Companion Note: On the Structural Uniqueness of K = 7 and Boundary Reduction in the Discrete Admissibility Framework

Keith Taylor VERSF Theoretical Physics Program

Version: Rigorous Clarification (companion to "Deriving the Bekenstein–Hawking Entropy Coefficient from Boundary Constraint Counting in a Discrete Admissibility Framework")

Abstract

This companion note strengthens the self-contained logical foundations of the entropy-coefficient derivation based on boundary constraint reduction in a discrete admissibility framework. We provide:

(i) A structured sketch of the exhaustion logic selecting the constraint count K = 7, including an explicit worked example demonstrating failure at K = 6 on a minimal two-simplex patch; **(ii)** A canonical taxonomy of the seven constraints as an algebraic minimal generating set, with formal clarification of the status of distinguishability and redundancy channels; **(iii)** A careful derivation of the coarse-graining map from adjacency-graph distance to emergent boundary area, yielding the two-Planck independence threshold $\ell_e = 2\ell_p$ and boundary cell area $A\_cell = 4\ell_p^2$, with an explicit statement of the geometric assumptions and their scope; **(iv)** Necessary and sufficient structural conditions for exact pairing of tangential constraints across shared faces; **(v)** A corrected interpretation of the $\mathbb{Z}_2$ redundancy as a local sign-gauge on representatives with a residual global symmetry after matching; **(vi)** An explicit comparison of boundary constraint reduction on null, spacelike, and timelike surfaces; **(vii)** A clarification of the staging structure relating the toy example in Section 3.3 to the constraint matrix in Appendix C of the main paper.

The note is intentionally technical and focuses on the constraint algebra and counting logic. It does not modify the original manuscript but makes its logical dependencies and assumptions explicit and checkable.

Table of Contents

Notation and Scope

We use ℓ_p for the Planck length and ℓ_e for the independence (decoupling) threshold. A boundary "tile" (effective cell) denotes the minimal horizon patch supporting one independent irreversible binary commitment as seen by an external observer. A "simplex" denotes an elementary cell of the simplicial foam; in the boundary context, simplices are triangular. Variables X_f denote face data on a shared face f, and $s \in \{+1, -1\}$ denotes a binary orientation label. All equations are written in plain text with Unicode notation.

---

1. K = 7: Structured Exhaustion Sketch

The entropy derivation assumes a simplicial substrate in which each elementary simplex carries a finite set of independent admissibility constraints. Let K denote the number of independent constraint channels required for (a) local admissibility and (b) globally extensible gluing under small perturbations. The companion paper [20] provides the full derivation; here we give a concise, checkable sketch that exhibits why $K = 7$ is the minimal consistent value and, critically, demonstrate an explicit failure at $K = 6$.

1.1 Constraint Channels as a Minimal Generating Set

The seven channels used in the entropy paper can be grouped into three algebraic classes:

Class A — Local well-posedness (2 channels): Local closure and bulk redundancy elimination. These ensure that each simplex is internally self-consistent and that the independent constraint count is not inflated by spurious degrees of freedom.

Class B — Embedding and gluing (4 channels): Three neighbor compatibility conditions (one per shared face of a boundary triangle) and orientation/chirality consistency. These ensure that simplices can be assembled into a global tiling without contradiction.

Class C — Non-degeneracy (1 channel): Distinguishability commitment. This ensures that each simplex contributes at least one independent, stably distinguishable label not fixed by neighbor matching — without which the physical configuration space collapses and entropy is ill-defined.

The claim is that these seven channels form a *minimal generating set* for the constraint algebra under gluing. "Minimal" means that removing any one channel produces a strict pathology: underdetermined extensions, loss of perturbative robustness, or unphysical overcounting. "Generating" means that any complete, consistent constraint system for simplicial admissibility can be expressed as a combination of these channels and their algebraic consequences; no independent channel outside this set is needed.

Different descriptive decompositions of these seven channels are possible (e.g., splitting "local closure" into sub-components), but any such decomposition must be related to the canonical one by an invertible recombination that preserves the total rank $K = 7$. The constraint count is a property of the algebra, not of a particular labeling.

1.2 Why K < 7 Fails: Explicit Demonstration at K = 6

We demonstrate the failure concretely on a minimal patch: two triangular boundary simplices σ and σ' sharing a single face f.

Setup. Each simplex carries tangential face data on its boundary-facing edges. After null reduction (which freezes the normal closure channel), simplex σ has independent tangential variables (x_a, x_b), and simplex σ' has (x_a', x_b'). The shared face f imposes the matching constraint $x_a = x_a'$ (up to orientation sign, absorbed into the variable definition).

At K = 7, the full constraint set for each simplex consists of:

Channel	**σ**	**σ′**
C1: Local closure (tangential)	$C_\sigma(x_a, x_b) = 0$	$C_{\sigma'}(x_a', x_b') = 0$
C2–C4: Neighbor compatibility	M_f: $x_a = x_a'$; plus tangential compatibility on remaining faces	Same structure
C5: Orientation consistency	ε_σ fixed	$\varepsilon_{\sigma'}$ fixed
C6: Distinguishability	$s_\sigma \in \{+1, -1\}$ independent	$s_{\sigma'} \in \{+1, -1\}$ independent
C7: Bulk redundancy elimination	Removes one spurious d.o.f.	Same

The system is fully determined: matching fixes the shared variable, closure fixes one internal variable per simplex, orientation is determined, and each simplex retains one independent binary label. The boundary data uniquely determines the admissible bulk extension (up to the binary commitment), and the microstate count is well-defined.

Now remove one channel (K = 6). Consider dropping the distinguishability commitment C6. Without C6, the binary label s_σ is no longer required to be independently specified. This creates two distinct pathologies depending on the global context:

Pathology 1 — Gauge ambiguity in the microstate count. If s_σ is not required to carry independent information, then configurations differing only in s_σ may or may not be physically distinct — the theory does not decide. The microstate count Ω becomes dependent on an interpretive choice (are $s_\sigma = +1$ and $s_\sigma = -1$ the same state or different states?), which is precisely the gauge-like ambiguity that the distinguishability channel is designed to prevent.

Pathology 2 — Collapse of the configuration space. Alternatively, if the missing channel causes s_σ to be determined by the remaining constraints (because there are now fewer independent conditions protecting it), then σ contributes zero independent degrees of freedom. The configuration space collapses: the boundary no longer supports independent commitments, and the entropy density is trivially zero. This is the discrete analogue of a system with more equations than unknowns after one symmetry-breaking condition is removed.

Explicit counting for the two-simplex patch:

At K = 7: Variables = $\{x_a, x_b, x_a', x_b', s_\sigma, s_{\sigma'}\}$ (6 variables). Independent constraints: matching (1) + closure on σ (1) + closure on σ' (1) + orientation on σ (1) + orientation on σ' (1) + redundancy elimination (absorbed into defining the independent set) = 5 constraints on 6 variables, leaving 1 free variable per simplex after the shared face is accounted for. Each simplex retains one binary commitment. Total admissible configurations: $2^2 = 4$.

At K = 6 (drop C6): The binary labels s_σ, $s_{\sigma'}$ are no longer constrained to be independently meaningful. The remaining 5 constraints on the 4 continuous variables $\{x_a, x_b, x_a', x_b'\}$ yield a system that is generically overdetermined (5 equations, 4 unknowns, rank ≥ 4). The binary labels either float freely (gauge ambiguity: Ω is undefined) or are determined by the continuous sector (collapse: $\Omega = 1$). Neither outcome yields a well-defined, nontrivial microstate count.

This demonstrates concretely that K = 6 fails to support a well-defined boundary entropy.

Alternative K = 6 failure (dropping a compatibility channel). If instead one neighbor compatibility constraint is dropped (say, the tangential matching on one face), then two adjacent simplices are no longer required to agree on their shared-face data. The boundary admits configurations where adjacent tiles carry contradictory face assignments. In the language of the main paper, this is an underdetermination failure: multiple inequivalent bulk completions share the same macroscopic boundary data, because the missing matching condition allows internal ambiguity. The entropy acquires a spurious contribution from unphysical (non-gluing) configurations, and the microstate count is inflated beyond the physical value.

1.3 Why K > 7 Fails: Overdetermination and Loss of Robustness

At K > 7, additional independent constraints must be imposed beyond the closure + embedding + non-degeneracy set. We demonstrate the failure at K = 8.

Setup. Suppose an eighth independent constraint C8 is imposed on each simplex. Since the seven channels of K = 7 already form a complete generating set — meaning they determine all local admissibility conditions and gluing requirements — C8 must be algebraically independent of the existing seven.

Consequence. An independent constraint that is not a consequence of the existing algebra imposes a new restriction on the configuration space. For the two-simplex patch, the K = 7 system already has rank equal to the number of continuous variables minus one binary degree of freedom per simplex. Adding C8 introduces one additional equation per simplex. The constraint system becomes:

Rank = (original rank) + 2 (one per simplex) on the same variable set.

For generic face data, the augmented system has no solution: the intersection of constraint surfaces in configuration space is empty on a set of positive measure. Admissible configurations survive only on a measure-zero tuned subset of boundary data.

This is the overdetermination pathology: the theory predicts that generic boundary configurations are inadmissible, which contradicts the physical requirement that horizons of arbitrary (macroscopic) area exist and carry entropy.

Robustness failure. Even if admissible configurations exist at K = 8 for special boundary data, arbitrarily small perturbations of the boundary data generically move the system off the admissible set. Admissibility becomes non-robust: it depends on fine-tuning, violating the physical requirement that the coherence predicate is stable under small perturbations.

1.4 Summary of the Exhaustion Logic

The exhaustion argument has the following structure:

1. Enumerate the independent functional roles that constraints must serve: local consistency, face matching, orientability, non-degeneracy, and redundancy quotienting.
2. Show that each role requires at least one independent constraint channel (Sections 1.1–1.2).
3. Show that no role requires more than the channels assigned to it, because additional channels within any role are either algebraically dependent on the existing ones or introduce overdetermination (Section 1.3).
4. Conclude that K = 7 is the unique value satisfying completeness, consistency, and robustness simultaneously.

The full argument, including the treatment of higher-dimensional simplices and non-boundary contexts, is given in [20]. The sketch above demonstrates the logic on the minimal boundary patch and exhibits concrete failure modes at K = 6 and K = 8.

---

2. Canonical Status of the Seven-Constraint Taxonomy

A potential concern is that "bulk redundancy elimination" and "distinguishability commitment" appear meta-structural rather than geometric. This section clarifies their algebraic role and establishes that the seven-channel decomposition is canonical in a precise sense.

2.1 Constraints versus Quotient Conditions

In discrete gauge theories, it is standard to distinguish two operations on configuration space:

Constraints are equations that restrict which representatives are admissible (e.g., closure, matching). They reduce the dimension of the configuration space.

Quotients are equivalence relations identifying representatives that describe the same physical state (e.g., gauge redundancy). They reduce the dimension of the *physical* configuration space without changing the admissible set.

A well-defined state count requires both: constraints select admissible configurations, and quotients identify physically equivalent ones. The independent degree-of-freedom count is:

d.o.f. = (variables) − (constraint rank) − (quotient rank).

The "redundancy elimination" channel (constraint #7) belongs to the quotient category: it removes one spurious degree of freedom that would otherwise inflate the naive constraint count. It is not a subjective bookkeeping trick but a required algebraic operation for defining the physical configuration space. Omitting it produces a gauge-dependent state count — precisely the pathology that makes entropy ill-defined.

2.2 Distinguishability Commitment as a Non-Degeneracy Axiom

The distinguishability commitment (constraint #6) is a *non-degeneracy condition on the configuration space*. In formal terms: it requires that the quotient configuration space (after imposing all constraints and gauge identifications) has dimension ≥ 1 per simplex. Equivalently, it requires that each simplex contributes at least one independent, stably distinguishable label not determined by its neighbors.

This is not a constraint "on the theory's ability to count" — it is a constraint on configurations. It excludes configurations in which a simplex is entirely determined by its boundary data, contributing no independent microstate. In algebraic terms, it prevents the physical configuration space from collapsing to a single equivalence class under gluing, which would make the entropy trivially zero.

The analogue in continuum field theory is the requirement that the field equation admits a non-trivial space of solutions modulo gauge. No one regards this as "meta-structural" — it is a property of the equation, not of the theorist's bookkeeping.

2.3 Minimality and Canonical Rank

The claim "$K = 7$ is canonical" means: among local constraint systems on a simplicial foam that satisfy

(i) nontrivial admissible configurations exist for generic boundary data, (ii) admissible configurations extend uniquely to bulk completions (determinacy), (iii) the admissible set is robust under small perturbations, and (iv) the quotient configuration space is well-defined and non-degenerate,

the number of independent constraint channels is exactly seven.

Alternative decompositions of the seven channels (e.g., merging two channels into a composite and splitting another) are possible but must be related to the canonical decomposition by an invertible linear recombination. The total rank is invariant under such recombinations. $K = 7$ is a rank, not a labeling.

---

3. Exact Pairing of Tangential Constraints: Structural Conditions

The entropy paper uses a pairing argument: tangential constraints on either side of a shared face are "the same constraint counted twice." This section states the conditions under which this is exact and characterizes the corrections when it is not.

3.1 Matching as Exact Identification

Let σ and σ' be adjacent boundary simplices sharing a face $f = \sigma \cap \sigma'$. Let $X_f^{(\sigma)}$ denote the face data as represented in simplex σ, and $X_f^{(\sigma')}$ the same face data as represented in σ'. The structural matching postulate is:

Exact matching: $X_f^{(\sigma)} + X_f^{(\sigma')} = 0$ (opposite orientations absorbed into the sign convention),

or equivalently $X_f^{(\sigma)} = -X_f^{(\sigma')}$.

Under exact matching, any constraint that depends on X_f only through the shared-face variable is automatically imposed on both sides: imposing it on σ fixes it on σ'.

3.2 Necessary and Sufficient Conditions for Pairing

Exact pairing — the elimination of a tangential constraint as an independent degree of freedom on both sides of a shared face — holds if and only if:

(P1) Shared-face identification is exact. The matching condition is an equality constraint with no slack variables: $X_f^{(\sigma)} + X_f^{(\sigma')} = 0$ exactly, not approximately.

(P2) The constraint depends only on shared-face data. The tangential constraint component in question is a function of X_f alone, with no dependence on simplex-internal data that differs between σ and σ' (no hidden internal reference frames).

If either condition is violated, pairing becomes approximate:

Violation of P1 (approximate matching). If $X_f^{(\sigma)} + X_f^{(\sigma')} = \delta_f$ for some mismatch variable $\delta_f \neq 0$, then the two tangential constraints are no longer identical. The paired constraint gains an independent correction proportional to δ_f, and the entropy density acquires a subleading correction controlled by the statistics of δ_f across the boundary.

Violation of P2 (internal-frame dependence). If the constraint depends on simplex-internal data beyond the shared face, then even exact matching of face data does not identify the constraints on opposite sides. Each simplex retains an independent constraint, and the pairing reduction fails for that channel.

In the discrete admissibility framework, both P1 and P2 are structural postulates — they follow from the definition of the simplicial foam as a complex with exact face identifications and locally defined constraints. They are not additional assumptions imposed to make the counting work; they are consequences of the framework's foundational axioms.

Robustness of the leading coefficient. The leading area-law coefficient is determined by the exact-matching framework (P1 and P2 satisfied). Corrections from approximate matching or internal-frame effects enter as subleading terms, potentially contributing to logarithmic or sub-area corrections discussed in the main paper's Section 7.1. The leading 1/4 coefficient is stable under small violations of exactness.

---

4. The $\mathbb{Z}_2$ Redundancy: Local Sign Gauge with Residual Global Symmetry

A critical question is whether the sign flip $s \to -s$ that removes one constraint is a global or local operation. The correct interpretation involves two stages: a local sign gauge that defines the per-simplex constraint rank, and a residual global symmetry that emerges after matching.

4.1 Local Representative Sign Gauge

In the framework, physical face data live in equivalence classes: X_f and $-X_f$ represent the same physical configuration (orientation reversal of a flux-type variable). The binary orientation label $s \in \{+1, -1\}$ encodes this sign choice. Selecting a specific representative from each equivalence class — choosing X_f rather than $-X_f$ — is a *local* gauge choice made independently at each simplex.

This local gauge freedom is what is counted in defining K = 7. At each simplex, one degree of freedom is removed because the sign of the representative is unphysical. This is not a constraint on configurations but a quotient identifying equivalent representatives:

Local gauge: $X_f \to -X_f$, $s \to -s$ (per-simplex representative choice).

The local gauge removal is part of the bulk constraint algebra and enters the K = 7 count via the "bulk redundancy elimination" channel (#7). It ensures that the seven independent channels correctly count physical degrees of freedom, not gauge copies.

4.2 Effect of Matching on the Gauge Group

When matching constraints are imposed on a connected boundary patch, the local sign choices on adjacent simplices become correlated. Matching requires $X_f^{(\sigma)} = -X_f^{(\sigma')}$, which ties the representative sign on σ to the representative sign on σ′. Propagating this across all shared faces on a connected boundary, the originally independent local sign choices reduce to a single global choice: the overall sign convention for the entire connected component.

After matching on a connected boundary:

- All local sign gauges are fixed relative to each other by the matching constraints.
- A single residual global $\mathbb{Z}_2$ symmetry remains: flipping all signs simultaneously.

- This residual global symmetry removes one additional degree of freedom from the *entire connected boundary*, not one per simplex.

4.3 Reconciling Per-Simplex and Global Counting

The two stages are logically distinct and must not be conflated:

Stage 1 (defining $K = 7$): Each simplex independently has its constraint rank reduced by one through local gauge fixing. This is part of the bulk algebra and determines $K = 7$. It occurs *before* matching is imposed and is independent of boundary topology.

Stage 2 (boundary reduction): After matching on a connected null boundary, the surviving degrees of freedom are further reduced by the residual global $\mathbb{Z}_2$. For a boundary with N independent tiles, this removes 1 degree of freedom from the total count, giving $N - 1$ rather than N independent commitments.

For macroscopic horizons ($N \to \infty$), the correction is $O(1/N)$ and negligible: the entropy is $S = N \ln 2 - \ln 2 \approx N \ln 2 = (A/4\ell_p^2) \ln 2$, with the $-\ln 2$ term absorbed into the subleading (constant) correction. The leading coefficient 1/4 is unaffected.

This resolves the apparent tension: the per-simplex counting in the $K = 7$ definition uses local gauge, while the boundary reduction uses the residual global gauge. Both are correct at their respective stages, and they produce consistent results for the leading entropy.

---

5. From Graph Distance to Geometric Area: $\ell_e = 2\ell_p$ and $A\_cell = 4\ell_p^2$

The entropy coefficient requires two ingredients: (i) one surviving binary commitment per independent boundary tile, and (ii) the geometric area of an independent tile. The first follows from boundary constraint reduction (Sections 3–4 of the main paper). The second requires a coarse-graining map from the constraint adjacency graph to emergent geometry.

5.1 Independence as a Constraint-Graph Property

Define the *constraint adjacency graph* G whose vertices are boundary simplices and whose edges represent shared-face matching constraints. Two boundary commitments are independent when there is no constraint path in G linking their associated degrees of freedom.

On a nearest-neighbor simplicial tiling, a single adjacency step implies a shared face and hence a direct constraint coupling via $M\_f = 0$. Consequently:

- At graph distance $d = 1$: the two simplices share a face, and their commitments are coupled. They are *not* independent.

- At graph distance d = 2: the two simplices share no face. Their commitments are coupled only indirectly, through an intermediate simplex. However, the intermediate simplex's commitment is itself fixed by its own constraint reduction; it does not propagate a correlation between the two endpoints. The commitments are independent.

The minimal adjacency distance for independence is therefore d_ind = 2.

5.2 Coarse-Graining Map: Graph Distance to Proper Length

To convert the graph distance d_ind = 2 to a proper length, we require a relationship between adjacency steps and emergent geometry. This relationship involves three elements:

Element 1: Planck-scale lattice spacing. In a simplicial foam modeling Planck-scale spacetime structure, the characteristic edge length is ℓ_p. This is the fundamental assumption of Planck-scale discreteness — the same assumption that underlies Regge calculus [17], dynamical triangulations [19], and spin-foam models. It is not an additional postulate of the present framework.

Element 2: Linear coarse-graining. At leading order, graph distance maps linearly to proper distance: d adjacency steps correspond to $d \cdot \ell_p$ in proper length. This is the simplest coarse-graining and is exact for regular lattices. For irregular foams, it holds on average over sufficiently large patches — precisely the regime relevant for macroscopic horizons.

Element 3: The independence length. Combining Elements 1 and 2:

$\ell_e = d\_ind \cdot \ell_p = 2\ell_p$.

This is the minimal proper separation at which boundary commitments decouple.

5.3 From Independence Length to Cell Area

The effective boundary cell area — the area of the minimal independent tile — is determined by the independence length ℓ_e and the geometry of the boundary tiling.

For a boundary tiling with characteristic cell shape, the cell area is:

$A\_cell = c\_geom \cdot \ell_e^2$,

where c_geom is a dimensionless geometry factor of order unity determined by the tiling type:

Tiling	c_geom	A_cell
Square	1	$4\ell_p^2$
Equilateral triangular	$\sqrt{3}/4 \approx 0.43$	$\sqrt{3}\ \ell_p^2 \approx 1.73\ell_p^2$
Hexagonal	$3\sqrt{3}/2 \approx 2.60$	$6\sqrt{3}\ \ell_p^2 \approx 10.4\ell_p^2$

The entropy paper adopts c_geom = 1, which corresponds to the statement that the effective independent boundary cell is characterized by a square patch of side ℓ_e. The following argument supports this choice within the framework:

Effective coarse-graining and the square convention. The boundary cell is not a literal geometric tile with a fixed shape; it is an effective region defined by constraint independence. The "area per independent commitment" is a coarse-grained quantity obtained by dividing the total boundary area by the number of independent commitments. For a regular simplicial tiling with independence distance d_ind = 2, the number of independent commitments per unit area scales as $1/\ell_e^2$ regardless of the underlying tile shape, because the constraint graph's independence structure is determined by adjacency distance, not by the embedding geometry of individual simplices.

To see this explicitly: on any regular 2D tiling with lattice spacing ℓ_p, the number of sites at mutual graph distance ≥ 2 in a region of area A is A/ℓ_e^2 up to boundary corrections, independent of whether the tiling is triangular, square, or hexagonal. The shape of the Voronoi cell around each independent site differs, but the *density* of independent sites is universal.

This universality fixes A_cell = $\ell_e^2 = 4\ell_p^2$ as the correct effective cell area, with c_geom = 1 understood as a statement about the density of independent commitments rather than the shape of a literal tile.

5.4 Scope and Precision of the Area Assignment

To be fully explicit about the status of this result:

The framework determines, without free parameters:

- The independence graph distance: d_ind = 2 (from the constraint structure).
- The independence length: $\ell_e = 2\ell_p$ (from Planck-scale discreteness + linear coarse-graining).
- The density of independent commitments: 1 per ℓ_e^2 = 1 per $4\ell_p^2$ (from the universality argument above).

The resulting cell area A_cell = $4\ell_p^2$ yields the entropy:

$S = (A / 4\ell_p^2) \ln 2$,

reproducing the Bekenstein–Hawking coefficient 1/4. The linear coarse-graining assumption (Element 2) is the only step that is not derived from the constraint algebra itself; it is a standard assumption shared with all discrete approaches to quantum gravity and is exact for the regular lattices used in the framework.

---

6. Why Null Boundaries Are Essential: Non-Null Comparison

Null geometry supplies a structurally essential step in the 7 → 1 boundary reduction: the freezing of the normal closure component. This section shows explicitly how the constraint counting changes on non-null boundaries, confirming that the 1/4 coefficient is specific to null horizons.

6.1 Null Boundary (Event Horizon)

On a null boundary, the normal direction is degenerate: the null normal vector lies in the tangent plane of the surface. Consequently:

- The normal component of the closure constraint $C_{\sigma\perp}$ carries no independent degree of freedom. It is automatically satisfied given admissible tangential boundary data.
- This freezes 1 constraint channel per simplex.

Combined with the 4 channels eliminated by tangential pairing and the 1 channel removed by gauge redundancy, the null boundary reduces 7 bulk constraints to 1 surviving binary commitment. This produces the entropy density:

s_{null} = 1 bit per $4\ell_p^2 \rightarrow S = A / 4\ell_p^2$ (in bits).

6.2 Spacelike Boundary

On a spacelike boundary, the normal direction is timelike and independent of the tangential directions. The normal closure component $C_{\sigma\perp}$ is *not* frozen — it remains an independent constraint that must be satisfied by the boundary data.

The boundary reduction therefore proceeds differently:

Constraint type	Null fate	Spacelike fate
Normal closure	Frozen (0 surviving)	Active (1 surviving)
Tangential closure (×2)	Paired (0 surviving)	Paired (0 surviving)
Normal compatibility	Subsumed (0 surviving)	Independent (1 surviving)
Tangential compatibility (×2)	Paired (0 surviving)	Paired (0 surviving)
Orientation + distinguishability	1 surviving after gauge	1 surviving after gauge

On a spacelike boundary, 3 independent constraints survive per effective cell (normal closure, normal compatibility, and the binary commitment), compared to 1 on a null boundary. The additional surviving constraints do not contribute to entropy in the same way — they impose conditions on the boundary data rather than supporting independent commitments. The entropy density on a spacelike boundary is framework-dependent and does not generically reproduce the Bekenstein–Hawking formula.

This is physically correct: spacelike boundaries are not horizons. There is no causal information-hiding, and the thermodynamic interpretation of boundary entropy as observer-inaccessible microstates does not apply.

6.3 Timelike Boundary

On a timelike boundary, the normal direction is spacelike and independent. As with the spacelike case, the normal closure does not freeze. Additionally, a timelike boundary does not function as a causal barrier: information can propagate across it in both directions. The projection argument — that bulk degrees of freedom are inaccessible and only boundary-supported configurations contribute to entropy — does not apply.

The $7 \to 1$ reduction and the resulting 1/4 coefficient are therefore specific to null boundaries. This is consistent with the physical expectation: the Bekenstein–Hawking formula applies to event horizons, which are null surfaces.

---

7. Staging of the Toy Example and Constraint Matrix

Appendix C of the main paper writes an explicit 4×8 matching matrix M for a four-simplex boundary patch and computes its rank (4) and null-space dimension (4). A potential misunderstanding is that this matrix should encode all reductions from the naive $4 \times 7 = 28$ bulk constraints to 4 surviving commitments. It does not. The overall reduction proceeds in stages, and the matrix captures one stage.

7.1 Staging of the Boundary Reduction

Stage 1 — Null freezing. Remove the normal closure channel from each simplex. This eliminates 4 constraint slots (one per simplex), reducing the system from 28 to 24 effective constraints. This stage is applied *before* writing the tangential constraint system and is not represented in the matrix.

Stage 2 — Normal compatibility absorption. Remove the normal compatibility channel from each simplex (subsumed by the frozen normal closure). This eliminates 4 more, reducing to 20 effective constraints on the tangential variables.

Stage 3 — Tangential matching (captured by the matrix). The remaining tangential closure and compatibility constraints are encoded in the matching matrix M acting on the 8 tangential variables (2 per simplex). The matrix has rank 4, so 4 of the 8 variables are determined by matching. The null space has dimension 4, corresponding to 4 independent tangential degrees of freedom — one per simplex.

Stage 4 — Gauge quotienting. The 4 independent tangential degrees of freedom are subject to the local $\mathbb{Z}_2$ representative gauge (Section 4). Each becomes a binary commitment. On a connected patch, the residual global $\mathbb{Z}_2$ removes 1 from the total count, giving 3 independent commitments for 4 simplices. For macroscopic boundaries ($N \gg 1$), this correction is negligible: $N - 1 \approx N$.

The matrix in Appendix C therefore verifies Stage 3 — the tangential matching reduction — in isolation. The full 28 → 4 reduction is the composition of all four stages. Stating this staging explicitly prevents the misreading that the 4 × 8 matrix should directly encode all constraint eliminations.

---

8. Summary of Conditional Dependencies

For clarity, we state the complete set of inputs on which the 1/4 coefficient depends:

Input 1: K = 7 (constraint count per simplex). Derived in [20] from closure consistency and exhaustive elimination of alternatives. Sections 1–2 of this note provide a checkable sketch with explicit failure modes at K = 6 and K = 8.

Input 2: $\ell_e = 2\ell_p$ (independence threshold). Derived in [21] from the minimal graph distance for constraint decoupling (d_ind = 2) combined with Planck-scale lattice spacing. Section 5 of this note provides the coarse-graining argument and the universality of the effective cell area.

Input 3: Linear coarse-graining (graph distance maps linearly to proper length). This is a standard assumption in discrete quantum gravity, shared with Regge calculus and dynamical triangulations. It is exact for regular lattices and holds on average for irregular foams in the macroscopic limit.

Structural postulates: Exact face matching (P1), constraints depend only on shared-face data (P2), and the boundary tiling admits a well-defined constraint adjacency graph. These are consequences of the simplicial foam axioms, not additional assumptions.

Given these inputs, the derivation produces:

- 7 bulk constraints per simplex → 1 surviving boundary commitment (Section 3 of main paper),
- Effective cell area $A\_cell = 4\ell_p^2$ (Section 4 of main paper),
- Entropy $S = (A / 4\ell_p^2) \ln 2$ with the coefficient 1/4 fixed.

No parameter tuning, no holographic postulate, and no matching to known results is required.

References

References [1]–[29] are as in the main paper. This note references specifically:

- [17] T. Regge, "General relativity without coordinates," Nuovo Cimento 19, 558–571 (1961).
- [19] J. Ambjørn, J. Jurkiewicz, and R. Loll, "Dynamically triangulating Lorentzian quantum gravity," Nuclear Physics B 610, 347–382 (2001).
- [20] K. Taylor, "Closure consistency and constraint counting in simplicial quantum gravity," VERSF Theoretical Physics Program (2025).
- [21] K. Taylor, "The two-Planck scale: Distinguishability limits and correlation lengths in discrete spacetime," VERSF Theoretical Physics Program (2025).