}) \cong \mathbb{Z}^R \oplus (\mathbb{Z}/p^s)^{\,b+1}, \tag{1}$$

$$H_1(Y; \mathbb{Z}) \cong \mathbb{Z}^R \oplus (\mathbb{Z}/p^s)^{\,b} \oplus \mathbb{Z}/p^{\,s-1}. \tag{2}$$

Remark 97 (Constructed inputs)**.** The source complexes are pure. The full cube complexes have maximum cell dimension 2 but also carry graph-only cells, and the anchored inputs need not be full Boolean-interval complexes of another height-four lattice. A suitable pair exists for each fixed k, and the pair may depend on k. The features and the allowed invariant observation are exactly those of Appendix D, with no additional input data. The finite abstract coefficient diagrams of Appendix H also agree in the shallow category specified there, so Corollary 98 applies to the enlarged observation.

Corollary 98 (Joint observation forces chance accuracy)**.** *Let (X, Y) be as in Definition 96 and let O be any observation assembled from permitted model outputs at the fixed order bound k together with the homology-module isomorphism types of (R2) and (R3), and from nothing else. Then $O(X) = O(Y)$; consequently, on the uniform binary task on $\{X, Y\}$ with X labelled 0 and Y labelled 1, every binary classifier whose conditional output distribution depends on the input only through O has expected accuracy exactly $1/2$, over the uniform hidden test draw, the classifier's coins and any shared learned parameters. In particular this holds for every deterministic rule of this form.*

Proof. Every constituent of O agrees on X and Y by (R1), (R2) or (R3), so $O(X) = O(Y) =: o$. The shared parameters θ are fixed before the draw, and the coins ω are independent of the hidden identity. Put $q(\theta) = \Pr_\omega[A(o, \theta, \omega) = 1]$; the common input o makes this the same probability on both sides. The expected accuracy at θ is $\frac{1}{2}[(1 - q(\theta)) + q(\theta)] = \frac{1}{2}$. Averaging over θ gives $1/2$. Deterministic rules have $q(\theta) \in \{0, 1\}$. □

Proposition 99 (Exact depth-s readout, uniform over the fixed (p, s) family)**.** *Fix p and $s \geq 2$ and let (X, Y) be any matched pair of Definition 96, with any $R \geq 0$ and any permitted $b \geq 1$. For a finite complex K, define the readout as follows:*

$$\ell_s(K) = \log_p \big| H_1(K; \mathbb{Z}/p^s) \big|, \tag{3}$$

$$d_s(K) = \ell_s(K) \bmod s, \tag{4}$$

$$\text{prediction}(K) = \begin{cases} 1, & d_s(K) \neq 0, \\ 0, & d_s(K) = 0. \end{cases} \tag{5}$$

Then (i) $\ell_s(X) = s(R+b+1)$ and $\ell_s(Y) = s(R+b)+s-1$; (ii) $d_s(X) = 0$ and $d_s(Y) = s-1 \neq 0$, so the single rule (3)–(5), depending on no parameter beyond p and s, labels every pair of the family exactly; (iii) $\ell_j(X) = \ell_j(Y) = j(R + b + 1)$ for every $1 \leq j < s$, where ℓ_j is (3) with s replaced by j.

Proof. Both complexes are connected, so $H_0(\,\cdot\,; \mathbb{Z})$ is free, the Tor term of the universal coefficient theorem (Hatcher, 2002) vanishes and

$$H_1(K; \mathbb{Z}/p^m) \cong H_1(K; \mathbb{Z}) \otimes \mathbb{Z}/p^m, \qquad K \in \{X, Y\},\ m \geq 1. \tag{6}$$

With $\mathbb{Z} \otimes \mathbb{Z}/p^m \cong \mathbb{Z}/p^m$ and $\mathbb{Z}/p^a \otimes \mathbb{Z}/p^m \cong \mathbb{Z}/p^{\min(a,m)}$, equations (1)–(2) and (6) give (i) at $m = s$, and (iii) at $m = j < s$, where $\min(s, j) = \min(s - 1, j) = j$. Hence $\ell_s(X) \equiv 0$ and $\ell_s(Y) \equiv s - 1 \pmod s$ with $1 \leq s - 1 < s$, which is (ii); neither residue depends on R or b, so one rule serves the whole fixed (p, s) family, and (iii) agrees with (R3), the exploited discrepancy first appearing exactly at coefficient depth s. □

Remark 100 (What the readout uses)**.** The rule reads only the isomorphism type of $H_1(K;\mathbb{Z}/p^s)$. Computing this module is a global operation on the whole input, so it lies outside the permitted model profile of Appendix D. By (R3), O stops strictly below coefficient depth s.

Remark 101 (Natural-HOMP scope and related results)**.** Appendix G proves the natural higher-order message-passing result under its own feature and signature conditions. The model correspondence follows Akbari et al. (2026). Eitan et al. (2025) showed that higher-order message passing does not in general determine homology. The contribution here is the controlled realization of Theorem 94, joint in field, power depth and fixed order.

F THE MARKED INTEGRAL FILLING BRIDGE

Proposition 102 (Marked observation and filling bridge)**.** *Let p be prime, $s \geq 2$, $q = p^s$, $a = p^{s-1}$, and let $r \geq 2$. Let H be a connected simple cubic base graph with $\mathrm{tw}(H) > 6r$, and let E^0 and E^a be the two affine instances built on H as in Appendix A. Let $X(E^0)$ and $X(E^a)$ be the associated full-cube complexes, marked by six unary predicates $P_0, \dots, P_5$, each interpreted by a single 0-cell, and let z_X be the marked query 1-cycle carrying the common scalar a. Then*

(i) $J_{\mathrm{mark}}(X(E^0)) \equiv_{C^r} J_{\mathrm{mark}}(X(E^a))$*;*

(ii) z_X bounds a 2-chain in $X(E^a)$ over $\mathbb{Z}$ and over $\mathbb{Z}/p^s$, and bounds no 2-chain in $X(E^0)$ over $\mathbb{Z}$ or over $\mathbb{Z}/p^s$;

(iii) z_X bounds in both complexes over every ring $\mathbb{Z}/p^j$ with $1 \leq j < s$, over every field of characteristic p including every extension field, and over every field of characteristic different from p.

The six prescribed marks are interpreted by zero-arity tags on the affine source and by singleton 0-cells in the full cube. Preservation of C^r equivalence is established for exactly these six marks.

Proof. Part (i) is (18); parts (ii) and (iii) are proved on the simplicial source in Section F.3 from the certificate maps of Sections F.1–F.2.1 and transported to the full cube by Sections F.5–F.5.2. □

F.1 CONVENTIONS AND AN EXACT DISK CALCULATION

Write $G = \{t\} \sqcup U$. The base graph is the wedge of the oriented three-edge circles $\gamma_g = [o, A_g] + [A_g, B_g] + [B_g, o]$ $(g \in G)$. An ordered pair denotes the directed edge as a signed chain and an ordered triple the oriented triangle; if a stored private-edge basis uses the opposite descriptor orientation, both its coefficient and its cochain value acquire that sign, so the formulas below apply to the specified descriptor-oriented matrices and put no order on the observable structure.

For a relator instance ρ, let $y_0, \dots, y_N = y_0$ be its positive-word walk, with cyclic indices on the private vertices, and use the following abbreviations:

$$e_i = [y_i, y_{i+1}], \quad d_i = [a_i, a_{i+1}], \quad v_i = [a_i, y_i], \quad z_i = [a_i, y_{i+1}], \quad k_i = [h, a_i].$$

The e_i may revisit the same base edge; the private edges in the other four families are distinct exactly as proved in Lemma 25, even at $N = 3$ or with repeated source coordinates. The three prescribed oriented triangles have the following boundaries:

$$\partial A_i = d_i + v_{i+1} - z_i, \qquad \partial B_i = v_i + e_i - z_i, \qquad \partial C_i = k_i + d_i - k_{i+1}. \tag{7}$$

With $D_\rho = \sum_{i=0}^{N-1}(B_i - A_i + C_i)$, equation (7) telescopes to

$$\partial D_\rho = \sum_i (e_i + v_i - v_{i+1} + k_i - k_{i+1}) = \sum_i e_i, \tag{8}$$

so a power disk satisfies $\partial D_{g^q} = q\gamma_g$. For an equation relator with residue $0 \leq c < q$, put $\delta_\rho = \mathbf{1}[c > 0]$ and $d = (-c \bmod q) = -c + q\delta_\rho$; its positive word and raw presentation column obey $\partial D_\rho = \sum_j \gamma_{x_j} + d\gamma_t = \Phi(A_\rho) + q\delta_\rho\gamma_t$, where A_ρ is a column, generators being rows throughout. Thus the *raw equation chain* is

$$F_\rho = D_\rho - \delta_\rho D_{t^q}, \qquad \partial F_\rho = \Phi(A_\rho), \tag{9}$$

and $F_{g^q} = D_{g^q}$ for a power column. Every baseline coefficient and all power columns are retained: the correction is zero when $c = 0$ and is one negative t^q disk when $c > 0$.

F.2 Constructive equivalence of compact solutions and integral fillings

The map $\Phi : \mathbb{Z}^G \to Z_1(B_G;\mathbb{Z})$, $b \mapsto \sum_g b_g\gamma_g$, is an isomorphism, hence injective: the boundary equations at A_g, B_g force equal coefficients on the three edges of each circle, and different circles have disjoint edges. This concerns cycles supported on the base graph and assumes nothing about an arbitrary cycle in the larger complex.

If $Ax = b$, define $U_x = \sum_\rho x_\rho F_\rho$ over all columns; by (9), $\partial U_x = \Phi(b)$, an explicit integer linear map from a compact solution to a simplicial filling. Conversely expand an arbitrary simplicial 2-chain uniquely as $U = \sum_\rho \sum_i (\alpha_{\rho i}A_{\rho i} + \beta_{\rho i}B_{\rho i} + \chi_{\rho i}C_{\rho i})$ and suppose ∂U is supported on the base graph. Suppressing ρ, the coefficients of the private z_i, d_i, v_i, k_i in (7) must vanish:

$$-\alpha_i - \beta_i = 0, \qquad \alpha_i + \chi_i = 0, \qquad \alpha_{i-1} + \beta_i = 0, \qquad \chi_i - \chi_{i-1} = 0. \tag{10}$$

The first three give $\alpha_i = \alpha_{i-1}$ and $\beta_i = \chi_i = -\alpha_i = m_\rho$ for one integer m_ρ, and the fourth then holds. No other relator has any of these private edges, so nothing further can cancel these equations. This gives the following chain identity:

$$U = \sum_\rho m_\rho D_\rho. \tag{11}$$

With m_{t^q} the power-disk coefficient, extract the raw coefficients as follows:

$$x_\rho = m_\rho \ (\rho \text{ an equation}), \qquad x_{g^q} = m_{g^q} \ (g \neq t), \qquad x_{t^q} = m_{t^q} + \sum_{\rho \text{ equation}} \delta_\rho m_\rho. \tag{12}$$

By (8), $\partial U = \Phi(Ax)$; if it equals $\Phi(b)$, injectivity of Φ gives $Ax = b$. Hence (9), (11) and (12) are mutually inverse on fillings of base-supported cycles, and the positive-word change of relator coordinates is an integral triangular change with unit diagonal. A supplied base-supported 1-chain can be checked to be a cycle and its three equal edge coefficients read off to obtain b; no such unchecked coordinate promise is made for arbitrary 1-cycles in S.

F.2.1 Modular obstructions

Fix $M \geq 2$, work in $R = \mathbb{Z}/M$, and let a compact obstruction satisfy $wA = 0$, $wb \neq 0$. Set $\eta([o, A_g]) = w_g$ and $\eta([A_g, B_g]) = \eta([B_g, o]) = 0$, so $\eta(\gamma_g) = w_g$. For each disk put $h_i = \eta(e_i)$, $\lambda_0 = 0$, $\lambda_i = \sum_{j<i} h_j$. The power columns give $qw_g = 0$; the equation columns with (9) give $\sum_i h_i = 0$ for every positive word, hence $\lambda_N = \lambda_0 = 0$ in R, as required for cyclic indices. Extend the cochain by the following assignments:

$$\eta(v_i) = 0, \qquad \eta(z_i) = \eta(d_i) = h_i, \qquad \eta(k_i) = \lambda_i. \tag{13}$$

The evaluations on (7) are $h_i + 0 - h_i = 0$, $0 + h_i - h_i = 0$ and $\lambda_i + h_i - \lambda_{i+1} = 0$, including the last cyclic triangle. All triangles of S are accounted for and private edges of different disks are disjoint, so η is a 1-cocycle with $\eta(\Phi(b)) = wb \neq 0$.

Conversely, from any R-valued simplicial 1-cocycle with $\eta(\Phi(b)) \neq 0$, set $w_g = \eta(\gamma_g)$; evaluating on the power-disk boundaries gives $qw_g = 0$ and evaluating on (9) gives $wA_\rho = 0$ for every raw equation, so $wA = 0$ and $wb \neq 0$. Extension followed by restriction recovers w; restriction followed by extension need not recover an arbitrary cochain edge by edge, which is neither needed nor claimed.

All construction and extraction equations hold over any commutative coefficient ring, since (10) uses no division, and known signed permutations compose with these maps: if $A' = RAC$ and $b' = Rb$ with R, C signed permutation matrices, then $x = Cx'$ converts a solution and $w' = wR^{-1}$ a modular obstruction. Such coordinate access belongs to the target and certificate interface, not to the model's observable input.

F.3 The canonical query cycle and its coefficient-depth boundary

Fix $a = p^{s-1}$, $q = p^s$, $s \geq 2$. Let $\rho_{e,0}$, $\rho_{v,0}$ be the zero-index edge and vertex relators. In E^a every zero-index edge has $c = 0$, every zero-index vertex other than u has $c = 0$, and u has $c = a$. Hence the signed integral chain

$$U_a = \sum_{e \in E(H)} D_{\rho_{e,0}} - \sum_{v \in V(H)} D_{\rho_{v,0}} + D_{t^q} \tag{14}$$

has boundary

$$\partial U_a = \sum_{(v,e)} \gamma_{x_{v,e,0}} - \Big(\sum_{(v,e)} \gamma_{x_{v,e,0}} + (q-a)\gamma_t\Big) + q\gamma_t = a\gamma_t,$$

each incidence appearing once in each of the first two sums and cancelling exactly. Equation (14) is the positive-answer filling certificate; its coefficients need not all be nonnegative.

For E^0, take $M = q$ and

$$w_t = 1, \qquad w_{x_{v,e,i}} = i \pmod q. \tag{15}$$

Every power column evaluates to zero; a vertex column evaluates to $i + j + k - c = 0$ and an edge column to $i + j - c = 0$ modulo q, with exactly the nonzero baseline residues of Appendix A retained. Extending (15) by (13) gives $\eta(a\gamma_t) = a \neq 0$ in $\mathbb{Z}/q$. This forbids an integral filling, since an integral boundary would evaluate to zero modulo q, and equally forbids a filling over $R_s = \mathbb{Z}/q$; the chain (14) reduced modulo q shows the positive side still fills at depth s.

For every $1 \leq j < s$, $a = 0$ in $R_j = \mathbb{Z}/p^j$, so the query is literally the zero 1-chain on both sides and the zero 2-chain fills it; likewise over every field of characteristic p, including every extension field. Over any field of characteristic different from p, including characteristic zero, q is invertible and $\partial\big((a/q)D_{t^q}\big) = a\gamma_t$ on both sides. This gives the all-field and shallower fillability directly, without inferring it from equality of abstract homology types.

F.4 The six source marks and the variable budget of the counting translation

The unmarked source interpretation of Appendix A has a disjoint finite tagged domain of arity at most three, and, crucially, six distinct *zero-arity* old-face tags, each with the single empty tuple as its unique representation and domain true: V_0, V_1, V_2 for the vertices o, A_t, B_t and E_0, E_1, E_2 for the edges $\{o, A_t\}, \{A_t, B_t\}, \{B_t, o\}$. For an interpretation branch with object tag τ, the translation of one of these unary atoms is true if τ is its indicated tag and false otherwise: a fixed tag test using no source variable. No affine source element is selected, named or ordered; the tags index the interpretation and are absent from the observed gadget colours.

All old dimension, strict-subface and equality translations are unchanged, so the established counting translation pulls each C^{2r} sentence in the *marked* simplicial signature back to C^{6r} on the original affine source. In particular,

$$\mathrm{tw}(H) > 6r \implies I_{\mathrm{mark}}(S(E^0)) \equiv_{C^{2r}} I_{\mathrm{mark}}(S(E^a)). \tag{16}$$

The common affine-source interpretation with zero-arity domains proves (16) for the six prescribed marks, using the affine counting transfer of Atserias et al. (2007). This conclusion asserts neither preservation under arbitrary marking nor definability of these singleton marks in the unmarked simplicial structure.

F.4.1 Explicit cost-two full-cube lift and marker atoms

The full-cube tagged domain and unmarked atomic translations are those of Definition 64 and Lemmas 65 and 66; all cells are retained. For the tags used below, write $V(v) := \mathsf{v}(v)$ and $E(e) := \mathsf{E}(e)$ in the notation of Appendix C. The six observable predicates have the following translations, false on every tag not displayed:

$$\begin{array}{llll} P_0(V(v)) \iff V_0(v), & P_2(V(v)) \iff V_1(v), & P_4(V(v)) \iff V_2(v), \\ P_1(E(e)) \iff E_0(e), & P_3(E(e)) \iff E_1(e), & P_5(E(e)) \iff E_2(e). \end{array} \tag{17}$$

Every representation has at most two source coordinates and each atom uses only coordinates already supplied, so (17) adds no variables.

For ordinary counting, reserve a pair of source names for each target variable. A counting quantifier over the disjoint tags is a finite disjunction over allocations of its numerical threshold among those tags: a zero-arity tag contributes zero or one, a one-arity tag uses unary counting, and an arity-two tag uses Lemma 70, including its threshold-zero case. Branches reuse the reserved pair and rebinding shadows only that pair; there is no quotient or multiplicity of representations to correct. Hence every cubical C^r sentence has a source-incidence C^{2r} translation, including the unary marks, and with (16) this proves

$$J_{\mathrm{mark}}(X(E^0)) \equiv_{C^r} J_{\mathrm{mark}}(X(E^a)) \tag{18}$$

under the exact hypothesis $\mathrm{tw}(H) > 6r$; equivalently the direct composite reserves at most six affine-source names per cubical variable, and no increased treewidth threshold is used.

The source marks name the directed three-cycle, and (17) names, in order, its vertex/edge/vertex/edge/vertex/edge subdivision, so adjacent cube marks determine a unique cover edge of that cycle. The orientation is therefore the named traversal $P_0 \to P_1 \to \cdots \to P_5 \to P_0$, with common scalar a; no orientation of any other cell is observable.

F.5 Subdivision chain map, descent of fillings, and cochain transport

Write V_v, M_e, T_σ for the vertices of Q labelled by a source vertex, edge and triangle, and define the subdivision maps as follows:

$$s_0(v) = V_v, \qquad s_1([u,v]) = [V_u, M_{uv}] + [M_{uv}, V_v]. \tag{19}$$

For an oriented source triangle $\sigma = [v_0, v_1, v_2]$, with $M_{ij} = M_{\{v_i,v_j\}}$ and $T = T_\sigma$, take the oriented corner-square boundary traversals

$$Q_0 = [V_{v_0}, M_{01}, T, M_{20}], \qquad Q_1 = [V_{v_1}, M_{12}, T, M_{01}], \qquad Q_2 = [V_{v_2}, M_{20}, T, M_{12}], \tag{20}$$

the square notation meaning the oriented 2-cell whose boundary follows that loop, and set $s_2(\sigma) = Q_0 + Q_1 + Q_2$. Every radial edge between an edge midpoint and T occurs twice with opposite directions, so

$$\partial s_2(\sigma) = s_1\big([v_0,v_1] + [v_1,v_2] + [v_2,v_0]\big) = s_1 \partial\sigma, \tag{21}$$

and $\partial s_1 = s_0 \partial$ directly from (19). Extending linearly, with sign reversal for reversed orientations, gives chain maps in degrees zero through two over $\mathbb{Z}$ and over each coefficient ring under discussion (Hatcher, 2002); they agree on every shared source edge, however many triangles contain it.

The face labelled T in Figure 1 is σ, whereas T_σ denotes its corner vertex. The figure uses the oriented basis $q_A = Q_0$, $q_B = -Q_1$, $q_C = Q_2$ from (20). Thus its signed sum $F_2(T) = q_A - q_B + q_C$ is exactly $s_2(\sigma)$, with boundary $s_1\partial\sigma$.

In particular $s_1(\gamma_t) = \gamma_X$, exactly in the six marked positions, and the marked query is $z_X := a\gamma_X = a\, s_1(\gamma_t) = s_1(a\gamma_t)$. The positive full-cube filling is $s_2(U_a)$, the characteristic-different-from-p filling is $(a/q)s_2(D_{t^q})$, and their boundaries are z_X; the zero-chain fillings in characteristic p and at shallower depth remain zero.

F.5.1 Descent of every filling of a subdivided source query

The double-guard interval classification of Appendix B proves that every 2-cell of the full cube is one of (20) and that there are no cells above dimension two, all extra cells being vertices or edges; so any 2-chain in the full cube is already a 2-chain in Q.

Let W be such a chain with $\partial W = s_1(z)$ for a source 1-cycle z, in particular the query $a\gamma_t$. Within the three corners of one oriented triangle write the coefficients as c_0, c_1, c_2; the coefficients of the radial edges oriented $M_{01} \to T$, $M_{12} \to T$, $M_{20} \to T$, read from (20), are as follows:

$$c_0 - c_1, \qquad c_1 - c_2, \qquad c_2 - c_0. \tag{22}$$

Each such edge belongs only to the two indicated corners of this same source triangle, and $s_1(z)$ contains no radial edge, so (22) forces $c_0 = c_1 = c_2 = n_\sigma$ for every triangle, over any coefficient ring. With $V = \sum_\sigma n_\sigma \sigma$ we get $W = s_2(V)$, and (21) yields $s_1(\partial V - z) = 0$. Since the two halves of each source edge involve the unique midpoint M_e and the half-edge sets of different source edges are disjoint, s_1 is injective on source 1-chains; therefore $\partial V = z$. This is exact descent of any filling of a subdivided query, with no homotopy and no assumption that a filling avoids extra edges: extra edges cannot occur in the boundary of a 2-chain at all.

F.5.2 Explicit transport of any source 1-cocycle

Let η be a source 1-cocycle with values in a commutative ring R, including the obstruction (13), (15) with $R = \mathbb{Z}/q$. Choose one orientation $u \to v$ for each source edge, solely for certificate construction, and define the half-edge cochain values:

$$\eta_Q([V_u, M_e]) = \eta([u,v]), \qquad \eta_Q([M_e, V_v]) = 0. \tag{23}$$

Equivalently, with $b_e(x) = \eta_Q([V_x, M_e])$ for either endpoint x, one has $b_e(u) = \eta([u, v])$, $b_e(v) = 0$, $b_e(u) - b_e(v) = \eta([u, v])$. For each oriented triangle $\sigma = [v_0, v_1, v_2]$ choose local potentials $f_\sigma(v_0) = 0$, $f_\sigma(v_1) = \eta([v_0, v_1])$, $f_\sigma(v_2) = \eta([v_0, v_1]) + \eta([v_1, v_2])$; the cocycle condition on σ gives $f_\sigma(y) - f_\sigma(x) = \eta([x, y])$ for every directed edge of that triangle. Consequently

$$f_\sigma(M_e) = f_\sigma(x) + b_e(x) \tag{24}$$

is independent of which endpoint x of $e \subset \sigma$ is used, the potential difference being exactly the opposite endpoint's needed correction. Set $f_\sigma(T) = 0$ and use the following assignment:

$$\eta_Q([M_e, T_\sigma]) = -f_\sigma(M_e). \tag{25}$$

All four edges of any corner square now carry differences of the local values f_σ, whose sum around the oriented square is zero, so η_Q is a cubical 1-cocycle. The half-edge values (23) are globally fixed and the midpoint-to-triangle edges (25) are private to that incident pair, so two triangles impose no inconsistent assignment on a shared edge; no division or field assumption is used.

Equations (23) and (19) give $\eta_Q(s_1(c)) = \eta(c)$ for every source 1-chain c. Finally, define the full-cube cochain as follows:

$$\eta_X(e) = \begin{cases} \eta_Q(e), & e \text{ an edge of } Q, \\ 0, & e \text{ an extra edge of the full cube.} \end{cases} \tag{26}$$

Every 2-cell has its entire boundary in Q, so (26) is a cocycle on the full cube, explicitly accounting for all extra edges rather than deleting them. Its evaluation on z_X is $a \neq 0$ in $\mathbb{Z}/q$ for E^0, proving both integral and depth-s nonfillability in the full cube. Conversely, restriction of any full-cube cocycle to Q, followed by $\eta_S(e) = \eta_X(s_1(e))$, is a source cocycle by (21), with the same query evaluation.

The complex contains every cell of the full Boolean-interval cube of the double-guard height-four lattice and has the common free rank $R = f_0(S) + 3f_1(S) + f_2(S) - 2$ in integral H_1. The full cube need not be pure two-dimensional. The common free summands preserve the equal all-field and shallow homology module types. The construction uses no anchor or broadcast extension.

F.6 Marked observation and model consequences

Given $k \geq 1$, set $r = k + 3$ and choose the connected cubic base with $\mathrm{tw}(H) > 6r = 6(k + 3)$; then (18) gives C^{k+3} equivalence in the exact marked full-cube signature, which for $k = 1$ is the bound C^4 of the local four-channel recurrence of Theorem 89 and for $k \geq 2$ the bound of the paired recurrence of Theorem 84. The finite unary query marks enter the shared initial bit colours and atomic types, and the formulas and common-function factorization of Appendix D admit those finite predicates without any new variable, under the model correspondence to Akbari et al. (2026).

The compared marked pair has equal complete frontend observations under those profiles and equal permitted neural invariant readouts at every finite depth with common parameters and shared feature vocabulary, by Theorem 92. By Sections F.3 and F.5 its integral and depth-s labels are opposite, while its previously established global field and shallow homology types are identical and every corresponding field or shallow query-fillability Boolean is true on both sides, so supplying only those identical side data preserves equality of the observation.

A decoder whose *entire* input is this common observation together with the identical side data has the same output distribution on the two instances, while the two correct binary decisions are opposite; on their uniform mixture its average decision success is therefore at most $1/2$, equal to $1/2$ if it always decides, by the two-point calculation of Corollary 98. Correct decision with a valid certificate is a subset of correct decisions, so certified success obeys the same bound, and abstention or invalid certificates cannot improve it.

The corresponding statement for natural HOMP architectures is Corollary 110 of Appendix G, conditional on equivalence in this same signature, including all six unary marks, and on a common fixed initialization with an invariant readout.

G NEURAL INTERFACES: NATURAL HOMP AND SMCN

The constructed pair's counting equivalence transfers to natural-neighborhood higher-order message passing (HOMP) of Eitan et al. (2025, Section 3, equations (1)–(3)). The interface applies to the same constructed pair and its established homology groups.

Throughout, a *structure* is finite, with equality, one binary strict full-containment predicate F and unary rank predicates $R_0, \dots, R_d$ forming a fixed partition. The vocabulary also includes finitely many unary *label predicates*. Every element, called a *cell*, satisfies exactly one R_r. Each cell's *complete initial label* is its truth pattern on the label predicates. Put $L_{\preceq}(x, y) := x = y \vee F(x, y)$. Each complete label is defined by its truth-pattern formula. For a finite bit vector of unary marks, each truth pattern conjoins the corresponding positive and negated mark atoms. The six unary query marks of Appendix F are one such vector. Each conjunction is quantifier-free in the single free variable x. A finite mark vector therefore encodes a complete initial label at no variable cost.

Definition 103. For every fixed pair of ranks r, t the four *natural binary neighborhood predicates* are

$$\begin{aligned} B_{rt}(x, y) &= R_r(x) \wedge R_t(y) \wedge L_{\preceq}(x, y), \\ A_{rt}(x, y) &= R_r(x) \wedge R_r(y) \wedge \exists z \big(R_t(z) \wedge L_{\preceq}(x, z) \wedge L_{\preceq}(y, z)\big), \\ B^T_{rt}(x, y) &= R_r(x) \wedge R_t(y) \wedge L_{\preceq}(y, x), \\ \mathrm{co}A_{rt}(x, y) &= R_r(x) \wedge R_r(y) \wedge \exists z \big(R_t(z) \wedge L_{\preceq}(z, x) \wedge L_{\preceq}(z, y)\big), \end{aligned}$$

the natural neighborhoods of Eitan et al. (2025) written over the present vocabulary.

Each predicate carries explicit rank guards on x, on y and on the mediator, so a channel is identically empty off its prescribed ranks and rank-specific message functions are well defined. No clause $y \neq x$ occurs and $L_{\preceq}$ is reflexive by its equality disjunct, so a cell may be its own neighbor in any channel whenever the guards permit. In A_{rt} and $\mathrm{co}A_{rt}$ the mediator is existentially quantified, so membership is Boolean: a cell joined to x by many mediators contributes exactly one copy of its colour, the point of divergence from the mediated recurrence of Appendix D. Only x, y occur freely and only z is bound.

Definition 104. Fix a finite ordered list $N_1, \dots, N_m$ of predicates from Definition 103. The *exact natural colour refinement* is defined as follows:

$$\begin{aligned} c_0(x) &= \big(\text{rank of } x,\ \text{complete initial label of } x\big), \\ c_{u+1}(x) &= \Big(c_u(x),\ \big(\{\!\{\, c_u(y) : N_i(x, y) \,\}\!\}\big)_{i=1}^{m}\Big). \end{aligned}$$

Each multiset ranges over the *distinct* cells y with $N_i(x, y)$, so multiplicity n records exactly n such cells of that colour.

Definition 105. A numerical HOMP model is *natural-neighborhood admissible* for $N_1, \dots, N_m$ if: (H1) the initial feature vector of a cell is a fixed function of its rank and complete initial label, common to the two inputs; (H2) every layer uses channel- and rank-specific message functions of the center and neighbor vectors, common to the two inputs; (H3) each neighbor multiset is aggregated by a permutation-invariant aggregation with a common convention on the empty multiset, the subsequent combination and activation being common; (H4) the readout is a common invariant function of the per-rank multisets of final vectors. Feature dimensions may differ from rank to rank, and no injectivity of any map is assumed.

Theorem 106. *Let X, Y be finite structures over the above vocabulary and fix any finite channel list from Definition 103. If $X \equiv_{C^3} Y$ over this vocabulary, then for every finite round u the whole-input histogram of the colours c_u agrees on X and Y, as does the rank-selected histogram of c_u for every rank r. Consequently every natural-neighborhood admissible model produces, after any finite number of layers, the same per-rank multiset of feature vectors and hence the same mathematical output on X and on Y.*

Proof. Induction on u: every round-u colour c has a C^3 formula $\varphi^u_c(x)$, with x its only free name, true exactly of the cells of colour c. At $u = 0$ take the rank atom conjoined with the truth-pattern conjunction of label atoms. For a record $(c, M_1, \dots, M_m)$ put $\varphi^{u+1}(x) := \varphi^u_c(x) \wedge \bigwedge_i \mathrm{Ex}_{N_i(x,y), M_i}$,

by Lemma 77 applied per channel with counted block the single name y, $\mathsf{C}_{\geq n}$ the ordinary counting quantifier, guard $N_i(x, y)$ and classes $\varphi_a^u(y)$. The class formulas are pairwise contradictory by induction. The empty-support case reads $\neg\exists y\, N_i(x, y)$: channel i is empty at x.

Three names suffice. The predicate $N_i(x, y)$ uses x, y and, in $A_{rt}, \mathrm{co}A_{rt}$, a mediator z bound within that conjunct. Before substituting φ_a^u, permute the whole pool $\{x, y, z\}$ so that its free name becomes y. This bijection preserves every binder–occurrence link and truth under the renamed assignment. Internal binders remain local and cannot capture the free x of $N_i(x, y)$. The outer counting quantifier on y binds both conjuncts. All conjunctions are finite.

For each $n \geq 0$ the sentence $\exists^{=n} x\, \varphi_c^u(x)$ is again C^3. Its unary counting quantifier adds no name, so every colour count agrees, including $n = 0$. Finitely many colours occur, giving equal whole-input histograms. Each rank-r histogram is their restriction because every colour retains rank. The sentence $\exists^{=0} x\, (x = x)$ also ensures that the two inputs agree on emptiness.

The layer vectors factor through c_u as in Theorem 92: (H1) at $u = 0$; and $c_{u+1}(x)$ determines $c_u(x)$ and each channel multiset of round-u colours, hence by (H2) the messages, by (H3) their aggregation under the common empty convention, and the common combination and activation, with no injectivity used and rank-specific dimensions harmless. Then (H4) on the equal per-rank histograms gives equal outputs. □

Remark 107. The conclusion concerns exact mathematical invariant functions at every finite depth; numerical comparisons use the tolerances of Section 5 to accommodate floating-point reduction order. The result covers any fixed finite tensor-diagram schedule that unrolls into updates of the form (H2)–(H3). Unused channels may be omitted or may send the same neutral message on both inputs. Cumulative colours cover identity updates and residual retention. Global summaries passed to the decoder must also agree. Each finite round is treated separately. A round with no further split yields an equitable partition that certifies all later rounds for that recurrence and pair.

Local attention. For a permitted neighborhood $N(x)$, a shared score $e(h_x, h_y)$ and value map $v(h_y)$ give

$$m(x) = \frac{\sum_{y \in N(x)} \exp(e(h_x, h_y)) v(h_y)}{\sum_{y \in N(x)} \exp(e(h_x, h_y))},$$

with the common zero convention when $N(x)$ is empty. The numerator and denominator are sums of allowed messages; normalization is a shared combination. Hence (H2)–(H3) cover this operation. On an equitable quotient, each class contribution is multiplied by its number of neighboring cells.

Proposition 108. *Let a cell of the complete Boolean-interval cube complex be the interval $[a, b]$ with $a \leq b$, represented by its vertex set $\{x : a \leq x \leq b\}$, with rank the geometric cell dimension. Then $\{x : a \leq x \leq b\} \subseteq \{x : c \leq x \leq d\}$ iff $c \leq a \leq b \leq d$, iff $[a, b]$ is a face of the cube $[c, d]$; and distinct cells have distinct vertex sets.*

Proof. If $c \leq a \leq b \leq d$, every x with $a \leq x \leq b$ satisfies $c \leq x \leq d$; conversely the inclusion puts a, b in $\{x : c \leq x \leq d\}$, so $c \leq a \leq b \leq d$. The faces of the cube $[c, d]$ are exactly the subintervals $[a, b]$ with $c \leq a \leq b \leq d$. The vertex set has unique minimum a and maximum b, recovering the interval; F excludes equality. □

Hence geometric full containment coincides with set containment of vertex sets, and since the identification is by cells rather than by dimension truncation, all cells of the full-lattice realization of Appendix B, including its additional graph-like cells, remain present.

Corollary 109. *Fix a prime p and $s \geq 2$ and consider the unmarked pair of Appendix D over the realization of Appendix B, with $\mathrm{tw}(H) > 6(k + 3)$ for $k \geq 1$, for which Appendices A and C give equivalence of the two full-lattice cube structures in ordinary C^{k+3}. Since $k + 3 \geq 4 \geq 3$ they are C^3-equivalent, and rank-only initial vectors satisfy (H1); hence every natural-neighborhood admissible HOMP model produces identical mathematical outputs on the two structures at every finite depth and for every parameter setting. All coefficient comparisons are inherited from Theorem 94, with the free summand of Appendix B: equality over every field and over $\mathbb{Z}/p^j$ with $j < s$, the integral and depth-s distinction over $\mathbb{Z}$ and $\mathbb{Z}/p^s$, and the common free summand of the complete lattice realization.*

Proof. Every C^3 sentence is a C^{k+3} sentence for $k+3 \geq 3$, so the hypothesis yields C^3 equivalence and Theorem 106 applies together with its factorization; the homological statements are those of Theorem 94. □

Corollary 110. *Assume the marked construction of Appendix F with a prime p, $s \geq 2$, $r \geq 2$ and $\mathrm{tw}(H) > 6r$, so that the two marked structures satisfy the ordinary C^r equivalence (18) in the expanded vocabulary containing exactly the six unary query marks, and assume in addition $r \geq 3$ (in practice $r = k + 3$ with $k \geq 1$). Then every natural-neighborhood admissible HOMP model over the marked vocabulary, with initial features a fixed common function of rank and of the truth pattern of the six marks, produces identical mathematical outputs on the two marked structures at every finite depth. The inherited content is that of Appendix F: the canonical query cycle $a\gamma_X$ with $a = p^{s-1}$ fills only on the twisted side over $\mathbb{Z}$ and over $\mathbb{Z}/p^s$, and fills on both sides over every field and every $\mathbb{Z}/p^j$ with $1 \leq j < s$, by the certificate maps already proved there.*

Proof. Since $r \geq 3$, (18) gives C^3 equivalence over the same vocabulary; the assumed initial features are exactly (H1) and the channels are those of Definition 103 over the same rank partition and containment relation, so Theorem 106 and its factorization apply unchanged. □

Remark 111. The factorization uses the stated vocabulary and counting premises. Proposition 122 covers the specified broadcast roles; Proposition 114 supplies the standard SMCN simulation. For comparison, Akbari et al. (2026) gives counting-logic upper bounds. A decoder that receives only the common observations and identical additional inputs has the same output distribution on both structures, whereas the integral and depth-s query labels are opposite. Its success on the uniform pair is therefore at most $1/2$, with equality for an always-binary rule. The same bound holds in Appendix F when success requires both a correct decision and a valid certificate.

Remark 112. For the finite K_4-based example of Section 5, equality follows from a direct equitable-partition certificate rather than the treewidth premise $\mathrm{tw}(H) > 6r$. The certificate uses Definition 104, Boolean neighborhood membership and, in the marked version, the six marks. The mediator-counting certificate of Appendix D concerns a different recurrence. For the complete cube pair with all six query marks, joint refinement with all 36 natural channels is stable and equal at round 13. A balanced equitable partition of the joint classes then certifies equality at all later rounds. This all-depth guarantee holds for the 36-channel natural recurrence and this finite pair.

Remark 113. Equal refinement colours give equal vectors inductively. Equal rank-selected multiplicities give equal pooled outputs. Sum, mean or max aggregation, rank-specific updates and residual maps are covered without an injectivity assumption.

G.1 The standard SMCN interface

We use the scalable multicellular network (SMCN) with subcomplex layer (SCL) updates from Eitan et al. (2025, Appendix D, Eq. (55)): tensor domains have total order at most two, updates use the stated natural neighborhoods, equivariant layers use the equality-orbit bases of their Eqs. (48)–(50), and nonlinear maps act pointwise. Inputs are common functions of the permitted unary labels and the Boolean masks for equality, full containment, natural adjacency and coadjacency, and their finite Boolean combinations. Readouts are the rank-selected sum, mean or zero aggregations followed by a shared MLP in their Eq. (95), or compositions of their invariant linear maps and pointwise maps. These operations define the interface below.

Proposition 114 (Order-four SMCN simulation)**.** *On the complete two-dimensional source and full-cube inputs, and their specified star augmentations, $\mathsf{Obs}_4(X) = \mathsf{Obs}_4(Y)$ implies equal mathematical outputs for this interface at every finite depth and width. The binary incidence initialization has finite, size-independent preparation. The distance initialization of Eq. (91) in Eitan et al. (2025), computed on the current $A_{0,1}$ vertex graph, is also determined by the stable profile.*

Proof. At completed network layers, a four-tuple u stores all required unary and ordered-pair feature slots on its coordinates. Temporary stages also store three-coordinate accumulators. Every slot is a common function of the current paired colour; typed slots are masked to zero outside their rank domains, including after nonlinear maps.

A sum over a replacement coordinate is available by selecting $\alpha = \beta = z$ in the paired contexts. Each z occurs exactly once. The shifted colour and extended atomic type supply the message and its

mask. A spare coordinate *position* is chosen outside the output positions even when output values coincide.

Adjacent-rank containment is atomic. Vertex-to-two-cell containment is equivalent to a two-step containment through an edge on these full inputs, so one round recovers it. A second round recovers natural adjacency and coadjacency by existence of a correctly ranked mediator, counting each neighbor once. This supplies the Boolean pair inputs. An incidence-sum initialization, such as Eq. (82), uses a further permitted aggregation; the subsequent Eq. (89) maps are pointwise. Each SCL term in Eq. (55) retains its output coordinates and sums one other coordinate, so the same operation implements it.

A general typed pair-to-pair orbit map requires both source coordinates:

$$(L_E h)(x,y) = \sum_{a,b} E(x,y,a,b)h(a,b), \qquad T(x,y,a) = \sum_{b} E(x,y,a,b)h(a,b).$$

Here E is a rank-compatible equality-orbit indicator. Store T after one replacement sum in the fourth position, then sum the third position in the next round. Apply this construction for every output-position pair, reusing the remaining positions as source slots. Diagonal output values are retained. Shared linear combinations and pointwise maps complete the layer. Lower-order cases use at most the same two sums. Induction therefore realizes every finite diagram through paired colours.

For readout, selecting $u_1 = u_3$ and $u_2 = u_4$ gives each ordered pair exactly once, including the diagonal; selecting all four coordinates equal gives each cell once. Rank restrictions recover the required joint feature multisets and mean denominators, with a common empty-set convention. Remark 86 and equal colour histograms therefore give equal outputs for the specified readouts.

For Eq. (91), initialize $R_0(x,y) = [x \preceq y]$ and iterate $R_{t+1}(x,y) = R_t(x,y) \vee \exists z\,(A_{0,1}(x,z) \wedge R_t(z,y))$. Recording first arrival gives vertex-to-cell distances within $n_0 - 1$ steps on the connected one-skeleton. Both inputs have the same n_0. Thus their stable profiles determine this initialization; its preparation can depend on input size. □

For joint coverage of a prescribed order k and this SMCN interface, apply the original construction at $K = \max\{k, 4\}$ with $\mathrm{tw}(H) > 6(K+3)$. The inherited counting equivalence supplies both observations on that pair. SMCN alone uses $K = 4$ and the sufficient threshold $\mathrm{tw}(H) > 42$. Section 5 evaluates natural HOMP and local attention; the SMCN conclusion follows from this simulation.

H Coherent shallow coefficient information

The two complexes in each source or full lattice cube pair admit a chain isomorphism over the shallow coefficient ring. Standard homological algebra then identifies their coherent coefficient information and its precision boundary.

H.1 Conventions and homology of the pair

Fix a prime p, an integer $s \geq 2$, a permitted observation order $k \geq 1$, and the pair of complexes of the primary construction. Put $q = p^s$, $a = p^{s-1}$ and $A = \mathbb{Z}/p^{s-1}$. Write ρ for the common free integral first-homology rank (denoted R earlier), $b \geq 1$ for the common source-base cycle rank and h for the common free integral second-homology rank; X, Y denote either side. We treat two cases separately: the connected pure simplicial *source pair* and its full two-guard height-four *lattice cube pair*. Both are taken with all cells retained. All complexes are the ordinary *unreduced* cellular complexes $C_*(K;\mathbb{Z})$ in degrees $0,1,2$ with $\partial_0 = 0$. In particular, C_0 is retained in full, without augmentation or quotient by a vertex (Hatcher, 2002).

The source pair has equal degreewise census $(c_0, c_1, c_2) = (v, e, f)$ on both sides with $\rho = 0$ (29, 31, 39, 40); the full cube pair has equal degreewise census $(v + 4e + f + 2,\ v + 8e + 4f,\ 3f)$ with $\rho = v + 3e + f - 2$, retaining every extra graph cell (52, 53). In both cases $h = v - e + f - 1$, $H_0 \cong \mathbb{Z}$, $H_2 \cong \mathbb{Z}^h$, homology vanishes above degree 2, and

$$H_1(X;\mathbb{Z}) \cong \mathbb{Z}^\rho \oplus (\mathbb{Z}/q)^{b+1}, \qquad H_1(Y;\mathbb{Z}) \cong \mathbb{Z}^\rho \oplus (\mathbb{Z}/q)^{b} \oplus \mathbb{Z}/(q/p).$$

Over A both sides have $H_i(K;A)$ *free as A-modules* of ranks $\beta = (\beta_0, \beta_1, \beta_2) = (1,\ \rho + b + 1,\ h + b + 1)$, zero above degree 2; this is the established decomposition at $j = s - 1$. The structural observation equality, the all-field homology equality, the deep and integral distinction and the marked-query facts keep their original hypotheses, in particular $\mathrm{tw}(H) > 6(k+3)$ for the stipulated finite-order structural statements; the counting and model bounds are those of Appendices C and D, the natural-HOMP interface is Appendix G, and the six-mark signature with the prescribed query $a\gamma$ is Appendix F.

H.2 Finite projective modules over the shallow ring

Lemma 115. *Every finitely generated projective A-module is free, of well-defined rank.*

A is local with maximal ideal (p) and residue field $\mathbb{F}_p$ (for $s = 2$, $A = \mathbb{F}_p$), so this is Stacks project authors (2026, Tag 00NX, Lemma 10.78.2), a finite projective module being a direct summand of a finite free module and free at maximal localisations; the rank is $\dim_{\mathbb{F}_p} P/pP$. The argument also applies when A has zero divisors.

H.3 Splitting of bounded free complexes with free homology

Lemma 116. *Let C be a bounded complex of finite free A-modules in degrees $0, \dots, d$ with every $H_i(C)$ finite free. Then, with $B_i = \operatorname{im} \partial_{i+1}$,*

$$C \ \cong\ \bigoplus_{i=0}^{d} H_i(C)[i] \ \oplus\ \bigoplus_{i=0}^{d-1} E_{i+1,i}(B_i),$$

where $H_i(C)[i]$ carries zero differential and $E_{i+1,i}(B)$ has B in degrees $i+1, i$ with differential id_B. Each B_i is finite free and, with $c_i = \operatorname{rank}_A C_i$, $\beta_i = \operatorname{rank}_A H_i(C)$, $u_i = \operatorname{rank}_A B_i$,

$$u_{-1} = 0, \qquad u_i = c_i - \beta_i - u_{i-1} \quad (0 \le i \le d), \qquad u_d = 0.$$

Hence two such complexes with equal (c_i) and equal (β_i) have equal unit-pair multiplicities and are A-chain isomorphic.

Proof. Put $Z_i = \ker \partial_i$, $B_{-1} = 0$; $\partial_0 = 0$ gives $Z_0 = C_0$. Inductively, assume B_{i-1} is projective. Then $0 \to Z_i \to C_i \xrightarrow{\partial_i} B_{i-1} \to 0$ splits, so Z_i is a finitely generated projective summand of C_i. The sequence $0 \to B_i \to Z_i \to H_i(C) \to 0$ splits because $H_i(C)$ is free. Thus B_i is finitely generated projective. Lemma 115 makes all Z_i, B_i finite free. Fixing both families of splittings gives degreewise $C_i \cong B_i \oplus H_i(C) \oplus B_{i-1}$ with ∂_i zero on the first two summands and the identity of the last onto the summand $B_{i-1} \subseteq C_{i-1}$, which is the asserted decomposition; $C_{d+1} = 0$ forces $B_d = 0$ and $B_{-1} = 0$ closes the other end. Taking ranks gives the recurrence; choosing bases then identifies both complexes with the same standard complex. This is a change of bases: every generator of every chain module is retained. □

H.4 The chain isomorphism for the full pair

Proposition 117. *Let X, Y be either the source pair or the full two-guard lattice cube pair, with their unreduced full cellular complexes. Then*

$$u_0 = c_0 - 1, \qquad u_1 = c_1 - c_0 - \rho - b, \qquad u_2 = 0 \ \ (\textit{equivalently } c_2 - u_1 = h + b + 1),$$

so $u_0 = v - 1$, $u_1 = e - v - b$ for the source ($\rho = 0$), and $u_0 = v + 4e + f + 1$, $u_1 = e + 2f - v - b$ for the full cube. Both $C_*(X;A)$ and $C_*(Y;A)$ are A-chain isomorphic to

$$\mathcal{S} = A[0] \oplus A^{\rho+b+1}[1] \oplus A^{h+b+1}[2] \oplus E_{1,0}(A^{u_0}) \oplus E_{2,1}(A^{u_1}),$$

and fixing $\theta_X : C_*(X;A) \to \mathcal{S}$, $\theta_Y : C_*(Y;A) \to \mathcal{S}$ and putting $F = \theta_Y^{-1}\theta_X$, $G = F^{-1}$ yields an actual A-chain isomorphism between the two full chain complexes, chosen once and for all.

Proof. Apply Lemma 116 to $C_*(K;A) = C_*(K;\mathbb{Z}) \otimes_{\mathbb{Z}} A$, bounded free in degrees $0, 1, 2$, with $\beta = (1, \rho + b + 1, h + b + 1)$ free. Equal degreewise chain ranks and equal free homology ranks then identify each complex with $\mathcal{S}$, whence F and $G = F^{-1}$. □

The maps θ_X, θ_Y, F and G exist after noncanonical basis choices. The result is an actual chain isomorphism.

H.5 Coherence over all coefficient modules

Theorem 118. *Fix $F, G = F^{-1}$ from Proposition 117 independently of the coefficient module. Let M range over* all *A-modules, including nonflat and infinitely generated modules. Morphisms are all A-linear maps, and coefficients use restriction of scalars along $\mathbb{Z} \to A$. Then:* (1) *$C_*(K;\mathbb{Z}) \otimes_{\mathbb{Z}} M \cong C_*(K;A) \otimes_A M$ and $\mathrm{Hom}_{\mathbb{Z}}(C_*(K;\mathbb{Z}), M) \cong \mathrm{Hom}_A(C_*(K;A), M)$ naturally;* (2) *$F \otimes 1_M$, $G \otimes 1_M$ are mutually inverse chain maps commuting with every $1 \otimes u$, $u : M \to N$, inducing natural isomorphisms $H_i(X;M) \to H_i(Y;M)$;* (3) *$G^*(\varphi) = \varphi \circ G$ is a cochain isomorphism with inverse F^* commuting with postcomposition by u, inducing natural isomorphisms $H^i(X;M) \to H^i(Y;M)$;* (4) *every short exact sequence $0 \to M' \to M \to M'' \to 0$ of A-modules induces degreewise exact sequences of chain and cochain complexes, and the above isomorphisms form commutative ladders with them, commuting with all connecting homomorphisms;* (5) *$H_i(K;M) \cong M^{\beta_i} \cong H^i(K;M)$ naturally, and consequently every connecting homomorphism of a short exact sequence lying wholly inside A-Mod vanishes in all degrees, in homology and cohomology.*

Proof. (1) Tensor associativity; and since $aM = 0$, every $\mathbb{Z}$-linear map to M kills $aC_*(K;\mathbb{Z})$ and factors uniquely A-linearly through $C_*(K;A)$. (2),(3) $GF = 1$ and $FG = 1$ are identities of A-linear chain maps, hence preserved by $- \otimes_A M$ and by $\mathrm{Hom}_A(-, M)$ for every M, since functors preserve inverse identities. Naturality follows because u acts on the other tensor factor or on the target of φ. The cohomology map $X \to Y$ is pullback along the inverse chain map G. (4) Each $C_i(K;A)$ is finite free, so both functors carry the coefficient sequence to a degreewise short exact sequence of complexes, and F, G give a morphism of such sequences; functoriality of the long exact sequence in short exact sequences (Stacks project authors, 2026, Tag `0117`, Lemma 12.13.12) yields the commuting ladders, connecting maps included. (5) Applying either functor to $\mathcal{S}$ gives M^{β_i} with zero differential plus acyclic pairs $E_{i+1,i}(M^{u_i})$. A coefficient sequence then yields degreewise $0 \to (M')^{\beta_i} \to M^{\beta_i} \to (M'')^{\beta_i} \to 0$, exact as a finite direct sum of copies, so the middle term already surjects onto the third in every degree and the long exact sequence forces every connecting map to vanish. □

For $s \geq 3$ the nonsplit sequence $0 \to \mathbb{F}_p \xrightarrow{\cdot p} \mathbb{Z}/p^2 \to \mathbb{F}_p \to 0$ lies inside A-Mod and still has vanishing connecting maps. The statement covers nonflat and infinitely generated coefficient modules.

H.6 Coherent abstract diagrams and the joint observation

Corollary 119. *Fix any shared finite collection of A-modules, A-linear coefficient maps and short exact sequences of A-modules, and form the diagram of homology and cohomology in the stated degrees with all induced and all connecting arrows. The fixed F of Proposition 117 identifies the whole diagram for X with that for Y. It commutes with every arrow, so the abstract diagram types agree. Consequently:* (1) *the tuple of these abstract coherent-diagram types, the inherited permitted structural observations and the all-field homology* types *is the same on both sides of the same original pair, under the inherited hypotheses, including $\mathrm{tw}(H) > 6(k+3)$;* (2) *any decoder whose entire input is that common tuple together with equal side information has the same output law on both sides, so on the uniform pair with opposite correct labels its success probability is $(P_0 + P_1)/2 \leq 1/2$, with equality for an always-binary rule, and valid-certificate outputs obey the same bound (Corollary 98);* (3) *the prescribed marked query is preserved as zero, $a\gamma \otimes m = \gamma \otimes am = 0$ for every A-module M and every $m \in M$.*

Proof. Theorem 118(2),(3) identifies every object naturally in M and (4) every coefficient and connecting arrow; all identifications come from the same F, G, hence commute at once. Adjoining the inherited structural observations and all-field homology types, equal under their original hypotheses, gives (1). Equal inputs give a decoder, randomized or not, the same law on both sides, so with opposite correct labels the average success is $\frac{1}{2}(P_0 + P_1)$ with $P_0 + P_1 \leq 1$ for mutually exclusive decisions; a valid certificate entails the corresponding correct decision, giving the same information-theoretic bound. (3) is the displayed computation. The depth-s information still separates the pair:

the integral first homology groups differ, and the readout of Appendix E with Proposition 40 distinguishes X from Y. □

H.7 THE PRESCRIBED NONZERO SHALLOW CYCLE

Proposition 120 (Pointed shallow-chain refinement)**.** *Fix a prime p, integers $s \geq 2$ and $k \geq 1$, put $q = p^s$, $a = p^{s-1}$ and $A = \mathbb{Z}/a$, and choose the marked source pair or the marked full two-guard cube pair on a connected simple cubic graph with $\mathrm{tw}(H) > 6(k+3)$. Let γ_X, γ_Y be the unscaled directed t-circles specified by the six marks of Appendix F, using their subdivisions in the cube case. There is an isomorphism of ordinary unreduced A-chain complexes*

$$F : C_*(X; A) \longrightarrow C_*(Y; A), \qquad F_1(\gamma_X) = \gamma_Y,$$

where the cycles in this display are reduced to A. Choose it once, independently of the coefficient module. Theorem 118 and Corollary 119 then apply to this fixed F over A.

For every A-module M, the maps

$$\iota_{K,M} : M \longrightarrow H_1(K; M), \quad m \longmapsto [\gamma_K \otimes m],$$
$$\mathrm{ev}_{K,M} : H^1(K; M) \longrightarrow M, \quad [\phi] \longmapsto \phi(\gamma_K)$$

are respectively split injective and split surjective, for $K = X, Y$. Splittings can be chosen naturally in M. Writing $G = F^{-1}$, the homology maps induced by $F \otimes 1_M$ and the cohomology maps induced by precomposition with G commute with these two maps, with the identity on M, and with all the coefficient and connecting arrows of Theorem 118. Consequently the two inputs have matching abstract finite coefficient diagrams even after these prescribed pointed maps are included. No raw cellular-basis alignment is supplied.

The integral classes have orders q and a, respectively. Their reductions in $H_1(-; \mathbb{Z}/p^j)$ both have order p^j for $1 \leq j < s$, so the distinguished shallow class is nonzero. The same marked task $a[\gamma_K] = 0$ still has opposite integral and depth-s answers. Appending the permitted structural observations and the original all-field homology types therefore retains the balanced-pair decoder bound under this enlarged, explicitly pointed shallow contract.

Proof. First track the distinguished generator, rather than just the abstract Smith factors. The shear of Lemma 14 leaves t fixed. The edge elimination of Lemma 15 also retains t. In Lemma 16, replace the distinguished vertex row by the sum of all vertex rows. It is then $-\epsilon t$. The other vertex rows contain no t term and eliminate the tree generators in terms of the chords. The remaining presentation is a direct sum of the b chord factors of order q and the factor generated by t, with relations qt and ϵt. Thus t itself generates the factor $\mathbb{Z}/\gcd(q, \epsilon)$, not a proper multiple of its generator.

The mapping-cone identification of Lemma 27 sends t to $[\gamma_t]$. The subdivision and inclusion maps of Appendix F send this class to the stated cube cycle; the free circles in Theorem 53 are a separate summand. Consequently the integral orders are q on X and a on Y. Naturality of the degree-one universal coefficient isomorphism, with free integral H_0, identifies the reduced class with the tensor of this generator. It has order $p^{\min(j,s)}$ on X and $p^{\min(j,s-1)}$ on Y. In particular over A it generates a free direct summand A of H_1 on each side.

Choose bases of the free A-homology modules beginning, in degree one, with $[\gamma_K]$. In the proof of Lemma 116, choose the section of $Z_1 \to H_1$ to send this basis vector to the actual cycle γ_K. Such a section exists by choosing lifts of the other free basis vectors. The remaining split sequences and bases are as in that lemma. Equal chain ranks and homology ranks give the same standard complex $\mathcal{S}$ of Proposition 117 and isomorphisms $\theta_K : C_*(K; A) \to \mathcal{S}$ that both send γ_K to its first isolated degree-one basis vector. Hence $F = \theta_Y^{-1}\theta_X$ has the stated cycle-preserving property. This uses the entire unreduced chain complex, not a quotient by a vertex.

Tensoring the mutually inverse maps with any M and applying $\mathrm{Hom}_A(-, M)$ preserves their inverse identities; no flatness assumption is needed. The distinguished isolated copy of A in $\mathcal{S}_1$ gives the natural splitting of $\iota_{K,M}$ and of $\mathrm{ev}_{K,M}$. The identities $F\gamma_X = \gamma_Y$ and $G\gamma_Y = \gamma_X$ give the following compatibility identities, respectively:

$$H_1(F \otimes 1_M)\iota_{X,M} = \iota_{Y,M}, \qquad \mathrm{ev}_{Y,M}(G^*[\phi]) = \phi(G\gamma_Y) = \mathrm{ev}_{X,M}([\phi]).$$

Theorem 118 supplies compatibility with every A-linear coefficient map and every short exact sequence internal to A-Mod, including the vanishing connecting maps. The choices were fixed before M, so all the arrows commute simultaneously.

The directed cycle carries exactly the six marks interpreted in Appendix F, so the corresponding counting equivalence applies. The integral order of the class determines the truth of $a[\gamma] = 0$; the filling and cocycle constructions there give the depth-s answers. The preceding maps identify the pointed shallow diagrams. The all-field homology types agree. Corollary 98 therefore gives the balanced-pair bound for this enlarged observation. □

The prescribed t-cycle is determined by the six query marks.

H.8 The coefficient precision boundary

Remark 121. The coefficient category A-Mod fixes the precision available to connecting maps. For $s = 2$ one has $a = p$ and $A = \mathbb{F}_p$, and the middle term of $0 \to \mathbb{F}_p \xrightarrow{\cdot p} \mathbb{Z}/p^2 \to \mathbb{F}_p \to 0$ is not annihilated by a. This sequence therefore lies outside A-Mod and outside the scope of Theorem 118. On a two-term integral torsion block with differential p^u, its first homological Bockstein is computed on a mod-p cycle lift by $p^{-1}\partial$ reduced mod p, i.e. p^{u-1}, which is nonzero for $u = 1$ and zero for $u \geq 2$. At $s = 2$ the relevant blocks for $b = 1$ are p^2, p^2 on one side and p^2, p on the other (in general $(p^2)^{b+1}$ against $(p^2)^b, p$). This Bockstein therefore has rank 0 on one side and rank 1 on the other, and it separates the $s = 2$ examples. For $s \geq 3$ the same sequence lies inside A-Mod, the pertinent differentials are divisible by p^2, and the Bockstein vanishes on both sides, in agreement with the theorem.

An A-chain isomorphism is an abstract module map. Accordingly, the decoder input of Corollary 119 receives homological information only through the allowed isomorphism types. Section 2 specifies the input contract; the prescribed marked cycle is treated in Proposition 120. The argument is the standard splitting of a bounded free complex with projective homology, followed by tensor and Hom naturality. Its specific inputs are the equal full complex censuses and the freeness of the A-homology. For coefficient-precision locality of connecting information, see Ghrist & Ding (2025, Lemma 5.9 and the proof of Theorem C, Section 5.3, version 1).

H.9 Broadcast and skeletal observations

Proposition 122 (Specified broadcast-star closure)**.** *For the marked source or full-cube pair of Proposition 120, apply Definition 79, retaining the six old marks, false on new cells, and the two common new roles. The augmented pair retains the permitted structural equivalence and opposite integral and depth-s filling labels. Its full A-chain complexes admit an isomorphism extending the fixed pointed map, so all admitted pointed coefficient diagrams still match.*

Proof. Put $N = c_0$. Choose an old root v_* and old edge chains with $\partial P_v = v - v_*$. Orient each new e_v from β to v. Keeping the old bases, set $d = v_* - \beta$, $f = e_{v_*}$ and $w_v = e_v - f - P_v$ for $v \neq v_*$. The inverse changes $\beta = v_* - d$, $e_{v_*} = f$ and $e_v = w_v + f + P_v$ for $v \neq v_*$ give integral bases, with $\partial f = d$ and $\partial w_v = 0$. Thus:

$$\Phi_K : C_*(K;\mathbb{Z}) \oplus E_{1,0}(\mathbb{Z}) \oplus \mathbb{Z}^{N-1}[1] \xrightarrow{\cong} C_*(\operatorname{Aug} K;\mathbb{Z}).$$

The first summand is the actual old inclusion. Counts become $(c_0 + 1, c_1 + N, c_2)$ and free shallow homology ranks become $(1, \rho + b + N, h + b + 1)$; the unit multiplicities are $u_0 + 1, u_1$. Reduce Φ_K to A and extend the fixed pointed map by $\widehat{F} = \Phi_{Y,A}(F \oplus 1 \oplus 1)\Phi_{X,A}^{-1}$. It preserves the actual cycle and its free homology summand, with choices fixed before coefficients; the proof of Theorem 118 gives the required naturality. The old filling solution sets are unchanged because no two-cells or old boundaries change. Old marks translate to P on the old tag and to false otherwise, so Theorem 81 preserves the marked widths $2r$ and r of (16) and (18), for $r = k + 3$. Equal added free summands preserve the all-field comparisons. □

Corollary 123 (Dimension-skeletal diagrams)**.** *For the original pair and its specified augmentation, the diagrams $\big(H_i(K^{(r)};M), H_i(K^{(r)} \hookrightarrow K^{(t)};M)\big)_{r \leq t}$, where $K^{(r)}$ is the full r-skeleton, agree naturally for every A-module M. The prescribed cycle maps agree at stages one and two.*

Proof. The skeleton complex equals the full cellular complex in degrees at most r and is zero above. The degree-preserving pointed isomorphism restricts to every such complex and commutes with inclusions. Tensoring preserves these inverse identities; literal cycle transport gives the pointed assertion. □

H.10 Finite coefficients and scalar precision

Corollary 124 (Finite-coefficient separation)**.** *Let $r_s(M)$ count the p-primary cyclic factors of a finite abelian coefficient group M whose orders are at least p^s. For the source or full-cube pair, as abelian groups:*

$$\frac{|H_1(X;M)|}{|H_1(Y;M)|} = \frac{|H_2(X;M)|}{|H_2(Y;M)|} = p^{r_s(M)}.$$

Both comparisons agree precisely when $r_s(M) = 0$. Coefficients $\mathbb{Z}/m$, $m \geq 1$, separate precisely when $v_p(m) \geq s$.

Proof. Lemma 39 and Theorem 53 give quotient factors M/p^jM in degree one and kernel factors $M[p^j]$ in degree two, with common factors M^ρ and M^h. On $\mathbb{Z}/p^e$, each quotient and kernel has order $p^{\min(e,j)}$; prime-to-p parts contribute only to the common factors. Comparing $j = s$ and $s-1$ proves the ratios. When $r_s(M) = 0$, the two multiplication images and the two kernels coincide, giving the isomorphisms. The zero group has order one and is included. □

Remark 125 (Scalar-summary tolerance)**.** The depth-s summaries in Proposition 99 are consecutive integers. If $|\widetilde{\ell}_s(K) - \ell_s(K)| < 1/2$, nearest-integer rounding before the modulo-s rule preserves the exact label. This open radius is sharp for a scalar-only decoder: the midpoint is compatible with both opposite labels under a closed error allowance of $1/2$.

H.11 Coherent mixed coefficients

Corollary 126 (Mixed coefficient diagrams)**.** *For the pair of Proposition 120, put $\widetilde{A} = A \times \mathbb{Z}[1/p]$ and $e = (1, 0)$. A fixed $\widetilde{A}$-chain isomorphism preserves the actual prescribed cycles. It identifies every finite abstract diagram assembled from $\widetilde{A}$-modules, linear coefficient maps, internal short exact sequences, their induced homology, cohomology and connecting arrows, and the naturally split pointed maps:*

$$eM \longrightarrow H_1(K;M), \quad m \longmapsto [\gamma_K \otimes m], \qquad H^1(K;M) \longrightarrow eM, \quad [\phi] \longmapsto \phi(\gamma_K).$$

Every finite abelian group with p-primary exponent dividing p^{s-1} has a canonical $\widetilde{A}$-action respecting all group homomorphisms. Every field is also a $\widetilde{A}$-module.

Proof. Use the existing pointed map over A. Flat localization at p gives equal free homology $(\mathbb{Z}[1/p], \mathbb{Z}[1/p]^\rho, \mathbb{Z}[1/p]^h)$. The splitting argument of Lemma 116 applies over this PID. The localized γ_K is a boundary, and its original cell coefficient ± 1 supplies a functional on C_1 taking it to one. Restricting the functional to the free boundary module splits off $\mathbb{Z}[1/p]\gamma_K$. Choose a boundary basis beginning with this cycle and lift it to C_2. Matching this contractible pair in the standard complexes gives a cycle-preserving localized isomorphism. Glue the two component maps over $\widetilde{A}$.

Tensoring this fixed map and pulling cochains back along its inverse give natural isomorphisms. Finite free chain groups give exact coefficient ladders and compatibility with internal connecting maps. Since $(1-e)\gamma_K$ is a boundary, evaluation lies in eM; the pointed splittings are those of the A-component. A finite group decomposes canonically into its p-primary and prime-to-p parts. The former is an A-module, while multiplication by p is invertible on the latter. Every group homomorphism commutes with these actions. Fields of characteristic p use the first factor, and all other fields use the second. □

H.12 An explicit family of cubic bases

Proposition 127 (An explicit sufficient witness-size bound)**.** *Fix a prime p, $s \geq 2$ and $k \geq 1$. For $L \geq 2$, the port-cycle grid with subdivided-K_4 caps has $n = 4L^2 + 20L$ vertices, is connected,*

simple and cubic, and satisfies $\mathrm{tw}(H_L) \geq L - 1$. *Taking* $L = 6(K + 3) + 2$ *gives source, full-cube and specified star-augmented witnesses with* $O((K+3)^2 p^{4s})$ *cells. Here* $K = k$, *or* $K = \max\{k, 4\}$ *when invoking the separate SMCN interface; this is a sufficient size bound.*

Proof. Replace each grid vertex by a private four-cycle of directional ports and join matching ports across grid edges. Exactly $4L$ ports lack neighbors. At each attach a private K_4 with one edge subdivided, joining its subdivision vertex to that port. Each cap adds five vertices; all degrees become three, with no loops or repeated edges. Connectivity is retained, and the edge count is $3n/2$.

Delete caps and contract port cycles to obtain the grid. Applying this branch-set map to decomposition bags preserves connected occurrences and never increases bag size. In the grid, crosses $\mathrm{row}_i \cup \mathrm{column}_j$ are connected and pairwise intersect. Fewer than L vertices miss a row and a column, hence miss a cross. Bags meeting a cross form a subtree: its vertex-subtrees intersect along its edges. These subtrees pairwise intersect and share a node by the subtree Helly property. Indeed, root the tree and choose the deepest nearest-root vertex among the subtrees; pairwise intersection forces this vertex into every subtree. The common bag hits every cross, so has size at least L.

Now $L - 1 > 6(K + 3)$, meeting Theorem 10 and the inherited counting budgets. Lemma 29 gives $g = 1 + 3qn$, $M = g + nq^3 + 3nq^2/2$, and $\Lambda \leq 3(q + 2)M$, hence $O(nq^4)$ source faces uniformly for $q = p^s \geq 4$. Theorem 53 gives linear full-cube counts; Definition 79 adds one vertex in total and one edge per old vertex. Finally $n = O((K + 3)^2)$. □

I Related work and computational methods

I.1 Comparison with related work

Table 3 identifies the established ingredients and the joint construction.

Table 3: Nearest-prior boundaries. The target is one matched pair, built after p, s, k are fixed, that combines full-incidence observation equality, all-field and shallow-module equality, and integral/depth-s separation. Rows separate inherited ingredients from the obligations of this construction.

Source and anchor	Inherited result	Construction in this paper
Eitan et al. (2025) §3, Eqs. (1)–(3); Thms. 4.2–4.3	HOMP with natural neighborhoods is blind to homology; the torus/disconnected-tori witness changes Betti numbers. Each neighboring cell is aggregated once.	Match all stipulated coefficients on connected pure 2D sources; natural HOMP follows via C^3, and the standard SMCN interface via order-four simulation.
Akbari et al. (2026) §3; App. B, p. 14; Thms. 4.1 and C.8	Local/paired CCWL recurrences, global signatures and the broadcast anchor; the forward bound uses pair-counting TC^{k+2}.	Ordinary-counting and full-input transfer for the specified profiles, then the matched coefficient/topology construction.
Atserias et al. (2007) Def. 1; Lems. 2–3	For any fixed nontrivial finite Abelian group, large enough treewidth gives counting-equivalent solvable/unsolvable affine systems; fixed $\mathbb{Z}/p^s$ is already permitted.	Integral presentation and bounded-arity realization of every face and containment. Fixed-p, s counting indistinguishability, Smith normal form and universal coefficients (Hatcher, 2002) are inherited.
Dawar et al. (2023) Thm. 7; Lem. 8 (author v2)	All-prime invertible-map equivalence for suitable ring-CFI pairs over $\mathbb{Z}/2^i$, with i chosen after k.	The prescribed-p, s, k full-incidence/coefficient pair, from a different relational input, with ordinary-counting equivalence.
Ghrist & Ding (2025) Defs. 5.1–5.2; Thm. 5.3	Arithmetic precision and valuation barcodes encode torsion lengths for DVR network sheaves; locality of connecting information is Lemma 5.9.	One full-incidence pair also carrying coherent A-Mod diagrams, by the standard splitting of a bounded free complex; all internal A-linear coefficient maps are included.
Krishnan & Rudman (2025), §4.2.1, Prop. 4.5; Athanasiadis (2012), Ex. 7.12; Kahle et al. (2021), Thm. 5.1	Boolean-interval cubes, cubical subdivision, and realization inside a selected cubical subcomplex.	A height-four lattice controlling *all* Boolean intervals and the common free-circle remainder while preserving full-face observations.

Homology entries compare module isomorphism types; Proposition 8 extends them to coherent diagrams inside A-Mod. The decoder receives abstract summaries without cellular-basis cross-alignment.

Torsion-based reconstruction (Walch, 2025) and the announced persistence separation from fixed-order Weisfeiler–Leman refinement (Kim, 2025) concern adjacent observation settings. Our coefficient channel uses abstract homology summaries rather than arbitrary filtrations.

I.2 Fusion-model implementation

Each cell has a three-dimensional rank indicator and six query-mark bits, collected in a_x. The 36 natural channels are ordered by family $(A, \mathrm{co}\, A, B, B^{\mathsf{T}})$, then by rank pair. Each channel contains distinct neighboring cells, including the self-neighbors allowed by its predicate. For channel j, let $\alpha^{\ell}_{xyj} = 1$ for sum aggregation; for attention it is the softmax over $y \in N_j(x)$ of the local query–key

dot product scaled by $1/\sqrt{w}$. Empty channels contribute zero. With width w, the updates are:

$$
\begin{aligned}
h_x^0 &= \tanh(a_x W_{\text{in}} + b_{\text{in}}),\\
S_x^\ell &= \sum_{j=1}^{36} \sum_{y \in N_j(x)} \alpha_{xyj}^\ell \tanh(h_y^\ell W_{\ell j} + b_{\ell j}),\\
v_x^\ell &= \tanh\big(h_x^\ell W_{\ell,\text{self}} + (S_x^\ell/36) W_{\ell,\text{mix}} + b_\ell\big),\\
h_x^{\ell+1} &= \tanh(v_x^\ell W_{\ell,\text{second}} + b_{\ell,\text{second}}),\\
p(K) &= \text{concat}_{r=0}^{2} \frac{\sum_{x:R_r(x)} h_x^T}{|\{x : R_r(x)\}|}, \qquad f(K) = [p(K); z(K)] w_{\text{head}} + b_{\text{head}}.
\end{aligned}
$$

An empty rank contributes a zero mean. For fixed width and seed, the two encoders share their input, value, update and affine-head parameters; attention has additional query/key projections. Gaussian weights have standard deviation equal to the inverse square root of their input dimension, and biases have standard deviation 0.05. Widths are $8, 16$, seeds 0–3, and evaluated depths $0, 1, 2, 4$. Parameters are shared across each pair and remain untrained.

The nine summary slots are the rational and $\mathbb{F}_2$ Betti vectors and the modular log-order vector, each in degrees $0, 1, 2$. The four coefficient-depth-two modes use zeros, field summaries, fields with $\boldsymbol{\ell}_1$, and fields with $\boldsymbol{\ell}_2$. These summaries are computed from signed physical boundaries and measured Smith factors. Fillings and mod-4 cocycles independently certify the targets; target labels and twist parameters do not enter the features.

Saved stable partitions are reused to reconstruct and verify equitable quotients separately from the complete cell inventories on both sides, using the same neighborhood core. An entry $n_j(c, d)$ counts physical neighbors in class d for a cell in class c. Sum aggregation weights each class by this count; attention normalizes count-weighted exponentials. Rank-wise pooling uses physical class sizes. Class-constant initialization and induction identify the quotient and full-cell recurrences.

I.3 Verification of the constructed inputs

Physical-boundary calculations verify the chain condition, integer elimination and its trace, and the residual Smith form. Source instances also agree with the Smith form of their complete affine presentations. Expanded-cell and quotient computations agree on small fixtures, including an independent literal-neighborhood attention calculation. CPU/GPU comparisons on small fixtures agree within the tolerance of Section 5; feature swaps and whole-channel deletion change the outputs. Repeated evaluation reproduces the exact topology summaries and numerical outputs.

Figure 1 uses rank-only refinement, with stable rounds 14 and 18 on the source and full-cube $K_4, q = 4$ pair; Table 1 uses the six-mark queries. The four $s = 2$ fusion realizations come from three source bases, and the cube realization depends on the K_4 source. The same eight width/seed parameter draws are reused across the grid.

Matched-head attribution. Matched-head controls reconstruct all $1{,}024$ records from saved pooled vectors, measured summaries and eight original affine heads. Full, neural-only and summary-only arms retain the same head and bias; parameters agree exactly and logit reconstruction errors are at most 2.91×10^{-11}. For paired differences $\delta f = f(X) - f(Y)$:

$$\delta f_{\text{full}} = \delta f_{\text{neural}} + \delta f_{\text{summary}}.$$

The observed full–summary gap residual is at most 7.28×10^{-12}; the neural-only gap is at most 6.19×10^{-15}. At coefficient depth two, only the H_1 and H_2 modular-summary slots differ, each by one. The separating gap is therefore the sum of their two head weights, up to numerical residuals. Each head is reused in 32 settings.

Depth-three source realization. A separately specified K_4 source pair uses $p = 2, s = 3$, twists $0, 4$ and query 4γ. Each full inventory has $(f_0, f_1, f_2) = (51{,}324, 194{,}691, 145{,}800)$. Physical-boundary elimination and the complete raw-presentation Smith calculation agree:

$$H_1(X; \mathbb{Z}) \cong (\mathbb{Z}/8)^4, \qquad H_1(Y; \mathbb{Z}) \cong (\mathbb{Z}/8)^3 \oplus \mathbb{Z}/4.$$

Both have free H_2 rank 2,432. Exact UCT decompositions from the measured factors give H_1 log-orders $(4, 8, 12, 12, 12)$ and $(4, 8, 11, 11, 11)$ at the fixed precisions $j = 1, \ldots, 5$. Integer filling and mod-8 cocycle certificates independently verify the opposite labels. Six-mark refinement stabilizes at round 12 with 11,799 classes per side. The quotient matrices are reconstructed from all cells in all 36 channels, using the same neighborhood core. Mod-4 cocycles evaluate to 1 on each unscaled γ, certifying a split copy of $\mathbb{Z}/4$ in shallow homology. These cochains are verification certificates, not decoder features.

Forward checks at coefficient depth three. The $s = 3$ grid reuses the stated encoders, parameter grid and eight heads: 64 paired settings give 320 records across zero, field, mod-2, mod-4 and mod-8 summaries. Quotients rebuilt from all cells agree in all 36 channel matrices, features and multiplicities. The first four modes match within tolerance; the maximum pooled gap is 2.42×10^{-15}. Mod-8 separates all 64 settings, with gaps 0.0302–0.4488. Same-head three-arm additivity and the two differing summary weights agree within 4.55×10^{-13}. Feature swaps and whole-channel deletion change the outputs; an independent calculation from the saved vectors reproduces all 320 affine rows. Reusing certified inputs and partitions, preparation takes 14.8 seconds with four CPU threads; CUDA float64 forwards take 7.5 seconds on an RTX 5070 Ti. These times exclude construction and homology computation.

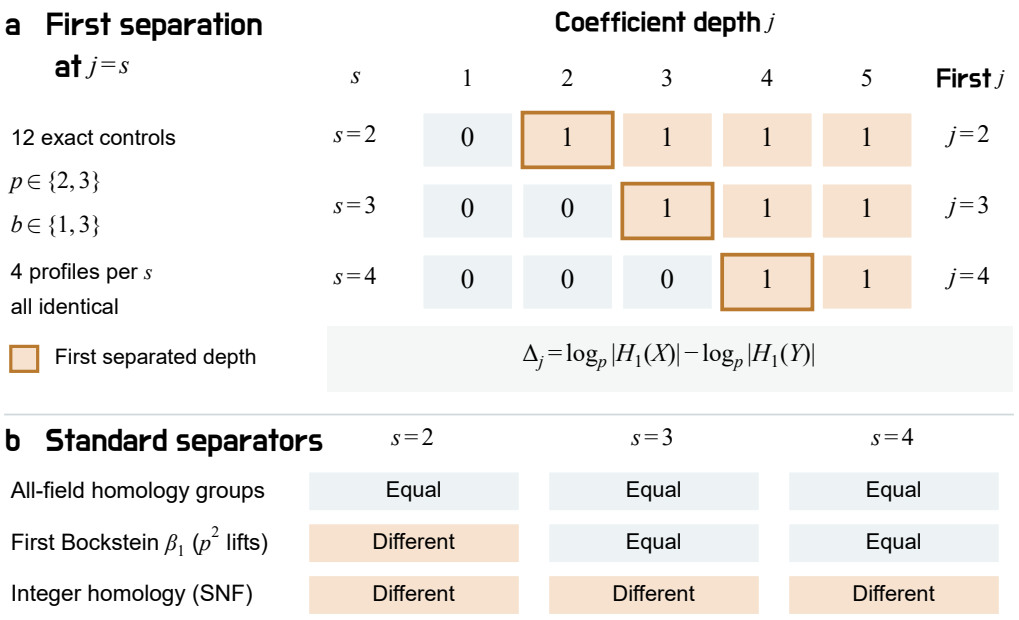

Figure 3: **The coefficient threshold and standard separators.** Twelve compact diagonal-chain controls use $p \in \{2, 3\}$, $s \in \{2, 3, 4\}$ and $b \in \{1, 3\}$. (a) For $j = 1, \ldots, 5$, the H_1 log-order gap over $\mathbb{Z}/p^j$ changes from 0 to 1 at $j = s$. (b) Diagonal algebra gives matching homology over every field. The first Bockstein uses mod-p^2 lifts and separates only $s = 2$, whereas integer SNF separates every pair.

I.4 Architecture and benchmark scope

General feature-filtration persistence (Horn et al., 2022), sheaf diffusion (Bodnar et al., 2022) and signed-incidence Hodge networks (Ebli et al., 2020; Roddenberry et al., 2021; Yang et al., 2022) require different observation assumptions. Feature lifting and hierarchical pooling in Cell Attention Networks (CAN) (Giusti et al., 2023) require a separate simulation. EXP, CEXP (Abboud et al., 2021) and SR25 (Balcilar et al., 2021) provide graph-discrimination benchmarks. Our constructed pairs provide certified coefficient-depth thresholds and filling labels for the validation in Section 5.