Geometric Closure, Simplicial Foam, and Renormalization Consistency

Keith Taylor VERSF Theoretical Physics Program

Summary. We ask whether the universe could fail to "close" geometrically if its fundamental constants took the wrong values. Beginning from a triangle-based coherence model of pre-geometric transport, we show that minimal consistency forces a U(1) phase structure and that large-scale closure must occur on the de Sitter horizon sphere. Topological integrality then requires that the number of Planck-scale cells on that surface multiplied by the electromagnetic phase per cell equals an integer. This yields a direct constraint relating α and Λ. When this constraint is combined with independent discrete mechanisms governing the coupling strength and vacuum regulation, the resulting winding number depends super-exponentially on the discrete constraint count K. Only K = 7 produces a winding number in the observed cosmological regime; nearby integers miss by 50–100 orders of magnitude. In this framework, the constraint count, the strength of electromagnetism, and the size of the universe are not independent—they are mutually restricted by geometric admissibility.

For the general reader. This paper explores a simple but radical idea: what if the basic constants of nature are not independent numbers, but must fit together for the universe to "close" geometrically? Imagine building space out of tiny triangular pieces. For those triangles to join consistently, they must satisfy a small set of internal rules. When those rules are written down carefully, they imply that space must carry a built-in phase symmetry—the same mathematical structure that underlies electromagnetism. Now zoom out. If the universe is expanding under the influence of a cosmological constant, it naturally has a spherical boundary—the de Sitter horizon. That sphere can be thought of as tiled with Planck-sized patches. A deep result from geometry says that the total "twist" of a phase field over a closed surface must add up to a whole number. In this framework, each Planck patch contributes a tiny amount of twist governed by the fine-structure constant. For the universe to close properly, the number of patches times that twist must equal an exact integer. When this global closure rule is combined with independent mechanisms that determine how the coupling strength depends on a discrete parameter K, the result is striking: the winding number becomes extraordinarily sensitive to K. Only one value—K = 7—places the total twist in the range implied by cosmology. Nearby values miss by tens of orders of magnitude. In other words, the strength of electromagnetism, the size of the universe, and the discrete structure underlying space are not arbitrary in this picture—they are tied together by the requirement that the geometry of the universe fits together without contradiction.

Contents

Abstract

We develop a unified geometric argument establishing that the fine-structure constant α and the cosmological constant Λ are jointly constrained by a single topological closure condition, and that the discrete constraint count K = 7 is the unique value compatible with the observed cosmological hierarchy at order-of-magnitude level.

The argument proceeds in five stages. First, we show that triangle coherence in a simplicial pre-geometric substrate, decomposed into K independent constraints, forces a minimal U(1) gauge redundancy through explicit exclusion of trivial, discrete, and non-abelian alternatives (§2). Second, we derive the global cell count N_Σ by identifying the de Sitter event-horizon 2-sphere as the unique closed surface satisfying homogeneity, isotropy, and asymptotic time-independence, discretized in Planck-area units (§3). Third, Chern–Weil integrality on the resulting U(1) bundle yields the closure condition $N_\Sigma \cdot \theta(\alpha) = 2\pi n$ with $n \in \mathbb{Z}$ (§4). Fourth, we show that renormalization-group consistency constrains the holonomy map θ(α) to be evaluated at a fixed reference scale, with the apparent higher-order coefficient c_2 revealed as a logarithmic, scale-dependent quantity rather than a universal constant (§5). Fifth, we show that when the closure equation is combined with the discrete coupling formula α(K) and the Two-Planck vacuum regulation Λ(K), the topological winding number n(K) is a super-exponential function of K, and only K = 7 places n in the cosmologically observed regime $\sim 10^{121}$, with K = 6 and K = 8 failing by 50–100 orders of magnitude (§6).

The combined architecture converts K = 7 from an enumerated input into a selected output, with α and Λ jointly constrained by an admissibility condition up to discrete topological class n and controlled corrections.

---

1. Introduction

The fine-structure constant $\alpha \approx 1/137$ and the cosmological constant $\Lambda \approx 1.1 \times 10^{-52}\ \text{m}^{-2}$ are among the most precisely measured yet least understood parameters in fundamental physics. In the Standard Model and general relativity, they enter as independent inputs. The present work investigates whether a pre-geometric consistency requirement—topological closure of a U(1) bundle over a cosmological horizon surface—can render them mutually constrained by an admissibility condition, and whether the discrete structure underlying this closure selects its own constraint count.

The argument is constructive. We do not assume U(1) gauge symmetry; we derive it as the minimal structure compatible with holonomy-based coherence in a simplicial substrate. We do not choose the cosmological horizon by hand; we show it is the unique 2-surface satisfying the symmetry and stability requirements of the closure argument. And we do not assume K = 7; we show it is the only integer for which the combined system of closure, coupling, and vacuum regulation produces a topological winding number in the observed cosmological range.

We state explicitly the logical status of each element. The K-constraint decomposition of triangle coherence, the identification of coherent holonomy with electromagnetic phase, and the leading-order uniformity of per-cell holonomy are structural inputs (assumptions). The U(1) minimality, the Chern–Weil integrality, the RG stability constraint, and the K = 7 selection are derived consequences. The de Sitter surface selection is a derived consequence of stated symmetry requirements, conditional on the late-time cosmology approaching a de Sitter attractor. The paper stands or falls on whether the structural inputs are physically reasonable and whether the derived consequences are internally consistent.

Main claim. The closure condition $N_\Sigma \cdot \alpha_0 = n \in \mathbb{Z}$ does not uniquely determine α and Λ; it restricts them to a discrete admissibility family labeled by the integer n. Closure alone does not select K. What selects K is the combination of the closure equation with two independent discrete mechanisms: the coupling formula α(K) from simplicial coherence counting, and the vacuum regulation Λ(K) from Two-Planck dimensional transmutation. The resulting winding number n(K) is super-exponentially sensitive to K, and only K = 7 places n in the observed cosmological regime. The claim is therefore not "we derive α and Λ from first principles" but "the discrete coherence structure, the electromagnetic coupling, and the cosmological constant are jointly constrained by geometric admissibility, and only one value of the constraint count is consistent with the observed hierarchy."

1.1 Assumption ledger

For rapid reference, we collect here the standing assumptions (SA), bridge hypotheses (BH), and structural inputs (SI) that the paper asks the reader to grant. Each is developed in the section indicated.

Standing assumptions (formal axioms of the coherence framework):

- **SA1** (Gauge invariance). Coherence predicates depend on gauge-equivalence classes of transport data, not on representative choices. (§2.1)
- **SA2** (Robustness). Sufficiently small continuous perturbations of edge data do not change coherence status, away from a measure-zero set of critical configurations. (§2.3)
- **SA3** (Minimality). The per-cell coherence predicate introduces no independent dimensionless parameters beyond the single holonomy-strength coupling. The holonomy normalization is controlled by one real degree of freedom. (§2.3)
- **SA3′** (Scalar completeness under local gluing). For any triangle Δ shared by two tetrahedra τ_1 and τ_2, the loop-closure constraint C4 and the co-face compatibility constraint C7 must be decidable using only: (i) the transport data on τ_1, (ii) the transport data on τ_2, and (iii) a single real scalar $s(H_\Delta)$ computed from the triangle holonomy. No additional local comparison data (choice of internal frame, stabilizer alignment, maximal-torus section, normal-bundle identification) may be introduced beyond the existing embedding constraints C5–C7. SA3′ is not an additional assumption beyond K = 7 minimality; it is the operational content of the statement that C4 contributes exactly one independent continuous constraint channel. (§2.3)

Operational justification of SA3′. SA3′ encodes an information restriction: the coherence test across a shared face must be decidable from gauge-invariant loop data accessible locally on the two incident tetrahedra. Allowing additional alignment/section data (e.g., internal frame matching or stabilizer identification) is equivalent to introducing extra locally propagating degrees of freedom beyond the transport variables already present. In the minimal programme, such extra structure is not "free bookkeeping"; it would constitute new physics and would introduce additional independent parameters or fields. The U(1) minimality result is therefore conditional on the absence of such alignment fields: if nature supplies additional local comparison structure, larger gauge groups could be admissible, but that would be a different (non-minimal) framework.

- **SA4** (Finite distinguishability). The space of holonomy classes is compact, excluding non-compact gauge groups. Physically: a finite-resolution observer cannot distinguish infinitely many holonomy states, so the holonomy target space must be bounded. (§2.3)
- **SA5** (One-parameter holonomy spectrum). The physically distinguishable near-identity holonomy states relevant to C4 form a connected 1-parameter family (a single "holonomy angle" degree of freedom), consistent with C4 contributing exactly one independent continuous constraint channel. (§2.2)

Bridge hypotheses (identifications connecting the abstract framework to physics):

- **BH1** (Holonomy–coupling identification). The per-cell holonomy normalization on the closure surface is identified with the electromagnetic fine-structure constant at a physical reference scale. Justification: (i) under SA3, the framework has exactly one dimensionless U(1) coupling; (ii) electromagnetism is the only unbroken long-range U(1) gauge interaction in late-time cosmology (hypercharge is broken; hidden U(1)s would require new fields, violating SA3); (iii) the identification is therefore the unique SA3-compatible bridge between the abstract holonomy normalization and infrared physics. (§4.3)
- **BH2** (Planck-area tiling). The Planck area $\ell_P^2 = \hbar G/c^3$ is the correct cell area for discretizing the U(1) bundle integral on the closure surface. We use ℓ_P^2 as the minimal universal covariant area scale built from $\hbar$, G, c; alternative order-unity area gaps (e.g., $4\ell_P^2$, or the loop-quantum-gravity area gap $\gamma\sqrt{3}\ \ell_P^2$) shift n by a multiplicative O(1) factor and do not affect the K-selection, which depends on super-exponential sensitivity. (§3.3)
- **BH3** (de Sitter attractor). The late-time cosmology approaches a de Sitter fixed point ($w \to -1$ asymptotically), so that the de Sitter event-horizon sphere is the canonical closure surface. (§3.2)

Structural inputs (additional hypotheses used in the K-selection argument of §6):

- **SI1** (Discrete coupling formula). $\alpha(K) = 2K / [2^K(2K + 1)]$ with $N_{loop} = 2K$. (§6.2)
- **SI2** (Two-Planck vacuum regulation). $\Lambda = 8\pi\ell_P^2 / \xi^4$. (§6.3)

- **SI3** (Dimensional transmutation). $\xi(K) = \ell\_em \cdot \exp[(4/K)(2^K - 1/p\_c)]$ with $\ell\_em = 2\ell\_P$. (§6.3)

Derived consequences (proven from the above):

- U(1) gauge minimality (Theorem 1)
- Unique closure surface selection (Proposition, §3.2)
- Joint admissibility condition on α and Λ (Theorem C)
- RG stability via reference-scale anchoring (§5)
- K = 7 selection at hierarchy level (Theorem 2)

---

2. Triangle Coherence and the K = 7 Constraint Structure

2.1 Definitions

Throughout, K denotes the discrete constraint count. The simplicial 2-complex on which transport is defined is denoted $\mathcal{K}$.

Definition 1 (Transport system). Let $\mathcal{K}$ be a simplicial 2-complex with vertex set V and oriented edge set E. A *transport system* assigns to each oriented edge (i → j) an element U_ij in a topological group G, with $U\_ji = U\_ij^{-1}$.

Definition 2 (Triangle holonomy). For an oriented triangle Δ = (i, j, k) with boundary edges (i → j), (j → k), (k → i), the *holonomy* is H_Δ := U_ij U_jk U_ki ∈ G.

Definition 3 (Gauge reparameterization). A local change of representatives is a map g : V → G acting as $U\_ij \mapsto g\_i\, U\_ij\, g\_j^{-1}$. Two transport systems related by such a map are *gauge-equivalent*.

Definition 4 (Edge admissibility). An oriented edge e = (i → j) is *admissible* if it satisfies a predicate $\mathcal{A}(e) \in \{0, 1\}$ that tests whether the edge supports well-defined transport. An edge may fail admissibility due to degenerate simplex geometry (zero-length edge), forbidden local configurations, or loss of distinguishability between adjacent vertices. Admissibility is a precondition for transport: if $\mathcal{A}(e) = 0$, the transport element U_ij is undefined or degenerate.

Definition 5 (Coherence predicate). A triangle Δ is *coherent* if its edge data and holonomy satisfy a set of conditions invariant under gauge reparameterization. The coherence predicate depends on the gauge-equivalence class of {U_ij}, not on a choice of representatives.

2.2 The seven independent constraints

We now enumerate the independent conditions that a triangle $\Delta = (i, j, k)$ must satisfy to be coherent within a 4-dimensional simplicial foam. The constraints fall into three groups: edge admissibility, loop closure, and embedding consistency.

Edge admissibility (C1–C3). Each of the three oriented edges of Δ must be admissible:

- C1: $\mathcal{A}(e_{ij}) = 1$ (edge ij supports well-defined, invertible transport).
- C2: $\mathcal{A}(e_{jk}) = 1$ (edge jk supports well-defined, invertible transport).
- C3: $\mathcal{A}(e_{ki}) = 1$ (edge ki supports well-defined, invertible transport).

These are independent because each edge may fail admissibility independently—for example, one edge of a triangle may become degenerate (zero length, coincident vertices, or singular local geometry) while the other two remain well-defined. Together they contribute 3 binary constraints. Note that admissibility is not automatic: in a pre-geometric substrate where edge data encode relational structure, degeneracy is a genuine failure mode, not a mathematical triviality.

Loop closure (C4). The holonomy $H_\Delta = U_{ij} U_{jk} U_{ki}$ must lie in a distinguished coherence class $\mathcal{C} \subset G$:

- C4: $H_\Delta \in \mathcal{C}$.

This is a gauge-invariant condition (for abelian G, H_Δ is fully gauge-invariant; for non-abelian G, the conjugacy class of H_Δ is gauge-invariant). It is independent of C1–C3 because edges may be individually admissible while their composed product violates the closure bound. C4 contributes 1 constraint parameterized by a single real degree of freedom (the holonomy angle, for compact G near the identity).

Embedding consistency (C5–C7). In a 4-dimensional simplicial foam, each triangle Δ is a 2-face shared by adjacent tetrahedra (3-simplices), which are in turn faces of 4-simplices. Coherence of Δ requires compatibility with this higher-dimensional embedding:

- C5 (Orientation compatibility): The orientation of Δ, as inherited from each adjacent tetrahedron, must be consistent. In a 4-dimensional simplicial complex, a triangle generically borders at least two tetrahedra; consistent orientability of the triangulation requires that the induced orientations agree up to sign across each shared face. This is one binary constraint.
- C6 (Normal transport consistency): The transport data on the edges of Δ must be compatible with transport data along edges connecting Δ's vertices to vertices of adjacent tetrahedra not contained in Δ. This ensures that curvature localized on Δ (the deficit angle in the Regge calculus interpretation) is consistent with the local geometry. This is one constraint per triangle, testing whether the holonomy of Δ is compatible with the deficit angle implied by the dihedral angles of adjacent 4-simplices.

- C7 (Co-face compatibility): The holonomy of Δ must be consistent with the transport environments of the two tetrahedra sharing Δ as a common face. Specifically, the two tetrahedral holonomy environments—the products of face holonomies around each tetrahedron—must agree on the value of H_Δ up to gauge equivalence. C7 must be decidable using only data on τ_1, τ_2, and their shared face Δ; it may not require gauge fixing over a larger neighborhood. This is a *local* gluing constraint between the two tetrahedra adjacent to Δ, not the global Bianchi identity over all faces of a 4-simplex. (The global Bianchi relations, which link holonomies of all 10 faces of a 4-simplex, are derived in §3.1 as independent closure channels at the 4-simplex level; they are emergent consequences of the per-triangle constraints, not identical to C7. See Appendix B.5 for further clarification.)

The three embedding constraints C5–C7 are independent of each other (orientation, normal transport, and co-face compatibility test different geometric data) and independent of C1–C4 (a triangle may have admissible edges with coherent loop closure yet fail to embed consistently in the higher-dimensional foam).

Total: K = 7 independent coherence constraints per triangle.

2.3 Derivation of the minimal U(1) gauge redundancy

Any transport system on the 1-skeleton admits local gauge reparameterizations $U_{ij} \mapsto g_i U_{ij} g_j^{-1}$ that do not alter the relational content. This defines a gauge redundancy valued in G. We now show that the minimal compact gauge group consistent with the K = 7 coherence structure is U(1), by systematically excluding all alternatives.

Lemma 1 — Trivial holonomy is excluded

Statement. If G is trivial ($|G| = 1$), then $H_\Delta = e$ for all triangles and all edge data. The constraint C4 is automatically satisfied and cannot constitute an independent coherence condition.

Proof. If $|G| = 1$, every transport element equals the identity: $U_{ij} = e$ for all edges. Therefore $H_\Delta = e \cdot e \cdot e = e$ for every triangle, regardless of edge configuration. Any predicate of the form "$H_\Delta \in \mathcal{C}$" is satisfied trivially and carries no information. This means C4 cannot serve as an independent failure mode, contradicting the K = 7 decomposition in which C4 is a distinct, non-redundant constraint. ■

Lemma 2 — Discrete gauge groups are excluded

Statement. If G is a discrete group (e.g., $G = \mathbb{Z}_n$) and the transport data admit continuous perturbations in a connected neighborhood, then any gauge-invariant holonomy map is locally constant. Consequently, holonomy either carries no nontrivial information or loses robustness under refinement.

Proof. Fix a triangle Δ and consider the holonomy map $\Phi : \mathcal{U} \to G$ assigning H_Δ to the edge data in a connected neighborhood $\mathcal{U}$ of the transport variables $\{U\_ij\}$. The map Φ is continuous because it is a finite product of group multiplication and inversion operations, which are continuous in any topological group.

If G is discrete, its only connected subsets are singletons. The continuous image of a connected set under a continuous map is connected. Therefore $\Phi(\mathcal{U})$ is a singleton: holonomy is constant throughout $\mathcal{U}$.

This means that holonomy classes cannot encode a continuous response to microscopic variation of edge data. To obtain nontrivial variability one must either (i) permit discontinuous jumps in holonomy class, violating robustness (SA2), or (ii) introduce additional structure beyond the discrete holonomy labels, violating minimality (SA3).

Therefore, no discrete group can serve as the minimal gauge redundancy for a coherence predicate that is both robust and non-trivially sensitive to holonomy. ■

Lemma 3 — Non-abelian gauge groups violate scalar-complete local gluing

Statement. Assume SA3′ (scalar completeness under local gluing). Then the holonomy gauge group governing C4 must be abelian.

Proof. Let Δ be a triangle shared by two tetrahedra τ_1 and τ_2. Each tetrahedron determines a representative holonomy element $H\_\Delta^{(1)}$ and $H\_\Delta^{(2)}$ computed from its local transport representatives. Gauge invariance (SA1) implies that any quantity entering the coherence predicate must be invariant under local reparameterizations; for a general compact Lie group, the natural invariant content of a single element is its conjugacy class.

Under SA3′, the shared-face consistency check (C7) must be decidable using only a single real scalar $s(H\_\Delta)$ computed from H_Δ, with no additional alignment structure permitted.

If G is non-abelian, then conjugacy classes do not canonically specify a unique representative element: any representative choice requires additional structure (e.g., a section to a maximal torus, an internal frame choice, or stabilizer alignment). Concretely, even in the smallest non-abelian case SU(2), a near-identity element is parameterized by an angle θ and an internal axis $\hat{n}$, and the conjugacy class discards the axis data. Two tetrahedral environments can therefore yield holonomies with the same class parameter (same θ) but different internal axes. Determining whether these represent the same glued transport system across τ_1 and τ_2 requires specifying how the internal axes (or stabilizers) are identified across the shared face—precisely the extra local gluing data SA3′ forbids. Equality of conjugacy class does not imply equality of group element under the gauge transformations admissible on the shared face; in non-abelian groups, matching class parameters does not define a canonical glued transport configuration.

Equivalently: any map $s : G \to \mathbb{R}$ that is invariant under conjugation collapses the degrees of freedom transverse to conjugacy classes. For non-abelian G, these collapsed degrees of freedom are exactly what controls whether two locally computed holonomies are compatible as the same glued object, rather than merely class-equivalent. Enforcing full co-face compatibility therefore requires additional independent comparison structure, contradicting SA3′.

For an abelian group, every element is its own conjugacy class. The scalar $s(H_\Delta)$ (the phase angle) completely determines H_Δ, and no alignment structure is needed to compare holonomies across the shared face. Co-face compatibility is decidable from s alone.

Hence G cannot be non-abelian. Therefore G must be abelian. ■

Corollary 3 — Minimal compact connected abelian choice is U(1)

Statement. If, in addition, the holonomy spectrum relevant to C4 is one-parameter (SA5), then the minimal compact connected abelian group consistent with the holonomy sector is U(1).

Proof. Any compact connected abelian Lie group is a torus $T^m \cong U(1)^m$. SA5 restricts the holonomy sector to a single continuous parameter, hence $m = 1$, so $G \cong U(1)$. If the near-identity holonomy sector were m-parameter ($m > 1$), then either (i) C4 would require m independent continuous constraint channels, contradicting $K = 7$, or (ii) the discarded parameters would reappear in gluing compatibility, violating SA3′. ■

Theorem 1 — Minimality of U(1)

Statement. Under SA1 (gauge invariance), SA2 (robustness), SA3 (minimality), SA3′ (scalar-complete local gluing), SA4 (finite distinguishability), and SA5 (one-parameter holonomy spectrum), U(1) is the minimal compact gauge redundancy consistent with the $K = 7$ coherence structure, and is the unique minimal compact connected choice.

Proof. By Lemma 1, trivial G is excluded (C4 would not be independent). By Lemma 2, discrete G is excluded (holonomy would be topologically rigid under SA2). By Lemma 3, non-abelian compact Lie groups are excluded (scalar-complete local gluing under SA3′ forces abelianity). By Corollary 3, SA5 forces the abelian group to be one-dimensional: $G \cong U(1)$. Non-compact groups are excluded by SA4: finite distinguishability requires the space of holonomy classes to be compact (a finite-resolution observer cannot distinguish among unboundedly many holonomy states, so the holonomy target must be bounded). ■

Remark (Structural inevitability). The emergence of U(1) here is not a statement about gauge symmetry in general, but about the minimal algebraic structure capable of supporting (i) a single continuous holonomy constraint and (ii) local scalar-complete gluing without additional comparison data. Within that class, U(1) is the unique compact connected realization.

---

3. From Local Simplicial Combinatorics to the Global Cell Count

3.1 Local structure: the 2-skeleton of the 4-simplex

A 4-simplex σ^4 with vertex set {0, 1, 2, 3, 4} has the following face counts:

Dimension	Name	Count
0	vertices	C(5,1) = 5
1	edges	C(5,2) = 10
2	triangles	C(5,3) = 10
3	tetrahedra	C(5,4) = 5
4	4-simplex	C(5,5) = 1

Curvature in 4-dimensional Regge calculus resides on 2-faces (triangles): each triangle carries a deficit angle encoding the local curvature. Therefore the 10 triangular faces of σ^4 are the fundamental curvature carriers.

The triangular holonomies are not all independent. The second Betti number of the 2-skeleton $\mathcal{K}^2$ of σ^4 counts the independent 2-cycles—closed 2-chains within the triangulation that impose relations among the face holonomies. Using the Euler characteristic:

$\chi(\mathcal{K}^2) = |V| - |E| + |F| = 5 - 10 + 10 = 5$

$\mathcal{K}^2$ is connected, so $\beta_0 = 1$. Every 1-cycle in the 1-skeleton (the complete graph K_5) bounds a 2-chain in $\mathcal{K}^2$, because any cycle in K_5 decomposes into 3-cycles (triangles), and every triangle of K_5 is a 2-face of $\mathcal{K}^2$. (Equivalently, $\mathcal{K}^2$ is the 2-skeleton of a contractible simplex and is therefore simply connected; see e.g. Munkres, *Elements of Algebraic Topology*, §5.) Therefore $H_1(\mathcal{K}^2; \mathbb{Z}) = 0$, giving $\beta_1 = 0$. From the Euler relation:

$\beta_0 - \beta_1 + \beta_2 = \chi(\mathcal{K}^2)$ $1 - 0 + \beta_2 = 5$ $\beta_2 = 4$

These 4 independent 2-cycles provide 4 *global* Bianchi-type closure relations among the 10 face holonomies of the 4-simplex. They emerge from the topology of the full 2-skeleton and constrain how face holonomies combine across the entire 4-simplex. This is distinct from the per-triangle co-face compatibility C7, which is a local constraint between the two tetrahedra sharing a given triangle (see §2.2 and Appendix B.5).

The total geometric information content of the local 4-simplex is therefore characterized by:

N_loop = 10 (curvature-carrying faces) + 4 (independent global closure relations) = 14

This establishes that geometric information in the UV (simplicial) regime propagates on 2-cells and is subject to closure constraints arising from the topology of the 2-skeleton.

3.2 Selection of the IR closure surface

The local face-based structure must have a global counterpart if geometric closure is to extend from the UV to the IR. We require a closed 2-surface Σ that serves as the domain for a global closure condition. The following requirements constrain the choice:

1. **Closure**: Σ must be a compact, oriented 2-manifold without boundary, so that Chern–Weil integrality applies.
2. **Homogeneity and isotropy**: Σ must be invariant under the spatial symmetry group SO(3) of the background cosmology, ensuring the closure condition respects the cosmological principle.
3. **Asymptotic time-independence**: Σ must define a fixed geometric scale in the late-time limit, so that the closure condition is not tied to a particular cosmological epoch.
4. **Determination by Λ alone**: Σ must be parameterized solely by the cosmological constant, without dependence on initial conditions, foliation choice, or observer location.

Proposition. In a late-time accelerating FRW cosmology approaching a de Sitter fixed point (BH3), the unique (up to isometry) closed 2-surface satisfying conditions 1–4 is the de Sitter event-horizon cross-section $\Sigma \cong S^2$ with areal radius $r_\Lambda = \sqrt{3/\Lambda}$.

Justification. The particle horizon and last-scattering surface are time-dependent: their comoving radii grow with conformal time and therefore violate condition 3. The apparent horizon is foliation-dependent, violating condition 4. The de Sitter event horizon, by contrast, exists as an invariant causal boundary in the asymptotic de Sitter phase, with constant areal radius $r_\Lambda = c/H_\Lambda = \sqrt{3/\Lambda}$ (in natural units). It is a 2-sphere, hence SO(3)-homogeneous, satisfying condition 2. It is compact and orientable, satisfying condition 1. And it is determined entirely by Λ, satisfying condition 4. No other standard cosmological 2-surface meets all four requirements simultaneously.

Conditionality. This selection is conditional on BH3 (late-time de Sitter approach). The closure surface is taken in the asymptotic de Sitter regime; the present epoch may be viewed as approaching this limit, with corrections suppressed by deviations of w from −1 and by finite-time approach to the attractor. The closure condition is therefore an IR asymptotic statement, not a claim that today's horizon is already exact. If the dark energy equation of state deviates from $w = -1$ at late times—for example, in quintessence or phantom models—the canonical closure surface must be revised, and the closure condition would take a different form. This constitutes an additional falsifier: observation of $w \neq -1$ asymptotically would invalidate the specific surface selection used here.

3.3 The global cell count

Discretizing the closure surface Σ in Planck-area units $\ell_P^2 = \hbar G/c^3$ (BH2). We use ℓ_P^2 as the minimal universal covariant area scale built from ħ, G, and c; alternative order-unity area gaps (such as those appearing in loop quantum gravity, where the minimal area eigenvalue is $\gamma\sqrt{3}$

$\ell_P^2/2$ with Barbero-Immirzi parameter $\gamma \sim 0.24$) would shift N_Σ and hence n by an O(1) factor, which does not affect the K-selection (the inter-K gaps are 10^{50+}).

$A_\Sigma = 4\pi r_\Lambda^2 = 4\pi \cdot (3/\Lambda) = 12\pi / \Lambda$

$N_\Sigma = A_\Sigma / \ell_P^2 = 12\pi / (\Lambda \ell_P^2)$

With $\Lambda \approx 1.1 \times 10^{-52}$ m^{-2} and $\ell_P \approx 1.616 \times 10^{-35}$ m, this gives $N_\Sigma \approx 1.3 \times 10^{123}$.

The relationship between the local count ($N_{loop} = 14$ per 4-simplex) and the global count ($N_\Sigma \sim 10^{123}$) is structural, not numerical. Both are enumerations of 2-dimensional geometric carriers: triangular faces locally, Planck-area cells globally. The UV combinatorics determines the *type* of geometric information carrier (2-cells with holonomy); the IR discretization determines how many such carriers tile the closure surface. This is a dimensional homology—a shared face-based combinatorial structure at vastly different scales—not an arithmetic identity.

---

4. Chern–Weil Closure Condition

4.1 Topological constraint

For a principal U(1) bundle over a closed oriented 2-surface Σ, the Chern–Weil theorem yields the integrality condition:

$(1/2\pi) \int_\Sigma F = n \in \mathbb{Z}$

where F is the curvature 2-form and n is the first Chern number (c_1) of the bundle—the total winding number. This is a topological invariant: it cannot change under continuous deformations of the connection. (We use n rather than χ for the Chern number to avoid confusion with the Euler characteristic $\chi(\mathcal{K}^2)$ used in §3.1.)

4.2 Discretization

Decompose Σ into N_Σ cells, with per-cell holonomy θ_i. The discretized Chern–Weil condition is:

$(1/2\pi) \Sigma_i \theta_i = n \in \mathbb{Z}$

This exact integrality holds regardless of whether the θ_i are uniform across cells. In the isotropic leading-order approximation, $\theta_i = \theta$ for all i, giving:

$N_\Sigma \cdot \theta = 2\pi n$

4.3 Holonomy–coupling identification

Define the dimensionless per-cell holonomy normalization $\hat{\alpha}$ by $\theta := 2\pi\hat{\alpha}$. The closure condition becomes:

$N\_\Sigma \cdot \hat{\alpha} = n$

At leading order, we identify $\hat{\alpha}$ with the electromagnetic fine-structure constant evaluated at a reference scale: $\hat{\alpha} := \alpha_0 \equiv \alpha(\mu_0)$. This is bridge hypothesis BH1. Its justification rests on three converging lines of reasoning:

Minimality. Under SA3, the coherence predicate introduces exactly one dimensionless U(1) coupling controlling holonomy normalization. The identification $\hat{\alpha} = \alpha_0$ does not introduce new parameters but equates the framework's sole dimensionless coupling with a physical one.

Infrared universality. The closure surface is the de Sitter horizon—a macroscopic, late-time structure. The holonomy that accumulates over this surface must couple universally to phase transport of all charged matter crossing the horizon. In late-time cosmology, electromagnetism is the *only* unbroken, long-range U(1) gauge interaction. The Standard Model hypercharge U(1)_Y is broken by electroweak symmetry breaking and does not survive as a long-range force below the weak scale. Any hypothetical hidden U(1) would require new charged fields, violating SA3 (minimality: no additional independent parameters). Therefore, the only U(1) coupling that can govern macroscopic phase accumulation on the closure surface is the electromagnetic fine-structure constant.

Uniqueness under constraints. Combining SA3 (one coupling) with the requirement of infrared universality (macroscopic, unbroken, long-range) leaves exactly one candidate: α. The identification is therefore the unique SA3-compatible bridge between the framework's abstract holonomy normalization and known physics.

Normalization and reference scale. The identification $\theta = 2\pi\hat{\alpha}$ fixes the normalization: the per-Planck-cell holonomy is 2π times the dimensionless coupling. Why $\hat{\alpha}$ rather than e, $4\pi\alpha$, or some other function of α? The answer is that $\hat{\alpha}$ must be small (near-identity holonomy is required for robustness under SA2; large holonomy per cell would place the system near the coherence boundary), and $\hat{\alpha}$ must be the *dimensionless* ratio that enters the holonomy exponent directly: $\theta = 2\pi\hat{\alpha}$ is the unique parameterization in which $\hat{\alpha}$ is the holonomy fraction of 2π per cell, with no arbitrary numerical prefactors. A hidden U(1) with its own coupling would introduce additional charged fields and an independent coupling constant, violating SA3. Finally, the reference scale μ_0 is not a free choice: RG invariance of n (§5) requires θ to be evaluated at a fixed physical reference, and the macroscopic IR meaning of phase transport on the closure surface anchors this to the Thomson limit $\alpha_0 \equiv \alpha(\mu \to 0)$. Any other scale choice yields the same n after RG conversion (§5.7), confirming that the physical content is scale-independent.

Normalization lemma. Under SA3 (no extra dimensionless parameters), the near-identity holonomy angle must scale linearly with the unique small dimensionless coupling: $\theta(\hat{\alpha}) = 2\pi\, c\, \hat{\alpha} + O(\hat{\alpha}^2)$, where c is a numerical constant. The constant c cannot encode new physics without introducing a new dimensionless parameter, which SA3 forbids. Defining $\hat{\alpha}$ as the fraction of a full 2π rotation per cell fixes c = 1 by construction: $\hat{\alpha} = \theta/(2\pi)$. This is not a physical assumption but a canonical parameterization—the unique normalization convention in which $\hat{\alpha}$ directly measures the holonomy fraction per cell. The higher-order corrections $O(\hat{\alpha}^2)$ are precisely the RG-generated terms analyzed in §5, whose coefficients are scale-dependent rather than universal (§5.5).

With this identification:

$N_\Sigma(\Lambda) \cdot \alpha_0 = n \in \mathbb{Z}$

This is the closure equation jointly constraining α and Λ. Since $N_\Sigma = 12\pi/(\Lambda \ell_P^2)$, it reads:

$12\pi\, \alpha_0 / (\Lambda\, \ell_P^2) = n$

The left-hand side must be an integer. This is a non-trivial admissibility condition: not all pairs (α, Λ) are permitted. The closure equation constrains α and Λ; it does not uniquely determine them but restricts them to a discrete family of admissible pairs labeled by n.

With standard inputs ($\alpha_0 \approx 1/137.036$, $\Lambda \approx 1.1 \times 10^{-52}$ m^{-2}), the inferred winding number is $n \approx 10^{121}$.

4.4 Why the coupling constant appears in the topological integral

In standard gauge theory, topological invariants such as Chern numbers classify bundles but do not depend on coupling constants—the coupling appears in the action, not in the topology. The present framework departs from this in a specific and controlled way, and the departure deserves explicit justification.

In conventional treatments, the curvature 2-form F is a geometric object defined by the connection, independent of the coupling. The Chern integral $(1/2\pi)\int F$ yields an integer that characterizes the bundle topology. The coupling g enters only when one writes the action $S \sim (1/g^2)\int F \wedge \star F$, which weights the dynamics but does not affect the topology.

In the closure framework, the situation is different because the connection is not specified independently of the coupling. The per-cell holonomy θ_i is the *only* continuous datum characterizing the U(1) bundle on each Planck-area cell (SA3, SA5). There is no separate "connection" specified first and "coupling" inserted afterward; the only continuous datum in the holonomy sector is the per-cell curvature/holonomy scale, and bridging it to the measured electromagnetic coupling identifies that datum with α_0. When we write $\theta = 2\pi\hat{\alpha}$, we are not inserting a coupling into a topological formula; we are recognizing that the total curvature integral $(1/2\pi)\Sigma_i\, \theta_i$ is built from per-cell contributions whose magnitude is controlled by the holonomy normalization.

The analogy is to a lattice gauge theory in which the plaquette action is $U_p = \exp(i\theta_p)$, and the topological charge is $Q = (1/2\pi)\Sigma_p\, \theta_p$. On the lattice, the coupling $\beta = 1/g^2$ controls the typical magnitude of θ_p, and therefore controls the value of Q in any given configuration. The coupling does not appear in the *definition* of Q, but it determines which configurations dominate—and in a closure framework where only one configuration (the coherent one) is admitted, the coupling directly sets Q.

In short: the coupling enters the topological integral not because topology depends on coupling in general, but because in this framework the per-cell holonomy normalization is the only continuous datum in the holonomy sector, and the Chern number is the sum of those per-cell curvature contributions.

To state this precisely: the first Chern number does not depend on coupling constants in general; it depends on the curvature of the specific connection realized on Σ. In the present framework, the per-cell curvature/holonomy normalization is the only continuous datum admitted in the holonomy sector (SA3, SA5), so specifying the physical connection on Σ is equivalent to specifying this normalization. The appearance of α is therefore not an insertion of dynamics into topology, but a bridge identifying the unique dimensionless holonomy-strength parameter of the realized connection with the measured electromagnetic coupling at a fixed reference condition.

---

5. Renormalization-Group Stability of the Closure Condition

5.1 The problem

The topological charge n is an integer and therefore RG-invariant—it cannot depend on the renormalization scale μ. However, the electromagnetic coupling α(μ) runs with μ. If the closure condition involves α, how can n remain scale-independent?

This section shows that the resolution is natural: the holonomy θ must be evaluated at a fixed physical reference scale μ_0, and any re-expression in terms of α(μ) at a different scale introduces logarithmic, scale-dependent corrections that precisely compensate the running.

5.2 Why a scale-independent polynomial ansatz fails

A naive approach would write $\theta(\alpha) = 2\pi(\alpha + c_2\alpha^2 + c_3\alpha^3 + \cdots)$ with μ-independent coefficients and demand:

$$d/d \ln \mu \, [N\_\Sigma \cdot \theta(\alpha(\mu))] = 0$$

Since N_Σ is μ-independent, this requires $\theta'(\alpha) \cdot \beta\_\alpha(\alpha) = 0$. But in QED, $\beta\_\alpha \neq 0$ for $\alpha \neq 0$, so this forces $\theta'(\alpha) = 0$—implying θ is constant, which contradicts θ depending on α at all.

The error is in assuming the coefficients $c_2, c_3, \ldots$ are μ-independent. They cannot be, because α(μ) is a scheme-dependent quantity while n is scheme-independent. The correct resolution involves evaluating α at a fixed reference scale.

5.3 QED one-loop running

At one loop, the QED beta function is:

$d\alpha/d \ln \mu = \beta\_\alpha(\alpha) = (2/3\pi)\, b\, \alpha^2 + O(\alpha^3)$

where $b = \Sigma\_f\, Q\_f^2$ is the sum of squared charges over active fermion species. This is a schematic, piecewise-effective beta function: the fermion content changes across mass thresholds, and b is understood as the value appropriate to the scale μ. We are not treating electroweak mixing or scheme-dependent subtleties, which would affect sub-leading corrections but not the qualitative structure of the argument.

The fermion content at representative scales:

- At Thomson limit ($\mu_0 \to 0$): only the electron contributes, giving $b = Q\_e^2 = 1$.
- At $\mu \approx M\_Z \approx 91.2$ GeV: all Standard Model charged fermions contribute. Quarks give 3 colors × $[(2/3)^2 + (1/3)^2]$ × 3 generations = 3 × (5/9) × 3 = 5. Charged leptons give 1^2 × 3 generations = 3. Total: b = 8.

Integrating the one-loop equation from reference scale μ_0:

$1/\alpha(\mu) = 1/\alpha(\mu_0) - (2b/3\pi) \ln(\mu/\mu_0) + O(\alpha)$

Define $\alpha_0 \equiv \alpha(\mu_0)$ at the Thomson limit. This is a fixed, physical, RG-invariant quantity once μ_0 is specified: it does not change when one varies μ.

Scope of the RG treatment. We use one-loop QED running to demonstrate the structural point: RG invariance of n forces θ to be referenced to $\alpha(\mu_0)$. A precision computation of the $\alpha(0) \leftrightarrow \alpha(M\_Z)$ matching—including hadronic vacuum polarization contributions, electroweak mixing effects, and scheme dependence—is a technical exercise deferred to future work and does not change the conceptual constraint. The qualitative conclusion (logarithmic compensation, scale-independence of n) is robust under higher-order corrections; what changes at higher loops is the numerical precision of the α_0 extraction, not the structure of the argument.

5.4 Resolution: θ is anchored to α_0

The closure condition is a macroscopic, global constraint. The natural choice is to evaluate the holonomy at the low-energy reference scale where macroscopic phase accumulation is defined:

$\theta = 2\pi\, \hat{\alpha}$ with $\hat{\alpha} := \alpha_0$

Then $n = N\_\Sigma\, \alpha_0$ is RG-stable by construction: it contains no dependence on μ because μ plays no role in the definition of α_0.

5.5 Re-expressing θ in terms of α(μ): the correct c_2

One may wish to express the reference holonomy in terms of α measured at a different scale μ. Inverting the one-loop running relation:

$\alpha_0 = 1 / (1/\alpha(\mu) + (2b/3\pi) \ln(\mu/\mu_0)) + O(\alpha^3)$

Expanding perturbatively for small $\alpha(\mu)$:

$\alpha_0 = \alpha(\mu) - (2b/3\pi) \ln(\mu/\mu_0) \cdot \alpha(\mu)^2 + O(\alpha(\mu)^3)$

Therefore, if one insists on writing $\theta = 2\pi(\alpha + c_2\alpha^2 + \cdots)$ in terms of α evaluated at scale μ, the quadratic coefficient is:

$c_2(\mu; \mu_0) = -(2b/3\pi) \ln(\mu/\mu_0)$

This is not a universal constant. It is logarithmic in the scale ratio, proportional to the one-loop coefficient, and changes when μ or μ_0 changes. For physically relevant separations, c_2 is naturally O(1–10), not the large values (e.g., $c_2 \approx -68$) that would appear if one incorrectly treated it as μ-independent.

5.6 Differential formulation

Define the RG-invariant holonomy $\theta(\mu) := 2\pi\alpha_0$ and write $\theta(\mu) = 2\pi\alpha(\mu)Z(\mu)$, where Z is a renormalization factor. RG invariance $d\theta/d \ln \mu = 0$ requires:

$d \ln Z / d \ln \mu = -\beta\_\alpha(\alpha)/\alpha$

At one loop, $\beta\_\alpha(\alpha)/\alpha = (2b/3\pi)\alpha$, which integrates to:

$Z(\mu) = \exp(-(2b/3\pi) \int \alpha \, d \ln \mu')$

This yields a logarithmic Z consistent with the perturbative expansion above. The compensation of running is logarithmic—a scheme/scale structure—not a polynomial correction.

5.7 Internal consistency test

The closure analysis provides a concrete, falsifiable check. Given any measurement of α at scale μ:

1. Convert to α_0 using the standard RG equation with appropriate piecewise fermion-threshold matching (b = 1 below the muon threshold, increasing stepwise as heavier charged fermions become active).
2. Compute $n\_inferred = N\_\Sigma \cdot \alpha_0$.
3. Verify that n_inferred is μ-independent.

The qualitative structure of this test is clear: the running of α between the Thomson limit and the Z-pole scale is well established experimentally ($\alpha(0) \approx 1/137.036$ vs. $\alpha(M\_Z) \approx 1/127.95$), and both must yield the same n after proper conversion. A full numerical demonstration requires careful piecewise threshold matching across all Standard Model charged fermion masses, which we defer to a dedicated analysis.

---

6. Discrete Selection of K = 7 from Closure Consistency

The preceding sections treated K = 7 as an enumerated structural input. We now show that, when the closure equation is combined with two additional mechanisms internal to the framework—a discrete coupling formula $\alpha(K)$ (SI1) and a vacuum regulation relation $\Lambda(K)$ (SI2–SI3)—the constraint count K = 7 is not merely assumed but *selected* as the unique integer compatible with the observed cosmological hierarchy. This is a selection theorem at hierarchy level, robust to O(1) uncertainties in structural prefactors.

6.1 Strategy

The closure equation $n = N_\Sigma \alpha$ relates the topological winding number to both Λ (through N_Σ) and α. If both α and Λ can be expressed as functions of the discrete parameter K, then n becomes a function of K alone (plus fixed structural constants). We will show that n(K) is super-exponentially sensitive to K, so that changing K by ±1 shifts n by 50–100 orders of magnitude. Only K = 7 places n in the regime consistent with observation.

The key requirement for this argument to be non-circular is that $\alpha(K)$ and $\Lambda(K)$ arise from *independent* physical mechanisms—not from fitting to the observed values of α and Λ. The coupling formula $\alpha(K)$ derives from discrete coherence counting on the simplicial 2-skeleton (§6.2). The vacuum regulation $\Lambda(K)$ derives from a Two-Planck saturation mechanism linking the cosmological constant to a coherence scale $\xi(K)$, which is itself determined by dimensional transmutation on the simplicial lattice (§6.3). These are logically independent chains of reasoning whose joint consistency at K = 7 constitutes a prediction, not a calibration.

Independence of the three input chains.

Input	Origin	What it determines	Depends on Λ data?	Free parameters
SI1: $\alpha(K)$	Local simplicial combinatorics: constraint counting on 2-skeleton, binary admissibility	Electromagnetic coupling at each K	No	None (K is the only input)
SI2: $\Lambda = 8\pi\ell_P^2/\xi^4$	IR vacuum saturation: Two-Planck regulation of vacuum energy density	Cosmological constant given ξ	No	None (dimensional analysis)
SI3: $\xi(K)$	Lattice dimensional transmutation: one-loop RG from $g_0^2 = 2^{-K}$ to continuum	Coherence/regulation scale at each K	No	$p_c \in [0.17, 0.30]$ (bounded)

SI1 is a purely combinatorial formula. SI2 is a dimensional relation. SI3 is a standard lattice-to-continuum transmutation with a single bounded structural parameter. None were fit to observation. The intersection of all three at K = 7—simultaneously producing $\alpha \approx 1/137$, $\Lambda \approx 10^{-52}$ m^{-2}, and $n \approx 10^{121}$—is the non-trivial content of the selection theorem.

6.2 The discrete coupling formula α(K)

In the coherence framework, the fine-structure constant at constraint count K is determined by the combinatorics of the simplicial 2-skeleton. The local loop count is $N_{loop} = 2K$ (the total number of curvature-carrying faces plus independent closure channels in the simplicial building block).

Remark on $N_{loop} = 2K$. For K = 7, the geometric calculation of §3.1 gives $N_{loop} = 10 + 4 = 14 = 2 \times 7$, confirming the relation in the case where it can be independently computed. The general formula $N_{loop} = 2K$ extends this to arbitrary K as a structural hypothesis (SI1): the loop content of the local building block scales linearly with the constraint count. This is an assumption of the discrete coupling program, not a derivation from first principles. (See §7.4 for further discussion.)

The coupling is then determined by the leading-order discrete closure condition:

$$\alpha^{-1}(K) = 2^K \cdot (2K + 1) / (2K)$$

This formula encodes the requirement that the holonomy normalization be consistent with K independent constraints distributed over $N_{loop} = 2K$ channels, with the factor 2^K reflecting the binary admissibility structure (each of K constraints has two states).

For K = 7:

$$\alpha^{-1}(7) = 128 \cdot 15 / 14 = 1920 / 14 \approx 137.14$$

This is within 0.08% of the measured value $\alpha^{-1} \approx 137.036$. The proximity is striking but should be interpreted cautiously: the formula is leading-order, and sub-leading corrections (analogous to higher-loop terms) have not been computed.

6.3 Two-Planck vacuum regulation and ξ(K)

The cosmological constant is related to a coherence regulation scale ξ through the Two-Planck saturation mechanism (SI2). The regulated vacuum energy density is $\rho_{vac} \sim \hbar c/\xi^4$, and Einstein's equation gives $\Lambda = 8\pi G\rho_{vac}/c^4$. Substituting $\ell_P^2 = \hbar G/c^3$:

$$\Lambda = 8\pi \, \ell_P^2 / \xi^4$$

This replaces the Λ-dependence in the cell count with a ξ-dependence:

$$N_\Sigma = 12\pi / (\Lambda \, \ell_P^2) = 12\pi / ((8\pi\ell_P^2/\xi^4) \cdot \ell_P^2) = (3/2) \cdot \xi^4 / \ell_P^4$$

The coherence scale ξ is determined by dimensional transmutation on the simplicial lattice (SI3, Route-M mechanism). The lattice coupling $g_0^2 = 2^{-K}$ runs to the continuum via:

$$\ln(\xi / \ell\_em) = (1 / 2b) \cdot (1/g_0^2 - 1/p\_c)$$

where $\ell\_em = 2\ell\_P$ is the emergence scale (the minimal resolved length in the Two-Planck framework; not to be confused with the electron Compton wavelength $\lambda\_C = \hbar/(m\_e\, c)$), $b = N\_loop/16 = 2K/16 = K/8$ is the one-loop lattice beta-function coefficient, and p_c is the percolation critical coupling. With $g_0^2 = 2^{-K}$:

$$\ln(\xi / \ell\_em) = (4/K) \cdot (2^K - 1/p\_c)$$

Therefore:

$$\xi(K) = \ell\_em \cdot \exp[(4/K)(2^K - 1/p\_c)]$$

The exponential dependence on 2^K makes $\xi(K)$ extraordinarily sensitive to K.

6.4 The winding number as a function of K

Combining the results of §6.2 and §6.3, the topological winding number is:

$$n(K) = N\_\Sigma(K) \cdot \alpha(K) = (3/2) \cdot (\xi(K)^4 / \ell\_P^4) \cdot \alpha(K)$$

This is now a function of K alone, with all other quantities being fixed physical constants (ℓ_P, ℓ_em, p_c). The dominant K-dependence enters through $\xi(K)^4$, which scales as $\exp[(16/K)(2^K - 1/p\_c)]$. Since 2^K grows exponentially in K while the prefactor 16/K varies slowly, n(K) is a super-exponential function of K.

6.5 Numerical consistency: computed values

Rather than relying on approximate prose estimates, we compute n(K) directly from the formulas of §6.2–6.4.

Numerical conventions. We take $\ell\_P = 1.616255 \times 10^{-35}$ m (2018 CODATA), $\ell\_em = 2\ell\_P$, and $\Lambda = 1.1 \times 10^{-52}\ m^{-2}$. All logarithms in the tables are base 10. The percolation threshold p_c is treated as a bounded structural parameter; we scan $p\_c \in [0.17, 0.30]$ as representative of standard site-percolation critical probabilities on comparable lattices (triangular lattice: $p\_c \approx 0.5$; face-centered cubic: $p\_c \approx 0.20$; bond percolation on comparable graphs spans the quoted range). See Appendix B.4.

Table 1. Computed values for p_c = 0.20 (representative Route-M percolation threshold).

K	$\alpha^{-1}(K)$	$\log_{10} \xi$ (m)	$\log_{10}(\xi/\ell\_P)$	$\log_{10}(\xi^4/\ell\_P^4)$	$\log_{10} n(K)$
6	69.3	−17.4	17.4	69.5	67.9

K	α^{-1}(K)	$\log_{10}$ ξ (m)	$\log_{10}$(ξ/ℓ_P)	$\log_{10}$(ξ^4/ℓ_P^4)	$\log_{10}$ n(K)
7	137.1	−4.0	30.8	123.3	121.3
8	272.0	+20.0	54.8	219.2	217.0

Table 2. Sensitivity of $\log_{10}$ n(K) to p_c across the Route-M percolation range.

K	p_c = 0.17	p_c = 0.20	p_c = 0.25	p_c = 0.30
6	66.8	67.9	69.0	69.8
7	120.5	121.3	122.3	123.0
8	216.2	217.0	217.8	218.4

The observed cosmological value is $\log_{10}$ n ≈ 121 (from N_Σ α_0 with measured inputs). The tables demonstrate:

- **K = 7 matches.** Across the full p_c range, $\log_{10}$ n(7) = 120.5–123.0, spanning the observed value.
- **K = 6 undershoots by ~53 orders of magnitude.** $\log_{10}$ n(6) ≈ 67–70.
- **K = 8 overshoots by ~96 orders of magnitude.** $\log_{10}$ n(8) ≈ 216–218.
- **p_c variation is negligible.** Across p_c = 0.17–0.30, the shift in $\log_{10}$ n is ~3 orders—far smaller than the ~50–100 order gaps between adjacent K values.

6.6 Robustness of the selection

The K-selection is robust against moderate uncertainties in the input parameters for two reasons.

Exponential leverage. The dominant factor ξ(K)^4/ℓ_P^4 depends on K through exp[(16/K) · 2^K], which changes by a multiplicative factor of order exp[16 · $2^K/K^2$] when K shifts by 1 near K = 7. No O(1) uncertainty in p_c, b, or the prefactors can compensate a gap of 10^{50} or more.

Asymmetric gaps. The gap is larger above K = 7 (~96 orders) than below (~53 orders), reflecting the accelerating growth of 2^K. This asymmetry reinforces the selection: K = 7 is not on the boundary of viability but sits squarely in a deep well of the $\log_{10}$ n landscape.

Structural rigidity. The selection depends on the exponential dependence on 2^K in the transmutation formula, not on the specific numerical prefactor. Any modification of SI3 that preserves the exp(const · 2^K) structure—for example, replacing the prefactor 4/K with 3/K or 5/K, or shifting the percolation threshold within physically motivated bounds—retains the selection rigidity. The inter-K gaps are set by the doubly exponential growth of ξ^4, which overwhelms any O(1) changes to coefficients.

6.7 Logical status of the selection theorem

Theorem 2 (K = 7 selection at hierarchy level). Assume: (i) the closure condition $n = N_\Sigma \alpha$ with $N_\Sigma = 12\pi/(\Lambda \ell_P^2)$; (ii) Two-Planck vacuum regulation $\Lambda = 8\pi\ell_P^2/\xi^4$ (SI2); (iii) the discrete coupling formula $\alpha(K) = 2K / [2^K(2K + 1)]$ (SI1); and (iv) the dimensional transmutation relation $\xi(K) = \ell_{em} \exp[(4/K)(2^K - 1/p_c)]$ (SI3). Then n(K) is a super-exponential function of K, and K = 7 is the unique integer for which n falls in the cosmologically observed regime $\log_{10} n \approx 121$.

What this proves (conditionally). If all four inputs are valid, then K = 7 is not a free parameter but a determined output of the closure system. The constraint count, the coupling constant, and the cosmological constant are jointly constrained by the admissibility condition, locked up to the discrete topological class n and controlled corrections.

What this does not prove. The theorem does not derive K = 7 from closure alone without the Two-Planck relation and the α formula. It does not establish that $n \sim 10^{121}$ is the correct target beyond order-of-magnitude consistency. And it does not prove exact integrality of n, which would require precision in α, Λ, and the O(1) prefactors far beyond current capability.

Independence of the input chains. The α(K) formula derives from the combinatorics of the simplicial 2-skeleton (constraint counting and binary admissibility structure). The ξ(K) formula derives from lattice dimensional transmutation (one-loop running from the lattice coupling $g_0^2 = 2^{-K}$ to the continuum). These are logically and physically distinct mechanisms. Their joint consistency at K = 7—producing both $\alpha \approx 1/137$ and $\Lambda \approx 10^{-52}\ m^{-2}$—is a non-trivial convergence, not a calibration.

---

7. Discussion

7.1 Summary of the logical chain

The architecture of the argument forms a closed sequence:

1. **K-constraint coherence** → existence of an independent holonomy constraint C4
2. **Exclusion lemmas** → minimal gauge redundancy is U(1)
3. **2-skeleton combinatorics** → geometric information propagates on 2-cells; $N_{loop} = 14$ at K = 7
4. **de Sitter surface selection** → unique IR closure surface $\Sigma \cong S^2$ (conditional on BH3)
5. **Planck-area discretization** → global cell count $N_\Sigma \approx 1.3 \times 10^{123}$
6. **Chern–Weil integrality** → $N_\Sigma \cdot \alpha_0 = n \in \mathbb{Z}$ (admissibility, not unique determination)

7. **RG stability** → holonomy anchored to reference scale; logarithmic compensation of running

8. **Discrete selection** → K = 7 uniquely compatible with $\log_{10} n \approx 121$; K ± 1 fails by 50–100 orders

Steps 1–7 establish the closure condition for general K. Step 8 closes the loop by selecting K = 7 as the unique solution of the full system at hierarchy level.

7.2 What is derived vs. what is assumed

Assumptions: SA1–SA5 (including SA3′), BH1–BH3, SI1–SI3 (see §1.1 for the complete ledger).

Derived consequences:

- U(1) as the minimal gauge redundancy (Theorem 1, via Lemmas 1–3 and Corollary 3).
- The de Sitter 2-sphere as the unique closure surface (Proposition in §3.2).
- The joint admissibility condition on α and Λ (the closure equation).
- The logarithmic, scale-dependent structure of higher-order holonomy corrections (§5.5–5.6).
- The μ-independence of n as an internal consistency test (§5.7).
- K = 7 as the unique admissible constraint count at hierarchy level (Theorem 2).

7.3 Scope of the α–Λ constraint

The closure equation jointly constrains α and Λ by an admissibility condition: they are restricted to a discrete family of pairs labeled by $n \in \mathbb{Z}$, but are not uniquely determined by closure alone. Note that $n \approx N_\Sigma \alpha_0 \approx 10^{123} / 137 \approx 10^{121}$, which is smaller than N_Σ itself by a factor of $\alpha_0 \approx 1/137$; the "cosmological hierarchy number" in this framework is therefore $\sim 10^{121}$, not the $\sim 10^{122}$ sometimes quoted for N_Σ alone. The remaining fundamental constants G, ℏ, and c enter through the Planck area $\ell_P^2 = \hbar G/c^3$, which sets the discretization scale. This is a parametric dependence—dimensional analysis, not a dynamical constraint. Whether a deeper formulation can promote the dependence on G, ℏ, and c from parametric to dynamical remains an open question.

A note on the hierarchy exponent. The cosmological constant problem is often stated as a $\sim 10^{122}$ discrepancy between the Planck-scale vacuum energy and the observed value. The winding number $n \approx N_\Sigma \alpha_0 \approx 10^{121}$ is smaller than $N_\Sigma \approx 10^{123}$ by a factor of $\alpha_0 \approx 1/137$, placing it one to two orders below the cell count itself. This is not a discrepancy; it is a direct consequence of the closure equation. The "10^{120}–10^{122}" range quoted in the literature reflects different conventions for what is being counted (area in Planck units, vacuum energy ratio, etc.).

In the present framework, the relevant quantity is the Chern number $n = N_\Sigma \alpha_0$, and its value $\log_{10} n \approx 121$ is a derived output, not a target.

7.4 The N_loop = 2K coincidence

For K = 7, the geometric calculation of §3.1 independently yields $N_{loop} = 14 = 2 \times 7$. The discrete coupling formula assumes $N_{loop} = 2K$ as a general relation. Whether this linear relationship holds for all K, or is a special property of K = 7 tied to the specific structure of the 4-simplex 2-skeleton, is an open question with two possible resolutions:

- If $N_{loop} = 2K$ holds only at K = 7, this would constitute an additional, independent selection mechanism further constraining the framework. The coincidence itself would then require explanation.
- If $N_{loop} = 2K$ holds generally, its origin should be derivable from the combinatorics of K-constraint simplicial complexes—a direction for future work.

Either resolution strengthens the framework; neither undermines it.

7.5 Relation to existing frameworks

The use of simplicial decompositions connects to Regge calculus and dynamical triangulations. The U(1) bundle structure and Chern–Weil integrality are standard tools in gauge theory and fiber bundle geometry. The RG analysis uses standard one-loop QED (treated schematically with piecewise thresholds; electroweak mixing and scheme subtleties are deferred). The dimensional transmutation from lattice to continuum coupling is standard in lattice gauge theory. The novelty lies in combining these elements into a single closure condition that links UV coherence combinatorics to IR topological constraints, yielding relationships among α, Λ, and K that are absent in each framework individually.

7.6 Falsifiability

The framework makes the following testable predictions:

1. **Integrality (convergence form)**: $N_\Sigma \cdot \alpha_0$ must be an integer. In practice, this means: as measurements of α and Λ improve, the quantity $n_{inferred} = 12\pi\alpha_0/(\Lambda \ell_P^2)$ should converge toward a stable integer within the propagated uncertainty from Λ and α_0. Current uncertainties in Λ (~few percent) dominate and permit n to be determined only at the order-of-magnitude level ($\sim 10^{121}$). The prediction is that future precision will not reveal n drifting away from integrality but converging toward it.
2. **Scale-independence**: n inferred from α measured at different scales μ must agree after RG conversion with proper threshold matching.

3. **Robustness under refinement**: the inferred n should be stable under changes in discretization conventions (e.g., replacing ℓ_P^2 with $\gamma\sqrt{3}\ \ell\_P^2$), up to O(1) multiplicative factors that do not affect K-selection.
4. **K-selection**: no self-consistent solution of the full system exists for $K \neq 7$ within the observed cosmological hierarchy.
5. **α prediction**: the leading-order formula $\alpha^{-1}(7) \approx 137.14$ should be correctable to the measured value by computable sub-leading terms.
6. **de Sitter conditionality**: observation of $w \neq -1$ asymptotically would require revision of the closure surface and potentially falsify the specific form of the admissibility condition.

If improved measurements of Λ yield a value of $N\_\Sigma \cdot \alpha_0$ that is not close to an integer, or if a self-consistent solution at $K \neq 7$ is found, the framework is falsified. By "close" we mean within the propagated 1σ uncertainty from Λ (dominant) and α_0; at present this uncertainty is order-percent in Λ, so the integrality prediction is testable only at order-of-magnitude level ($\sim 10^{121}$). The prediction sharpens as cosmological measurements improve. To be concrete: current Planck/DESI-level constraints place Λ at ~5% precision, translating to a ~5% uncertainty in n — i.e., $\log_{10} n \approx 121.0 \pm 0.02$. This is far too coarse to test exact integrality (which would require knowing n to ±1 out of $\sim 10^{121}$), but it is more than sufficient for K-selection: the nearest competitor ($K = 6$ at $\log_{10} n \approx 68$) is excluded by over 50 orders of magnitude.

Appendix A: Formal Theorem Statements

Theorem A (Minimal gauge redundancy). In any relational pre-geometric substrate where (i) coherent triangles carry K = 7 independent constraints including a single holonomy-based loop closure C4, (ii) coherence predicates are gauge-invariant (SA1), (iii) transport data admit continuous perturbations with robust coherence (SA2), (iv) no additional independent parameters are introduced beyond a single holonomy-strength coupling (SA3), (v) co-face compatibility and loop closure are decidable from a single scalar holonomy invariant without extra alignment structure (SA3′), (vi) holonomy classes form a compact space (SA4), and (vii) the near-identity holonomy spectrum is one-parameter (SA5), U(1) is the minimal compact gauge redundancy consistent with all stated assumptions, and is the unique minimal compact connected choice.

Proof: Lemmas 1–3 and Corollary 3 of §2.3, composed as Theorem 1.

Theorem B (Bundle extension). If coherent triangular faces percolate to form a connected macroscopic 2-complex, and the macroscopic geometry admits a closed oriented 2-surface Σ, then the local U(1) transport structure extends to a principal U(1) bundle over Σ.

Proof sketch. Principal U(1) bundles over a simplicial complex $\mathcal{K}$ are classified by the first Čech cohomology group $\check{H}^1(\mathcal{K}, U(1))$. On the 2-skeleton, a U(1) transport system defines transition functions on overlaps of vertex stars. The cocycle condition $g_{ij} g_{jk} g_{ki} = 1$ on triple overlaps is precisely the content of the triangle holonomy constraint C4 (with holonomy H_Δ encoding the cocycle failure). Coherence ($H_\Delta \in \mathcal{C}$ near the identity) ensures the transition functions satisfy the cocycle condition up to controlled corrections. Percolation ensures connectivity, allowing the local data to define a global Čech cocycle and hence a principal U(1) bundle over Σ. ■

Theorem C (Joint constraint). Let $\theta(\alpha)$ be the per-cell holonomy and $N_\Sigma = 12\pi/(\Lambda \ell_P^2)$ the Planck-area cell count. Chern–Weil integrality yields the admissibility condition:

$$N_\Sigma(\Lambda) \cdot \theta(\alpha) = 2\pi n, \; n \in \mathbb{Z}$$

This is a single geometric closure functional jointly constraining α and Λ to a discrete family of admissible pairs.

Theorem D (Scheme invariance). For any two renormalization schemes s, s′ related by an analytic redefinition $\alpha_{s'} = f(\alpha_s)$, topological invariance of n requires $\theta_s(\alpha_s) = \theta_{s'}(\alpha_{s'})$ when both are evaluated at a common physical reference condition.

Proof. Scheme changes are analytic reparameterizations of the coupling. Since $n = (N_\Sigma/2\pi)\theta$ is a topological integer, it must be invariant. Therefore θ transforms covariantly under scheme redefinitions: $\theta_s(\alpha_s) = \theta_{s'}(f(\alpha_s))$, which constrains θ to be either evaluated at a fixed physical reference or constructed from scheme-invariant combinations. ■

Theorem E (K = 7 selection at hierarchy level). Under the closure condition (Theorem C), Two-Planck vacuum regulation (SI2), the discrete coupling formula (SI1), and the dimensional transmutation relation (SI3), the winding number $n(K) = (3/2)(\xi^4/\ell_P^4)\alpha(K)$ is a super-exponential function of K. The unique integer K for which n falls in the cosmologically observed regime $\log_{10} n \approx 121$ is K = 7 (see Table 1, §6.5). This selection is robust to O(1) uncertainties in structural prefactors and to variation of p_c across the Route-M percolation range.

Appendix B: Supplementary Material

B.1 Euler characteristic calculation for the 2-skeleton

The 2-skeleton $\mathcal{K}^2$ of the standard 4-simplex $\sigma^4 = \{0,1,2,3,4\}$ consists of all vertices, edges, and triangles. Its Euler characteristic is:

$$\chi(\mathcal{K}^2) = 5 - 10 + 10 = 5$$

The 1-skeleton is the complete graph K_5. Every 1-cycle in K_5 bounds a 2-chain in $\mathcal{K}^2$: any cycle decomposes into 3-cycles (triangles), and every triangle of K_5 is a 2-face of $\mathcal{K}^2$. (Equivalently, $\mathcal{K}^2$ is the 2-skeleton of a contractible simplex and is therefore simply connected.) Therefore $H_1(\mathcal{K}^2; \mathbb{Z}) = 0$, giving $\beta_1 = 0$. With $\beta_0 = 1$ ($\mathcal{K}^2$ is connected), the Euler relation gives:

$\beta_0 - \beta_1 + \beta_2 = \chi(\mathcal{K}^2)$ $1 - 0 + \beta_2 = 5$ $\beta_2 = 4$

The 4 independent 2-cycles correspond to the 4 independent global Bianchi-type relations among the 10 face holonomies. This is consistent with the fact that the 4-simplex has 5 tetrahedral 3-faces imposing 5 relations, of which 4 are independent (the 5th being a linear combination of the others, reflecting the single global constraint from the 4-simplex itself).

B.2 Non-uniform holonomy

If holonomy varies cell-to-cell, the Chern–Weil discretization gives:

$n = (1/2\pi) \Sigma_{i=1}^{N} \theta\_i$

Integrality of n is preserved: it is a topological invariant independent of how the curvature is distributed among cells. The uniform case $\theta\_i = \theta = 2\pi\hat{\alpha}$ for all i is the isotropic leading-order approximation. Corrections from non-uniformity affect the relationship between n and the average coupling $\langle\hat{\alpha}\rangle$ but do not alter the integrality constraint itself. To leading order, $n = N\_\Sigma \langle\hat{\alpha}\rangle$, where $\langle\hat{\alpha}\rangle = (1/N\_\Sigma) \Sigma_i \hat{\alpha}\_i$ is the cell-averaged holonomy normalization.

B.3 Standing assumptions, bridge hypotheses, and structural inputs

See §1.1 for the complete assumption ledger.

B.4 Sensitivity of n(K) to the percolation threshold

See Tables 1–2 in §6.5 for the computed sensitivity analysis.

B.5 Distinction between local C7 and global Bianchi relations

To prevent confusion between the per-triangle constraint C7 (§2.2) and the global closure relations (§3.1):

- **C7 (local)**: For a single triangle Δ, tests whether the holonomy of Δ is compatible with the transport environments of the *two* tetrahedra sharing Δ as a common face. This is a local gluing condition that can be evaluated from data in the immediate neighborhood of Δ.
- **$\beta_2 = 4$ relations (global)**: Emergent constraints among the holonomies of all 10 triangular faces of a 4-simplex σ^4, arising from the topology of the full 2-skeleton. These relate faces that may not share any edges, and reflect the global 2-cycle structure of σ^4.

C7 is a necessary condition for the global relations to hold (inconsistent local gluing would obstruct global closure), but it is not sufficient: the global Bianchi relations carry additional topological content beyond pairwise co-face compatibility. There is no double-counting.

References

1. T. Regge, "General relativity without coordinates," *Nuovo Cimento* **19**, 558–571 (1961).
2. S. S. Chern, "A simple intrinsic proof of the Gauss-Bonnet formula for closed Riemannian manifolds," *Annals of Mathematics* **45**, 747–752 (1944).
3. D. Husemoller, *Fibre Bundles*, 3rd ed. (Springer, 1994).
4. M. Nakahara, *Geometry, Topology and Physics*, 2nd ed. (CRC Press, 2003).
5. M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory* (Westview Press, 1995).
6. J. Ambjørn, J. Jurkiewicz, and R. Loll, "Dynamical triangulations, a gateway to quantum gravity?" *J. Math. Phys.* **36**, 6299–6339 (1995).
7. M. Creutz, *Quarks, Gluons and Lattices* (Cambridge University Press, 1983).
8. D. Stauffer and A. Aharony, *Introduction to Percolation Theory*, 2nd ed. (Taylor & Francis, 1994).
9. J. R. Munkres, *Elements of Algebraic Topology* (Addison-Wesley, 1984).