Deriving the Bekenstein–Hawking Entropy Coefficient from Boundary Constraint Counting in a Discrete Admissibility Framework

Keith Taylor *VERSF Theoretical Physics Program*

General Reader Summary

Black holes are among the most extreme objects in the universe. In the 1970s, physicists Jacob Bekenstein and Stephen Hawking made a startling discovery: black holes have *entropy*—a measure of hidden internal disorder—and the amount of that entropy is proportional to the *surface area* of the black hole, not its volume. This was shocking because, for every other object in physics, entropy scales with volume. A box twice as big can hold twice as many disordered arrangements. Black holes break this rule.

Even more specifically, Bekenstein and Hawking showed that the entropy is one-quarter of the horizon area in Planck units (up to a choice of information units, bits vs nats). The formula is simple: $S = A / 4\ell_p^2$. The factor of 1/4 has been confirmed by multiple independent approaches, but nobody has been able to explain *why* it is 1/4 and not, say, 1/3 or 1/6, without inserting some adjustable number by hand.

This paper shows that the factor of 1/4 is not arbitrary—it is an unavoidable consequence of how space is built at the smallest scales. If space is made of tiny discrete building blocks (like atomic-scale tiles), and those tiles follow a specific set of internal consistency rules, then the surface of a black hole can only store information in a very particular way. Most of the internal rules of each tile become redundant or locked at the surface; only one binary choice—yes or no, on or off—survives per tile. And the natural size of each tile turns out to be exactly four Planck areas. Divide the total surface area by four Planck areas, and you get the number of independent yes/no choices. That count *is* the entropy, measured in bits. The 1/4 was never a free choice—it was baked into the structure of space itself.

Crucially, the rules that determine the tile size and the number of internal constraints were established in earlier work that had nothing to do with black holes. This paper applies those pre-existing rules to a black hole horizon and finds that the famous 1/4 falls out automatically. No knobs were turned. No numbers were adjusted to match the known answer.

Abstract

The Bekenstein–Hawking entropy formula $S = A / 4\ell_p^2$ is one of the most robust results in theoretical physics, yet the precise numerical coefficient 1/4 has proven notoriously difficult to derive from first principles. Most approaches recover the area scaling but require adjustable parameters, tuning, or auxiliary assumptions to fix the coefficient. In this paper, we establish the following conditional result: *if* the closure-consistency analysis of companion work fixes the constraint count at $K = 7$ per simplex, and *if* the decoupling-scale analysis fixes the independence threshold at $\ell_e = 2\ell_p$, *then* the Bekenstein–Hawking area coefficient 1/4 follows from boundary constraint reduction on null horizons. Space is modeled as a simplicial foam whose degrees of freedom are governed by these closure constraints. We demonstrate that, when restricted to a horizon-bounding surface, bulk constraints project out through a mechanism we enumerate explicitly: of the seven bulk constraints per simplex, one is frozen by causal structure at the null boundary, four are eliminated by closure pairing across adjacent boundary simplices, and one is removed by gauge redundancy, leaving exactly one independent binary commitment per effective boundary cell. The fundamental boundary cell area is shown to be $4\ell_p^2$, arising from the two-Planck independence threshold—the minimal simplex separation at which boundary commitments decouple given shared-face matching constraints. The framework therefore fixes the geometric cell count $N = A / 4\ell_p^2$, establishing the 1/4 coefficient. In natural units (nats), $S = N \ln 2$; in bits, $S_{\beta its} = N$. The coefficient is fixed independently of logarithm base. The present paper provides a self-contained derivation of the boundary reduction and area-coefficient implication; independent validation of the inputs $K = 7$ and $\ell_e = 2\ell_p$ is addressed in companion papers. This conditional derivation places the Bekenstein–Hawking coefficient on the same axiomatic footing as the cosmological constant within this framework.

Table of Contents

---

1. Introduction

For the general reader. When you drop a book into a black hole, all the information about that book—its words, its pages, its atoms—seems to vanish behind the event horizon. But physics says information cannot be destroyed. So where does it go? Bekenstein and Hawking showed that the black hole's *surface* keeps a record: every bit of swallowed information is encoded on the horizon, like writing on the skin of a balloon. The amount of information the surface can hold is measured by entropy, and their formula says it equals exactly one-quarter of the surface area in Planck units. This paper explains where that "one-quarter" comes from.

The discovery that black hole entropy scales with horizon area rather than volume was a profound shift in our understanding of gravitational thermodynamics [1]. The Bekenstein–Hawking result,

$S = A/4\ell_P^2$,

has since been confirmed across semiclassical gravity [2,3], Euclidean path integrals [4], string theory microstate counting [5], and loop quantum gravity [6,7]. The Wald entropy formula further generalizes the result to arbitrary diffeomorphism-invariant theories via Noether charge methods [16]. Despite this success, a persistent foundational issue remains: while the proportionality $S \propto A$ is generic, the numerical coefficient 1/4 is not [25].

In most approaches, the coefficient emerges only after the introduction of auxiliary assumptions—such as a choice of ultraviolet cutoff [8,9], an adjustable Immirzi parameter [7], or an implicit normalization of microscopic degrees of freedom [10]. This raises the question of whether the coefficient is truly fundamental, or merely an artifact of model-dependent bookkeeping.

In this paper, we establish the following conditional result:

Theorem (Conditional). *If* the closure-consistency analysis of [20] fixes the constraint count at K = 7 per simplex, and *if* the decoupling-scale analysis of [21] fixes the independence threshold at $\ell_e = 2\ell_P$, *then* the Bekenstein–Hawking area coefficient 1/4 follows from boundary constraint reduction on null horizons.

The derivation assumes a discrete admissibility framework in which:

1. The microscopic structure of space is simplicial.
2. The number of independent constraints per simplex is K = 7 (taken as input from [20]).
3. Entropy is identified with the count of admissible configurations.
4. Horizons are treated as constraint-bounding surfaces rather than dynamical objects.

No holographic postulate [10,13,14], entanglement entropy ansatz [8,9], or parameter tuning is assumed. The coefficient 1/4 arises purely from boundary constraint reduction and the two-Planck independence threshold.

Status. The present paper provides a self-contained derivation of the boundary reduction and the area-coefficient implication. The independent validation of the inputs K = 7 and $\ell_e = 2\ell_p$ is addressed in companion papers [20,21]. The strength of the present result is therefore conditional on those inputs; readers who accept the companion derivations obtain a parameter-free prediction, while those who do not can treat this paper as establishing a precise logical dependence.

---

2. Discrete Structure and Constraint Admissibility

We begin by summarizing the minimal structural elements required for the derivation.

For the general reader. This section introduces the idea that space is not smooth and continuous at the smallest scales, but instead built from tiny, indivisible building blocks—like how a mosaic is made of individual tiles. Each tile has to follow certain rules to fit properly with its neighbors. The number of rules per tile turns out to be exactly seven, determined by pure logic rather than measurement. This section explains what those rules are and why there must be exactly seven.

2.1 Simplicial Foam and Constraint Closure

Space is modeled as a simplicial foam composed of elementary cells (simplices) [17,18,19]. Each simplex carries a fixed number of internal constraints governing its admissible configurations. These constraints encode closure, compatibility with neighboring simplices, and internal consistency of distinguishable states.

For the general reader. A "simplex" is the simplest possible shape in a given number of dimensions: a triangle in 2D, a tetrahedron in 3D. A "simplicial foam" means space is tiled entirely by these simple shapes, packed together like bubbles. Each tile must satisfy internal rules ("constraints") to be a valid piece of space—just as a jigsaw puzzle piece must have the right shape to fit its neighbors.

Previous work [20,23,24] has shown that closure consistency uniquely fixes the number of independent constraints per simplex to

K = 7,

with no free parameters. This number is not postulated but follows from internal admissibility requirements: fewer constraints lead to underdetermined configurations, while additional constraints over-restrict the system and prevent consistent gluing.

2.2 Summary of Constraint Closure and the Origin of K = 7

Because the derivation of K = 7 is the foundation on which the present result rests, we provide an explicit summary of its logic. The full derivation appears in [20]; here we enumerate the constraint types and the argument that fixes their count.

Each simplex in the foam must satisfy the following independent constraints for admissible participation in a global tiling:

#	Constraint type	Role
1	Local closure	Ensures internal geometric consistency of the simplex
2	Neighbor compatibility (face 1)	Matching conditions on the first shared face
3	Neighbor compatibility (face 2)	Matching conditions on the second shared face
4	Neighbor compatibility (face 3)	Matching conditions on the third shared face
5	Orientation / chirality consistency	Fixes handedness relative to the global tiling
6	Distinguishability commitment	Ensures the simplex contributes an independent microstate
7	Bulk redundancy elimination	Removes one spurious degree of freedom from the bulk closure set, yielding 7 independent constraints per simplex

For the general reader. Think of each tile as having a checklist of seven rules it must pass: (1) its own internal shape must be self-consistent; (2–4) it must match properly with each of its three neighboring tiles; (5) it must agree on which way is "left" and which is "right" relative to the global pattern; (6) it must be distinguishable from its neighbors—it has to represent a genuinely different piece of information; and (7) one apparent rule is actually redundant and gets removed, like realizing two items on a to-do list are really the same task.

The count K = 7 is fixed by an exhaustive consistency argument:

- **K < 7**: The constraint set is incomplete. Configurations admit non-unique extensions, meaning the same macroscopic boundary data corresponds to multiple incompatible bulk completions. The theory is underdetermined.
- **K = 7**: The constraint set is complete and consistent. Every admissible boundary configuration extends to a unique bulk completion, and simplices can be glued into arbitrary global tilings without contradiction.

- **K > 7**: The constraint set is overdetermined. No admissible configurations survive for generic boundary data, and global tilings are forbidden.

For the general reader. Fewer than seven rules and the tiles are too loosely constrained—multiple incompatible arrangements could all claim to be valid, like a jigsaw with pieces that fit in more than one place. Exactly seven and everything fits perfectly. More than seven and the rules are so strict that no valid arrangement exists at all—the puzzle becomes unsolvable.

The present paper does not re-derive K = 7; it uses this as an established structural input and tests whether black hole entropy follows as a consequence.

2.3 Entropy as State Count

Entropy is defined in the standard statistical sense:

$S = \ln \Omega$,

where Ω is the number of admissible microscopic configurations consistent with macroscopic boundary conditions.

Crucially, entropy is not introduced as an emergent thermodynamic quantity but as a primary measure of distinguishability capacity. Geometry and dynamics arise from the admissibility structure, not the other way around.

For the general reader. Entropy counts the number of different microscopic arrangements that all look the same from the outside. A shuffled deck of cards has high entropy because there are trillions of possible orderings that all look like "a deck of 52 cards." A sorted deck has low entropy—there's essentially only one arrangement. Here, entropy counts the number of different ways the tiny tiles of space could be arranged while producing the same large-scale black hole.

2.4 Clarification: Folds, Boundary Tiles, and Dimensional Projection

Because the present derivation combines concepts introduced in earlier work with new boundary-specific constructions, we clarify here the relationship between folds, boundary commitments, and effective boundary cells ("tiles"), and address a common potential misunderstanding.

For the general reader. In earlier papers we introduced "folds" as the most basic stable information-bearing units in the framework: each fold can ultimately be forced into a definite yes/no outcome, becoming one irreversible bit. Folds are not assumed to be little chunks of area. By contrast, "tiles" in this paper are geometric patches on a horizon surface: each tile is the smallest boundary region that still carries one independent yes/no outcome after the horizon removes (or locks) most of the internal freedom. This section explains why the paper distinguishes those two ideas.

2.4.1 Folds as Stable Binary Commitments with Irreversible Potential

In the underlying framework and in prior work [23,24], a *fold* is a minimal stable unit capable of supporting a binary commitment. In operational terms, a fold carries a binary orientation variable

$s \in \{+1, -1\}$,

whose collective phase contributions generate ordinary quantum amplitudes in the continuum limit.

Crucially, a fold should be understood as a unit of stable distinguishability capacity: under reversible dynamics it can participate coherently (without generating entropy), while under appropriate constraints it can undergo an irreversible binary commitment (creating one unit of classical distinguishability). Thus, it is correct to say that one fold corresponds to one binary capacity for commitment, but it is not necessary to identify a fold with any fixed geometric area.

2.4.2 Boundary Projection and Loss of Independent Degrees of Freedom

When a bulk region is bounded by a null horizon, the external observer loses access to bulk distinguishability. The horizon acts as a projection on admissible configurations:

- Normal (radial) degrees of freedom are frozen by null degeneracy.
- Tangential closure and compatibility constraints pair across adjacent boundary simplices.
- Gauge redundancies remove additional apparent degrees of freedom.

As a result, many fold-level commitments that can be treated as independent in the bulk become correlated, paired, or redundant when viewed from the boundary. Although folds retain their binary orientation internally, these orientations can no longer be specified independently at the horizon.

In this sense, the horizon "collapses" bulk information—not by physically compressing space, but by eliminating independent distinguishability through causal and constraint-based projection.

2.4.3 Boundary Tiles as Emergent Geometric Cells

After projection, the entropy relevant to an external observer is governed not by the number of folds, but by the number of independent irreversible boundary commitments that remain distinguishable. These surviving commitments are supported on emergent geometric patches of the horizon, which we call *effective boundary cells* or *tiles*.

A tile is therefore defined as:

The minimal horizon patch that supports one independent boundary commitment after constraint projection and redundancy elimination.

Tiles are geometric objects; folds are not assumed to be geometric area elements. A tile typically corresponds to many correlated folds, with only one independent commitment surviving at the boundary.

2.4.4 Correlation Length and Effective Cell Area

Independence of boundary commitments occurs only once correlations between candidate commitments decay. Prior work [21] establishes that the minimal separation at which boundary commitments decouple is the two-Planck independence threshold,

$\ell_e = 2\ell_p$.

On a two-dimensional boundary this implies an effective cell area

$A^{cell} = \ell_e^2 = 4\ell_p^2$.

This area characterizes the size of an independent boundary tile, not the size of a fold. The appearance of $4\ell_p^2$ therefore reflects constraint-adjacency coarse-graining and boundary projection, rather than an intrinsic geometric size assigned to folds.

2.4.5 Summary of the Distinction

To summarize:

- **Folds** are minimal stable units capable of becoming irreversible binary commitments.
- **Tiles** are emergent geometric boundary cells on a horizon.
- One tile supports exactly one independent boundary commitment.
- Many folds contribute to each tile through correlation and constraint sharing.
- The area $4\ell_p^2$ characterizes tiles, not folds.

With this clarification, the entropy derivation can be read unambiguously: the Bekenstein–Hawking coefficient arises from counting independent boundary tiles, each carrying a single surviving binary commitment, rather than from assigning geometric size to individual folds.

---

3. Horizons as Boundary Constraint Surfaces

The central step in the derivation is recognizing that a black hole horizon functions as a constraint-bounding surface, not a bulk region.

For the general reader. This section is the heart of the paper. Imagine the black hole's surface as a wall. You cannot see or access anything behind the wall, so only the tiles *on* the wall itself can contribute to the information you can detect. The question becomes: of the seven rules each tile normally follows, how many still matter when the tile sits on this wall? The answer turns out to be just one. This section shows exactly how the other six rules become irrelevant at the surface—frozen, paired off with neighboring tiles, or redundant.

3.1 Projection of Bulk Constraints

Consider a spacetime region bounded by a horizon surface $\mathscr{H}$. Bulk simplices interior to $\mathscr{H}$ are causally inaccessible to an external observer. As a result:

- Bulk constraints do not contribute to externally distinguishable configurations.
- Only constraints residing on, or intersecting, the boundary remain relevant.

This projection eliminates volumetric degrees of freedom and leaves a purely boundary-supported entropy, in agreement with the area law [10,13,14].

3.2 Explicit Boundary Constraint Reduction

Each boundary simplex inherits all seven bulk constraints from the underlying simplicial foam. However, not all inherited constraints remain independent on a horizon. Before tracing the fate of each constraint type, we introduce minimal notation to make "pairing" and "gauge redundancy" precise.

3.2.1 Schematic Constraint Definitions

Let σ be a boundary simplex with faces f. We define:

Face data. Each face f carries admissible data $X_f \in \mathcal{X}$, where $\mathcal{X}$ is the space of admissible face assignments (e.g., induced metric data, connection holonomy, or closure flux—the specific interpretation is not required for the counting argument).

Orientation sign. Each face has an orientation $\varepsilon_{\sigma f} \in \{+1, -1\}$ encoding whether f is outward or inward oriented relative to σ.

Closure constraint. The closure constraint on σ is

$$C_\sigma := \Sigma_{f \subset \sigma}\, \varepsilon_{\sigma f}\, X_f = 0.$$

On a null boundary, this splits into normal and tangential components:

$$C_\sigma = C_\sigma^\perp + C_\sigma^\parallel.$$

Matching constraint. On a shared boundary face f = σ ∩ σ', the matching constraint is

$$M_f(\sigma, \sigma') := X_f^{(\sigma)} - X_f^{(\sigma')} = 0,$$

with the identification $X_f^{(\sigma')} = -X_f^{(\sigma)}$ when opposite orientations are used. This is the precise statement that the "conjugate/opposite-side" constraints are not independent: they are the same constraint written from two perspectives.

Gauge redundancy. Let G act on face data by $X_f \mapsto g \cdot X_f$, leaving physical equivalence classes $[X_f]$. For the binary fold orientation $s \in \{+1, -1\}$, the minimal relabeling symmetry is $\mathbb{Z}_2$:

$X_f \sim -X_f$ ($\mathbb{Z}_2$ relabeling).

This sign redundancy is consistent with the fold orientation variable introduced in Section 2.4.

With these definitions, "pairing" means double-counting of M_f across shared faces, and "gauge removal" means quotienting by G.

3.2.2 Constraint Fate on Null Boundaries

We now trace the fate of each constraint type explicitly.

Normal closure (1 constraint). The local closure constraint involves components both tangential and normal to the boundary. On a null horizon, the normal direction is degenerate: the outgoing null generator is fixed by the causal structure, and no independent normal degree of freedom remains. In the notation above, $X_f^{\perp}$ is determined by tangential data and the horizon generator, so $C_\sigma^{\perp} = 0$ is automatically satisfied given admissible boundary conditions. The normal component of the closure constraint is therefore *frozen*—it contributes no distinguishable state.

Why null boundaries freeze the normal constraint: A null surface is generated by null geodesics whose tangent vectors are simultaneously tangent to and normal to the surface (the null normal is contained in the tangent plane). This degeneracy means the normal closure component carries no independent information; it is determined entirely by the tangential data. On a spacelike boundary, by contrast, the normal direction is independent and the constraint would not freeze.

For the general reader. One of the seven rules governs how the tile relates to the direction pointing "inward" (toward the black hole's center). But a black hole horizon is special: light there is trapped at the boundary, moving along the surface rather than away from it. The inward direction collapses into the surface itself. So the rule governing that direction is automatically satisfied—it contains no choice, no information. That is one rule eliminated.

Tangential closure (2 constraints). The two tangential components of closure pair across adjacent boundary simplices. In the notation above, each tangential closure constraint $C_\sigma^{\parallel}$ on simplex σ and the corresponding constraint $C_{\sigma'}^{\parallel}$ on adjacent simplex σ' are related by the matching constraint $M_f = 0$ on their shared face. They encode the same shared-face matching condition from opposite sides. This pairing eliminates both as independent degrees of freedom: knowing one determines the other.

Pairing mechanism: Two adjacent boundary simplices share a codimension-1 face $f = \sigma \cap \sigma'$. The tangential closure condition on each side of this face is the same geometric requirement (matching of induced metric data) viewed from opposite simplices: $X_f^{(\sigma)} = X_f^{(\sigma')}$ via $M_f = 0$. These are not independent constraints but a single constraint counted twice.

For the general reader. Two of the rules describe how each tile lines up with its neighbors along the surface. But each such rule is shared between two tiles—tile A's "match my right neighbor" rule is the same as tile B's "match my left neighbor" rule, just seen from the other side.

They look like two rules, but they are really one rule counted twice. Both get eliminated as independent choices.

Neighbor compatibility (3 constraints). Of the three neighbor compatibility constraints, one involves the normal direction and is subsumed by the frozen normal closure ($Xf\perp$ already determined). The remaining two tangential compatibility constraints pair across shared boundary faces by the same mechanism as tangential closure: each encodes a shared-face matching condition $Mf = 0$ that is a single constraint counted from both sides.

Gauge redundancy (1 constraint). The orientation/chirality constraint and the distinguishability commitment together contain one gauge redundancy: the $\mathbb{Z}_2$ relabeling symmetry $Xf \sim -Xf$ that does not change the physical configuration. Removing this redundancy eliminates one constraint. This is consistent with the fold orientation $s \in \{+1, -1\}$: the sign of s can be globally flipped without changing physical predictions.

For the general reader. The handedness rule and the distinguishability rule overlap: they share a symmetry, like the fact that labeling tiles "A, B, C" contains the same information as labeling them "1, 2, 3." Renaming the labels does not change the physics. Removing this redundancy eliminates one more rule.

Surviving commitment (1 constraint). After all reductions, exactly one independent binary commitment remains per boundary cell: the distinguishability commitment, which records whether the cell is in one admissible state or the other.

The full accounting is summarized below:

Constraint type	Bulk count	Boundary fate	Surviving
Normal closure	1	Frozen (null degeneracy)	0
Tangential closure	2	Paired across shared faces	0
Neighbor compatibility (normal)	1	Subsumed by frozen normal	0
Neighbor compatibility (tangential)	2	Paired across shared faces	0
Orientation + distinguishability	2	One removed by gauge redundancy	**1**
Total	**7**		**1**

This reduction—from seven bulk constraints to one boundary commitment—is the structural origin of the area law. It is not assumed but derived from the causal and geometric properties of null boundaries acting on the admissibility constraint set.

For the general reader. Of the seven rules each tile follows, six become locked, paired, or redundant when the tile sits on a black hole's surface. Only one survives: a single binary choice, like a coin flip—heads or tails, yes or no, 0 or 1. Each surface tile contributes exactly one bit of information. This is the deep reason why black hole entropy is proportional to surface area: the surface is tiled by these minimal one-bit cells.

3.3 Worked Example: Toy Boundary Patch

To illustrate the constraint reduction concretely, consider a minimal boundary patch consisting of four triangular simplices sharing edges on a null surface $\mathscr{H}$.

Label the simplices σ_1, σ_2, σ_3, σ_4, arranged so that σ_1 shares an edge with σ_2 and σ_3, while σ_4 shares edges with σ_2 and σ_3.

Naive counting: Each simplex carries 7 constraints, so the patch begins with $4 \times 7 = 28$ constraint slots.

Normal freezing: Each simplex loses 1 normal closure constraint to null degeneracy. This removes 4 constraints, leaving 24.

Tangential pairing: The patch contains 4 internal shared edges. Each shared edge eliminates 2 tangential constraints (one closure, one compatibility) that are double-counted across the edge. This removes $4 \times 2 = 8$ constraints, leaving 16.

Normal compatibility: Each simplex has 1 normal compatibility constraint subsumed by the frozen normal closure, removing 4 more. This leaves 12.

Tangential compatibility pairing: The same 4 internal edges each eliminate 1 additional tangential compatibility constraint by pairing, removing 4 more. This leaves 8.

Gauge redundancy: Each simplex has 1 gauge redundancy among its orientation and distinguishability constraints, removing 4 more. This leaves 4.

Result: The 4-simplex patch supports exactly 4 independent binary commitments—one per cell—confirming the general reduction.

The eliminations above are organized by constraint type to avoid double counting; each pairing removal corresponds to a single shared-face constraint counted twice in the naive per-simplex tally. More generally, for a boundary tiling with N boundary simplices and E internal shared edges, the paired constraints scale with E, leaving a net of one independent commitment per simplex in the null limit.

Note that naive bulk counting would assign 2^{28} configurations to this patch. The boundary reduction yields $2^4 = 16$ admissible configurations, demonstrating that the constraint machinery does substantial non-trivial work.

For the general reader. To make this concrete: take four tiny triangular tiles on the black hole's surface. Naively, with seven rules each, they could have 2^{28}—over 268 million—possible configurations. But after accounting for frozen, paired, and redundant rules, only $2^4 = 16$ configurations survive. The rules do enormous work, stripping away almost all of the apparent freedom. What remains is exactly one binary choice per tile.

The toy example serves as a consistency check of the general reduction rules. Appendix C provides an explicit constraint matrix whose null-space dimension reproduces this boundary degree-of-freedom count for the minimal horizon patch.

---

4. Fundamental Boundary Cell Area

The effective boundary cell area is not set by the Planck length ℓ_p alone, but by the intrinsic two-Planck independence threshold of the framework.

For the general reader. We now know that each surface tile carries one bit of information. But how big is each tile? If tiles were the size of a single Planck area (the smallest meaningful length in physics, squared), the entropy would be $S = A / \ell_p^2$—four times too large. This section shows why the tiles are actually four Planck areas in size, which gives the correct factor of 1/4. The reason is that neighboring tiles share rules, so tiles right next to each other are not truly independent. You have to space them out by twice the Planck length before they carry independent information—making each effective tile four Planck areas.

4.1 The Two-Planck Independence Threshold

The fundamental length scale governing independent boundary commitments is

$\ell_e = 2\ell_p$.

This scale arises not as a geometric edge length, simplex face size, or thermal correlation length, but as the *minimal graph distance at which boundary commitments become independent* given shared-face constraints.

4.1.1 Independence as a Graph-Distance Criterion

Define an "independent boundary commitment" as a degree of freedom not fixed by face-sharing constraints. Two boundary simplices σ and σ' share a face $f = \sigma \cap \sigma'$, and the matching constraint $M_f = 0$ couples their admissible states. Consequently:

- A single ℓ_p-separated patch (one adjacency step) shares at least one constraint with its neighbor. The two patches are therefore not independent: their boundary commitments are correlated through M_f.
- Independence first occurs at the smallest separation where there is no shared constraint path—i.e., no shared face and no shared closure coupling.

On a nearest-neighbor simplicial tiling, this occurs at **two adjacency steps**. The minimal graph distance for independence is therefore

$d_{in}^{d} = 2$ (adjacency steps).

Each adjacency step corresponds to one Planck length in the simplicial lattice, so the independence threshold is

$\ell_e = 2\ell_p$.

4.1.2 Why Not $\sqrt{2}\ell_p$ or $3\ell_p$?

A common objection is that Euclidean distance on a triangular lattice might introduce factors like $\sqrt{2}$ or $\sqrt{3}$. However, ℓ_e is not a Euclidean distance; it is an adjacency distance in the constraint graph. The relevant measure is the number of shared-face hops required to reach a given simplex, not the embedding geometry of the lattice. Euclidean factors arise from embedding; the constraint structure depends only on adjacency.

The answer to "why not $3\ell_p$?" is equally direct: at three adjacency steps, the simplices are still independent (no shared constraint path), but independence already holds at two steps. The independence threshold is the *minimum* graph distance for decoupling, which is two.

For the general reader. Think of tiles connected by shared rules as nodes in a network, with an edge between any two tiles that share a rule. Two tiles directly connected by an edge are correlated—they cannot make independent choices. To find two tiles that are truly independent, you need to step two edges away: the minimum "network distance" at which no shared rule connects them. That network distance of 2 corresponds to a physical separation of $2\ell_p$. This means each independent "pixel" of the black hole's surface is $2\ell_p \times 2\ell_p = 4\ell_p^2$ in area.

The two-Planck scale was derived independently in [21] from admissibility and distinguishability limits, prior to and without reference to black hole entropy. Its appearance here is a consequence of the constraint structure, not a parameter tuned to match the known answer.

4.2 Boundary Area Quantization

The area associated with a single effective boundary cell is therefore

$A^{cell} = \ell_e^2 = (2\ell_p)^2 = 4\ell_p^2$.

Each such cell supports exactly one independent binary commitment.

---

5. Derivation of the Bekenstein–Hawking Entropy

For the general reader. All the pieces are now in place. Each independent surface tile is $4\ell_p^2$ in area and carries one binary choice. To find the total entropy, we simply count the tiles: divide the total surface area by the tile size. This gives $N = A / 4\ell_p^2$ tiles, each contributing one bit of information. When entropy is measured in bits, $S = N = A / 4\ell_p^2$—exactly the Bekenstein–Hawking formula. The factor of 1/4 was never put in by hand; it is the ratio of the Planck area to the tile area, which was determined by the structure of space itself.

Let A be the total horizon area. The number of effective boundary cells is

$N = A/A^{cell} = A/4\ell_p^2$.

Each cell supports one independent binary commitment, so the total number of admissible configurations is

$\Omega = 2^N$.

The entropy is

$S = \ln \Omega = N \ln 2 = A/4\ell_p^2 \ln 2$.

5.1 On Entropy Units and the Role of ln 2

The derivation fixes the number of independent boundary commitments:

$N = A/4\ell_p^2$,

which is the origin of the 1/4 coefficient. The numerical value of entropy depends only on the conventional choice of logarithm base.

If entropy is defined in nats, $S \equiv \ln \Omega$, then with binary commitments $\Omega = 2^N$ gives $S = N \ln 2$. If entropy is defined in bits, $S_{\beta its} \equiv \log_2 \Omega$, then $S_{\beta its} = N$.

Thus the framework determines the geometric density of degrees of freedom on the horizon: one independent commitment per $4\ell_p^2$. Converting between nats and bits is a unit choice analogous to converting temperatures between Kelvin and Rankine. The 1/4 coefficient is entirely geometric and independent of this convention.

This situation is not unique to the present framework. Bekenstein's original information-theoretic argument [1], Strominger and Vafa's string microstate counting [5], and the loop quantum gravity derivation [6,7] all face the same unit-choice step. The physical content of any such derivation is the count of independent degrees of freedom per unit area; the rest is convention.

For the general reader. A brief technical aside: when physicists report black hole entropy, they have to choose an information unit—like choosing between miles and kilometers. In "bits," each yes/no choice counts as 1, so the formula reads $S = A / 4\ell_p^2$. In "nats" (the unit preferred in statistical mechanics), a yes/no choice counts as about 0.693, so the formula picks up that factor. This is purely a unit conversion—the underlying physics is identical. The important result is that each $4\ell_p^2$ patch carries exactly one binary degree of freedom. That geometric fact is what the framework derives, and it is independent of which unit you use to report the answer.

---

6. Comparison with Existing Approaches

The present derivation differs from established approaches to black hole entropy in both mechanism and assumptions. A brief comparison clarifies the structural distinctions.

For the general reader. Several other theoretical frameworks have also derived the Bekenstein–Hawking formula, each using very different ideas. This section compares them to the present approach. The key difference: many previous methods either require adjusting a free parameter to get the right answer, or only work for special types of black holes. The present approach has no adjustable parameters and works for all black holes.

String theory microstate counting (Strominger–Vafa) [5]. The entropy is computed by counting BPS microstates of a dual weakly-coupled system. The area scaling and the coefficient 1/4 emerge for specific extremal and near-extremal black holes. The mechanism is interior microstate counting in a dual description. The present approach requires no duality, no supersymmetry, and applies to generic horizons.

Loop quantum gravity (Rovelli–Smolin, Ashtekar–Baez–Corichi–Krasnov) [6,7]. The area spectrum of the horizon is quantized, and entropy is computed by counting spin-network punctures consistent with a given total area. The coefficient 1/4 is obtained only by fixing the Immirzi parameter γ to a specific value ($\gamma = \ln 2 / \pi\sqrt{3}$ in the original counting). The Immirzi parameter is a free parameter of the quantization, analogous to a θ-angle, and its value is not determined by internal consistency alone. In contrast, the present framework has no free parameters: K and ℓ_e are both fixed by closure consistency. We note that subsequent work by Engle, Noui, and Perez [29] has argued that the dependence on the Immirzi parameter can be removed when the appropriate quantum group structure is imposed at the horizon. While this development strengthens the internal consistency of loop quantum gravity, it also underscores that the origin of the 1/4 coefficient remains subtle even within that framework, motivating approaches—such as the present one—in which the coefficient follows directly from boundary constraint structure.

Entanglement entropy (Bombelli–Koul–Lee–Sorkin; Srednicki) [8,9,27]. Entropy is identified with the entanglement between degrees of freedom inside and outside the horizon. The area scaling is generic, but the coefficient depends on a UV cutoff that must be matched to the gravitational coupling. The present approach requires no UV regularization: the discreteness of the simplicial foam provides a natural cutoff, and the coefficient follows without matching.

Euclidean path integral (Gibbons–Hawking) [4]. The entropy is computed from the on-shell Euclidean action. The coefficient 1/4 follows from the Einstein–Hilbert action with the standard normalization $1/16\pi G$. This is closest in spirit to the present approach, in that the coefficient traces to a structural feature of the theory (the action normalization), but it operates within the continuum and does not explain why the action has that normalization.

In each case, the coefficient 1/4 either requires an adjustable parameter, depends on a specific class of black holes, or traces to an unexplained normalization. The present framework fixes the coefficient from the constraint structure of spacetime itself, with no parameters to adjust.

---

7. Discussion and Implications

The derivation shows that the Bekenstein–Hawking entropy coefficient is not an arbitrary numerical artifact but a structural consequence of three elements:

- Fixed constraint count ($K = 7$) per bulk simplex,
- Boundary projection of admissibility at null surfaces,
- The intrinsic two-Planck independence threshold.

Each of these is determined independently of black hole physics. Their combination yields $N = A / 4\ell_p^2$ independent boundary commitments—i.e., the 1/4 geometric coefficient—without tuning.

Because the same admissibility framework independently fixes the cosmological constant Λ [22], the appearance of the same fundamental scales ($K = 7$, $\ell_e = 2\ell_p$) in both results is structurally significant. The entropy coefficient and Λ arise from the same microscopic axioms, placing black hole thermodynamics and vacuum energy on a unified footing. A framework that produces two widely separated physical predictions—one at the Planck scale (black hole entropy) and one at the cosmological scale (Λ)—from the same parameter-free inputs is highly constrained and offers concrete targets for falsification.

For the general reader. The same set of rules that explains black hole entropy also, in separate work, predicts the value of the cosmological constant—the mysterious energy that drives the accelerating expansion of the universe. These are two of the most important unsolved problems in physics, operating at completely opposite scales (the tiniest versus the largest structures in the universe). The fact that a single framework with no adjustable parameters addresses both is either a remarkable coincidence or a sign that the framework is capturing something real about how space works.

7.1 Predictions and Testability

The framework makes several predictions that distinguish it from parameter-dependent approaches:

1. **Universality of the coefficient.** Within the framework, the 1/4 coefficient holds for all null horizons—Schwarzschild, Kerr, charged, cosmological—because the boundary constraint reduction depends only on the null character of the surface, not on the details of the black hole interior.
2. **No logarithmic corrections from the constraint structure.** The constraint reduction fixes the leading area term with no free parameters. Subleading corrections, if present, must arise from additional dynamics not captured by kinematic admissibility counting—such as horizon fluctuations, matter content, or effective field theory loops [28].
3. **Quantized horizon area.** The area is quantized in units of $4\ell_p^2$. The minimum horizon area increment is $\Delta A = 4\ell_p^2$, corresponding to the addition of one boundary cell. Note that in the present framework, this quantization arises from the graph-distance structure of independent commitments—it is an effective cell size set by constraint sharing, not an

eigenvalue spectrum derived from a Hamiltonian. Whether the same scale appears as a hard spectral gap in a full dynamical theory remains an open question.

For the general reader. The framework makes testable predictions. First, the 1/4 factor should be the same for every type of black hole—spinning, charged, or otherwise. Second, the surface area of a black hole should come in discrete steps, like a staircase rather than a ramp. The smallest possible increase in surface area is exactly four Planck areas. These predictions are beyond current experimental reach, but they are precise and falsifiable, which is what separates a scientific theory from speculation.

---

8. Conclusion

We have established that the Bekenstein–Hawking entropy coefficient follows from boundary constraint counting in a discrete admissibility framework, conditional on two inputs: the constraint count $K = 7$ per simplex and the independence threshold $\ell_e = 2\ell_p$. The factor 1/4 arises inevitably from the combination of a single independent boundary commitment per effective cell and a fundamental cell area of $4\ell_p^2$. Given these inputs, the derivation requires no additional tuning, no free parameters, and no holographic postulate.

The derivation proceeds by explicitly reducing the seven bulk constraints per simplex to one surviving boundary commitment through null degeneracy, constraint pairing, and gauge redundancy—a reduction that can be verified constraint by constraint. The two-Planck graph-distance threshold that sets the cell area arises from shared-face constraints between adjacent simplices, not from reverse-engineering the known answer.

The logical structure is: companion papers [20,21] establish $K = 7$ and $\ell_e = 2\ell_p$ from closure consistency and decoupling analysis; the present paper shows that these inputs, combined with boundary constraint reduction on null horizons, yield the 1/4 coefficient. Readers who accept the companion derivations obtain a parameter-free prediction; those who do not can treat this paper as establishing the precise conditional dependence.

A framework that simultaneously addresses the cosmological constant and the black hole entropy coefficient from related axioms is highly constrained, produces concrete predictions, and merits continued scrutiny.

For the general reader. The 1/4 in the black hole entropy formula is not a number someone chose—it is a number that space itself demands, given certain rules about how space is built. When you tile a black hole's surface with the smallest independent building blocks allowed by those rules, each block is exactly four Planck areas in size and carries exactly one bit of information. Count the blocks, and you get the entropy. The rules that determine the block size were established in separate work that had nothing to do with black holes. The fact that they reproduce this famous result—with no adjustment—suggests that the rules may be telling us something true about the fabric of space.

Acknowledgments

[To be added]

References

[1] J. D. Bekenstein, "Black holes and entropy," *Physical Review D* **7**, 2333–2346 (1973).

[2] S. W. Hawking, "Particle creation by black holes," *Communications in Mathematical Physics* **43**, 199–220 (1975).

[3] S. W. Hawking, "Black holes and thermodynamics," *Physical Review D* **13**, 191–197 (1976).

[4] G. W. Gibbons and S. W. Hawking, "Action integrals and partition functions in quantum gravity," *Physical Review D* **15**, 2752–2756 (1977).

[5] A. Strominger and C. Vafa, "Microscopic origin of the Bekenstein–Hawking entropy," *Physics Letters B* **379**, 99–104 (1996). arXiv:hep-th/9601029.

[6] C. Rovelli, "Black hole entropy from loop quantum gravity," *Physical Review Letters* **77**, 3288–3291 (1996). arXiv:gr-qc/9603063.

[7] A. Ashtekar, J. Baez, A. Corichi, and K. Krasnov, "Quantum geometry and black hole entropy," *Physical Review Letters* **80**, 904–907 (1998). arXiv:gr-qc/9710007.

[8] L. Bombelli, R. K. Koul, J. Lee, and R. D. Sorkin, "Quantum source of entropy for black holes," *Physical Review D* **34**, 373–383 (1986).

[9] M. Srednicki, "Entropy and area," *Physical Review Letters* **71**, 666–669 (1993). arXiv:hep-th/9303048.

[10] R. Bousso, "The holographic principle," *Reviews of Modern Physics* **74**, 825–874 (2002). arXiv:hep-th/0203101.

[11] T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," *Physical Review Letters* **75**, 1260–1263 (1995). arXiv:gr-qc/9504004.

[12] S. Carlip, "Black hole entropy from conformal field theory in any dimension," *Physical Review Letters* **82**, 2828–2831 (1999). arXiv:hep-th/9812013.

[13] G. 't Hooft, "Dimensional reduction in quantum gravity," in *Salamfestschrift*, edited by A. Ali, J. Ellis, and S. Randjbar-Daemi (World Scientific, 1993). arXiv:gr-qc/9310026.

[14] L. Susskind, "The world as a hologram," *Journal of Mathematical Physics* **36**, 6377–6396 (1995). arXiv:hep-th/9409089.

[15] J. M. Maldacena, "The large N limit of superconformal field theories and supergravity," *Advances in Theoretical and Mathematical Physics* **2**, 231–252 (1998). arXiv:hep-th/9711200.

[16] R. M. Wald, "Black hole entropy is the Noether charge," *Physical Review D* **48**, R3427–R3431 (1993). arXiv:gr-qc/9307038.

[17] T. Regge, "General relativity without coordinates," *Nuovo Cimento* **19**, 558–571 (1961).

[18] R. M. Williams and P. A. Tuckey, "Regge calculus: A brief review and bibliography," *Classical and Quantum Gravity* **9**, 1409–1422 (1992).

[19] J. Ambjørn, J. Jurkiewicz, and R. Loll, "Dynamically triangulating Lorentzian quantum gravity," *Nuclear Physics B* **610**, 347–382 (2001). arXiv:hep-th/0105267.

[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).

[22] K. Taylor, "Deriving the cosmological constant from admissibility constraints in discrete spacetime," VERSF Theoretical Physics Program (2025).

[23] K. Taylor, "Void Energy-Regulated Space Framework: Foundations and axioms," VERSF Theoretical Physics Program (2024).

[24] K. Taylor, "Bit Conservation and Balance: An information-theoretic foundation for discrete spacetime," VERSF Theoretical Physics Program (2024).

[25] S. Carlip, "Black hole thermodynamics," *International Journal of Modern Physics D* **23**, 1430023 (2014). arXiv:1410.1486.

[26] D. N. Page, "Hawking radiation and black hole thermodynamics," *New Journal of Physics* **7**, 203 (2005). arXiv:hep-th/0409024.

[27] S. N. Solodukhin, "Entanglement entropy of black holes," *Living Reviews in Relativity* **14**, 8 (2011). arXiv:1104.3712.

[28] A. Sen, "Logarithmic corrections to black hole entropy: An infrared window into the microstates," *General Relativity and Gravitation* **44**, 1207–1266 (2012). arXiv:1108.3842.

[29] J. Engle, K. Noui, and A. Perez, "Black hole entropy and SU(2) Chern–Simons theory," *Physical Review Letters* **105**, 031302 (2010). arXiv:1006.0634.

Appendix A: Closure Consistency and the Constraint Count K = 7

A.1 Summary

The derivation of black hole entropy in the main text relies on the fact that each simplicial cell carries exactly $K = 7$ independent admissibility constraints. This appendix summarizes the origin and uniqueness of this count.

Each simplex σ carries face-associated variables X_f and a binary orientation label $s \in \{+1, -1\}$. The seven admissibility constraints—local closure, three neighbor compatibility conditions, orientation consistency, distinguishability capacity, and bulk redundancy elimination—are enumerated explicitly in Section 2.2.

A.2 Uniqueness of K = 7

The closure-consistency argument establishing the uniqueness of $K = 7$ proceeds by exhaustion:

- If $K < 7$, the constraint set is incomplete: multiple inequivalent bulk completions exist for the same boundary data (non-uniqueness).
- If $K = 7$, the constraint set is complete and consistent: generic boundary data admits a unique admissible bulk completion.
- If $K > 7$, the constraint set is overdetermined: no admissible configurations survive for generic boundary data.

Requiring both local admissibility and global extensibility therefore fixes $K = 7$ uniquely. The full derivation is given in Ref. [20] and is not reproduced here.

A.3 Role in the Present Paper

The present paper does not re-derive $K = 7$; it uses this as a foundational input and tests whether black hole entropy follows as a consequence. The boundary constraint reduction (Appendix B) and explicit rank calculation (Appendix C) demonstrate that the $7 \to 1$ reduction on null horizons is enforced algebraically given this input.

Appendix B: Boundary Constraint Reduction on a Null Horizon

B.1 Objective

This appendix provides a formal account of how the seven bulk admissibility constraints per simplex reduce to a single independent boundary commitment on a null horizon.

B.2 Boundary Variables and Constraint Decomposition

Let σ be a simplex intersecting a null horizon $\mathscr{H}$. The closure constraint decomposes as

$$C_\sigma = C_\sigma^\perp + C_\sigma^\parallel,$$

where $\perp$ and $\parallel$ denote components normal and tangential to $\mathscr{H}$.

Structural assumptions on constraint space. The counting arguments in this appendix assume only that the admissibility constraints are equality constraints with well-defined rank, and that the face variables $X_f \in \mathcal{X}$ admit a linear (or linearizable) decomposition into normal and tangential components relative to a null hypersurface. In particular, the argument requires that matching constraints identify face data by equality and that the relevant gauge redundancy is discrete (here $\mathbb{Z}_2$). The specific realization of $\mathcal{X}$—whether vector-valued fluxes, group-valued holonomies linearized near admissible configurations, or other simplicial data—does not affect the constraint count so long as these structural conditions hold.

B.3 Freezing of the Normal Closure Constraint

On a null hypersurface, the normal vector lies in the tangent space of $\mathscr{H}$. Consequently, the normal component of the closure constraint carries no independent degree of freedom:

$$C_\sigma^\perp = 0$$

is automatically satisfied for admissible boundary data.

One constraint is therefore frozen by null geometry.

B.4 Pairing of Tangential Closure Constraints

For two adjacent boundary simplices σ and σ' sharing a face f, both impose a tangential closure condition involving X_f. The matching constraint

$$M_f(\sigma, \sigma') = 0$$

identifies these two conditions as the same physical requirement.

The remaining face of a boundary simplex is shared with the causally inaccessible interior; this is precisely the face whose associated constraint lies in the normal direction and is already eliminated by null degeneracy (Section B.3). Consequently, only faces shared with other boundary simplices contribute tangential closure constraints subject to pairing.

Thus, two tangential closure constraints are paired and eliminated as independent degrees of freedom.

B.5 Reduction of Neighbor Compatibility Constraints

Of the three neighbor compatibility constraints:

- One involves the normal direction and is subsumed by the frozen normal closure,
- The remaining two tangential constraints are again paired across shared faces via the same matching equations.

This removes three additional constraints.

B.6 Gauge Redundancy and Binary Quotient

The gauge redundancy removed at the boundary is distinct from the bulk redundancy eliminated in defining the independent constraint count $K = 7$. The latter ensures independence of bulk constraints; the former arises only after restriction to the boundary and corresponds to a residual $\mathbb{Z}_2$ relabeling symmetry of the surviving boundary commitment.

The orientation and distinguishability labels admit a minimal relabeling symmetry

$s \mapsto -s,$

corresponding to a global $\mathbb{Z}_2$ transformation that leaves all closure and matching equations invariant.

Quotienting by this symmetry removes one redundant binary label, leaving a single physical equivalence class.

B.7 Surviving Boundary Commitment

After all reductions:

- 1 constraint is frozen,
- 4 constraints are eliminated by pairing,
- 1 constraint is removed by gauge redundancy,

leaving exactly

$7 - (1 + 4 + 1) = 1$

independent boundary constraint per effective cell.

This surviving constraint corresponds to one irreversible binary commitment, i.e., one unit of classical distinguishability contributing to entropy.

B.8 Independence Scale and Effective Cell Area

Boundary commitments associated with adjacent simplices share matching constraints and are therefore not independent. Independence occurs at the minimal adjacency distance where no shared constraint path exists.

In the simplicial adjacency graph, this distance corresponds to two adjacency steps, fixing the independence length

$\ell_e = 2\ell_p$.

On a two-dimensional boundary, this implies an effective independent area

$A^{cell} = \ell_e^2 = 4\ell_p^2$.

This area characterizes boundary tiles, not folds.

B.9 Boundary Reduction Result

Boundary Reduction Result. Admissibility and null boundary geometry reduce the bulk constraint system to exactly one independent irreversible binary commitment per boundary area $4\ell_p^2$.

This result supplies the structural basis for the Bekenstein–Hawking area coefficient derived in the main text.

Appendix C: Explicit Constraint Matrix for a Four-Simplex Boundary Patch

C.1 Purpose of This Appendix

Sections 3.2–3.3 argue that seven bulk admissibility constraints per simplex reduce to a single independent boundary commitment on a null horizon. Section 3.3 presents a toy boundary patch illustrating this reduction. The purpose of this appendix is to make that example fully explicit, by:

- defining concrete boundary variables,
- writing all constraint equations explicitly,
- exhibiting the associated constraint matrix, and
- computing its rank directly.

This eliminates any ambiguity about whether the reduction follows from the structure of the constraints or from a particular choice of variables.

C.2 Geometry and Boundary Variables

Consider a minimal boundary patch consisting of four triangular simplices σ_1, σ_2, σ_3, σ_4 arranged as in Section 3.3:

- σ_1 shares faces with σ_2 and σ_3,
- σ_4 shares faces with σ_2 and σ_3,
- σ_2 and σ_3 do not share a face,
- each simplex has one face oriented toward the bulk interior.

After null reduction (Section 3.2), each simplex contributes two independent tangential face variables. Denote these by

$(x_{1a}, x_{1\beta}, x_{2a}, x_{2\beta}, x_{3a}, x_{3\beta}, x_{4a}, x_{4\beta}) \in \mathbb{R}^8$.

No assumptions are made about the physical interpretation of these variables beyond linearity and equality matching.

C.3 Matching Constraints

Each shared boundary face imposes a matching constraint identifying the corresponding tangential variables. For the chosen patch, these are:

$x_{1a} - x_{2a} = 0$ $x_{1\beta} - x_{3a} = 0$ $x_{4a} - x_{2\beta} = 0$ $x_{4\beta} - x_{3\beta} = 0$

These four equations encode all tangential compatibility conditions between boundary simplices. The interior-facing faces are normal-direction faces and have already been eliminated by null degeneracy via the closure constraints.

C.4 Constraint Matrix

The matching constraints may be written compactly as a linear system

$M\mathbf{x} = 0$,

where $\mathbf{x} \in \mathbb{R}^8$ is the vector of tangential variables ordered as

$\mathbf{x} = (x_{1a}, x_{1\beta}, x_{2a}, x_{2\beta}, x_{3a}, x_{3\beta}, x_{4a}, x_{4\beta})^T$.

The constraint matrix M is then

$$
\begin{array}{lcccccccc}
 & x_{1a} & x_{1\beta} & x_{2a} & x_{2\beta} & x_{3a} & x_{3\beta} & x_{4a} & x_{4\beta} \\
M = [& 1 & 0 & -1 & 0 & 0 & 0 & 0 & 0\ ] \\
\quad [& 0 & 1 & 0 & 0 & -1 & 0 & 0 & 0\ ] \\
\quad [& 0 & 0 & 0 & -1 & 0 & 0 & 1 & 0\ ] \\
\quad [& 0 & 0 & 0 & 0 & 0 & -1 & 0 & 1\ ]
\end{array}
$$

C.5 Rank and Null Space

The sparsity pattern of M is fixed by the adjacency graph of the patch: each row corresponds to a shared-face matching constraint, and each nonzero entry occurs only where a given simplex carries the corresponding face variable. The matrix is not chosen; it is implied by the patch topology.

The four rows of M are linearly independent. Therefore,

$\text{rank}(M) = 4.$

Since the total number of variables is 8, the dimension of the null space is

$\dim(\ker M) = 8 - 4 = 4.$

Thus the solution space of admissible boundary configurations is four-dimensional.

C.6 Interpretation

Each of the four independent directions in the null space corresponds to one independent boundary degree of freedom, one per boundary simplex. The four-dimensional null space parameterizes continuous boundary configurations; in the full framework, each independent degree of freedom corresponds to a stable fold-capacity that is realized as a binary commitment upon irreversibility (Section 2.4). The four independent boundary degrees therefore yield $2^4 = 16$ admissible committed configurations for the patch. When irreversibility is imposed (e.g., by coarse-graining or measurement), each such degree of freedom yields a single binary commitment.

This confirms explicitly that, for the four-simplex patch,

4×7 bulk constraints $\rightarrow$ 4×1 independent boundary commitments.

The result follows from the rank of the constraint matrix and does not rely on narrative pairing arguments.

C.7 Why the Result Is Not "Built In"

Although the counting result appears simple, it is not imposed by construction. The dimension of the null space depends on:

- the existence of matching constraints as equality conditions,
- the fact that interior-facing faces are eliminated by null degeneracy,
- the topology of the boundary patch (which determines the number of matching constraints).

A different adjacency structure, or the absence of null freezing, would change the rank of *M*. The present calculation shows that given the admissibility constraints and null boundary structure defined in the main text, the reduction to one independent boundary commitment per simplex is enforced algebraically.

For *d*-dimensional face data ($X_f \in \mathbb{R}^d$), the constraint structure scales uniformly: the null space becomes 4*d*-dimensional, but the closure-consistency count K = 7 constrains the total per-simplex admissible freedom to a single binary commitment regardless of *d*, as the additional dimensions are absorbed by correspondingly higher-dimensional closure and matching constraints.

C.8 Conclusion

This appendix provides a concrete linear-algebraic verification of the boundary constraint reduction for a minimal four-simplex patch. The explicit constraint matrix has rank 4, leaving a four-dimensional null space corresponding to one independent boundary commitment per simplex.

This removes any ambiguity about whether the 7→1 reduction is merely restated in words: for this patch, it follows directly from the structure and rank of the constraint system.