A Geometric Closure Condition Linking Fundamental Constants

Keith Taylor VERSF Theoretical Physics Program

For the General Reader

Physics depends on a handful of fundamental constants — numbers like the speed of light, the strength of gravity, and the fine structure constant that governs how light interacts with matter. For over a century, these constants have been treated as separate, unexplained inputs: we measure them, plug them into our equations, and move on. Nobody knows why they have the values they do, or whether those values are connected to each other.

This paper argues that they are connected — tightly. The central idea is that the universe must be geometrically self-consistent. Imagine tiling the entire observable universe with the smallest possible units of area (each roughly 10^{-70} square metres, known as the Planck area). There are approximately 10^{122} such tiles. Each tile contributes a tiny amount of geometric "twist" — a phase rotation governed by the fine structure constant α. For the universe to close properly — for all those tiny local twists to add up to a coherent whole — the constants cannot take arbitrary values. They are locked together by a single equation.

The paper derives that equation, shows what it predicts numerically, and identifies how it can be tested. If correct, it means the fundamental constants are not free parameters. They are consequences of the universe being geometrically possible at all.

A note on "geometric closure": when we say the universe is "closed," we do not necessarily mean it curves back on itself like the surface of a sphere (though that is one possibility). We mean something more general — that the total geometry is internally consistent, that local structure and global structure are compatible, and that the mathematical description of the universe does not leave loose ends. Think of it like a jigsaw puzzle: every piece must fit with its neighbours, and the completed puzzle must form a coherent picture with no gaps and no overlaps. The closure condition in this paper is the mathematical statement of that requirement, applied to the geometry of spacetime itself.

Another way to think about it — perhaps more intuitive for a modern audience — is as a software programme. The closure condition is the compiler check: it tests whether a given set of constants produces a valid, executable universe. If the check fails, the programme does not run — there is no universe. The constants themselves are like dependencies in the code: they must all resolve against each other, and you cannot swap one out without breaking the build. The winding number χ — the large integer the paper predicts — functions as a checksum: a single integrity test that validates the whole structure. And the timeless geometric condition itself is the source code — static, complete, and non-temporal. What we experience as time, change, and physical evolution is the runtime: what it looks like from inside the execution.

So What?

Physics currently treats its fundamental constants — the speed of light, the strength of gravity, the fine structure constant, the cosmological constant — as independent, unexplained inputs. They are measured with extraordinary precision but fed into the equations by hand, with no account of whether they are mutually compatible. Worse, when we try to calculate some of them from first principles — particularly the cosmological constant — the answers are wrong by 120 orders of magnitude. This is not a small discrepancy. It is the worst prediction in the history of science.

The standard response has been to treat these constants as arbitrary: perhaps they vary across a vast multiverse, and we observe these particular values because they happen to permit our existence. This is the anthropic argument. It may be correct, but it explains nothing. It replaces "why these values?" with "why not?" and calls that progress.

The broader VERSF programme has developed independent structural derivations and constraints for α, Λ, c, and G individually. But deriving each constant separately is not enough. A set of individually valid constants could still be globally inconsistent — they could fail to fit together within a single coherent geometry, in the same way that individually well-made jigsaw pieces might not belong to the same puzzle.

This paper addresses that gap. It establishes a geometric admissibility condition: a single equation that the constants must jointly satisfy if they are to coexist within a closed, self-consistent spacetime. The equation produces a specific, falsifiable number — a topological winding number of approximately 10^{120} that must be an exact integer. That number can be checked against improving cosmological measurements. If it converges toward an integer, geometric closure is real and the constants are not merely individually constrained but collectively locked. If it does not, the framework is wrong — even if the individual derivations remain internally valid.

The stakes are straightforward. If the admissibility condition holds, physics is not modular. Gravity, electromagnetism, quantum mechanics, and cosmology are not separate sectors that happen to coexist in the same universe — they are structurally entangled. You cannot alter one constant without breaking all of them. The universe is not assembled from independent parts; it is a single geometric object whose local and global properties must be mutually consistent. That is a fundamentally different picture from the one modern physics operates with, where each force and each constant belongs to its own theoretical silo and the question of their mutual compatibility is never formally asked.

The deeper consequence is that if mutual consistency fully determines the constants, the universe is not contingent. It is not one configuration selected from a landscape of possibilities. It is the only configuration that closes. The values of the constants are not chosen, tuned, or anthropically filtered — they are geometrically necessary. This removes the need for a multiverse, for selection effects, and for any external principle that picks "this universe" over other possible ones. If closure plus the independent derivations fully fix the constants, there may be no other geometrically admissible configurations within this framework.

This paper establishes mutual constraint (admissibility), not uniqueness; demonstrating uniqueness would require deriving α and Λ within the same closure system with no remaining discrete freedom. What follows here is the admissibility condition and its consequences.

In plain terms: the universe is not a lucky accident. It is not one roll of the dice among countless others that happened to come up right for stars, chemistry, and life. It is the only way the geometry works. Everything that exists — every atom, every force, every scale from the subatomic to the cosmological — is locked in place by a single requirement: that the whole thing fits together. If this is correct, the question "why is there something rather than nothing?" gets a surprisingly concrete answer: because nothing else was geometrically possible.

And that includes you. If the constants could not have been otherwise, then stars were inevitable, chemistry was inevitable, and the conditions for life were not a coincidence to be explained away — they were built into the only geometry that works. If closure strongly constrains admissible constants, then life-permitting structure may be far less contingent than standard anthropic framing suggests.

A note on what this does and does not imply. The picture described here — a universe where nothing is arbitrary, where every part fits every other part, where the whole structure is determined by a single self-consistency requirement — has every characteristic we normally associate with design: coherence, necessity, parsimony, and the absence of free parameters. It is worth being honest about that resemblance. But the mechanism is not an external designer. It is internal geometric necessity. The universe does not look designed *by* something; it looks the way it does because no other configuration was geometrically possible. The coherence is self-imposed. Whether one finds that more or less remarkable than an external design is a philosophical question this paper does not attempt to answer — but it would be dishonest not to acknowledge that the question arises.

Once that question is asked — and answered with a specific equation — the consequences follow. The cosmological constant problem — the 120-order-of-magnitude embarrassment — is not a problem at all. It is a category error: the result of treating a geometric boundary condition as if it were a vacuum energy. On this interpretation, Λ does not correspond to a substance-like dark energy component. The constants are not arbitrary. No multiverse is required by this framework to explain the constants' mutual consistency. And the familiar cosmological hierarchy — the vast ratio between the smallest and largest scales in physics — is revealed not as a fine-tuning coincidence but as a topological invariant: the winding number of a geometrically closed universe.

Table of Contents

Abstract

We present a geometric consistency relation linking the fine structure constant (α), the gravitational constant (G), the cosmological constant (Λ), the speed of light (c), and the reduced Planck constant (ℏ). These constants are shown to be jointly constrained by a single, timeless closure condition: the requirement that local U(1) gauge holonomy and global geometric closure are mutually compatible across the full tiling of cosmological-scale spacetime by Planck-area cells. Time is not treated as a fundamental dimension. All relations are formulated in purely geometric and topological terms, with apparent temporal evolution understood as an emergent ordering of irreversible geometric commitments rather than as a primitive coordinate.

The closure condition identifies the topological winding number χ — a first Chern number — as the bridge between local gauge structure and global geometry. The form of the equation and the integrality of χ are derived from the Chern-Weil integrality theorem applied to the geometric U(1) bundle, with homogeneity following from de Sitter maximal symmetry. Under the leading-order holonomy function $\theta(\alpha) = 2\pi\alpha$ — derived from gauge invariance, analyticity, and the absence of additional cell-scale structure — the framework yields a specific numerical

prediction: $\chi \approx 9.52 \times 10^{120}$, an integer whose exact value is constrained by improving cosmological constraints on Λ and the associated horizon-scale inference. The framework is falsifiable, predicts correlated variation of constants, and connects directly to de Sitter horizon entropy.

Scope and Claims

What is derived in this paper:

- The geometric closure condition $N_\Sigma \cdot \theta(\alpha) = 2\pi\chi$ from the Chern-Weil integrality theorem, via discretisation of the U(1) flux integral over the closure surface
- The integrality of χ as a mathematical theorem (not a physical assumption)
- The homogeneity of per-cell holonomy from de Sitter maximal symmetry
- The leading-order holonomy function $\theta(\alpha) = 2\pi\alpha$ from gauge invariance, analyticity, and geometric parsimony
- The bridge hypothesis connecting geometric phase bookkeeping to the electromagnetic coupling
- The numerical prediction $\chi \approx 9.52 \times 10^{120}$ and the correlated variation constraint on constants
- The RG scheme-invariance test as a probe of the closure condition itself
- The Casimir-weighted non-abelian extension (exploratory, Appendix A)

What is assumed:

- The existence of a geometric U(1) bundle over the closure surface (VERSF substrate postulate)
- The coupling identification $\theta = 2\pi\alpha$ (argued in Section 3, not derived from first principles)
- The geometric identification $N = N_\Sigma = 4\pi N_\Lambda$ (Planck-area cells on the de Sitter horizon sphere)

What is not claimed:

- Derivation of the numerical values of α, G, Λ, c, or $\hbar$ (these exist elsewhere in the VERSF programme)
- A non-abelian topological theorem (the extension is exploratory)
- An alternative cosmological model (this is a consistency test, not a replacement for ΛCDM phenomenology)
- A complete microscopic derivation of the geometric U(1) bundle (a roadmap is provided)

---

1. Dimensionless Structure and the Role of *ℏ*

Any fundamental relationship between constants must be expressible in dimensionless form. Of the constants considered, $\alpha \approx 1/137.036$ is already dimensionless. The remaining constants G, Λ, and c cannot form a dimensionless combination among themselves: direct dimensional analysis in SI units shows that no product of powers of G ($L^3 M^{-1} T^{-2}$), Λ (L^{-2}), and c ($L T^{-1}$) eliminates both mass and length dimensions simultaneously. The minimal additional constant required is $\hbar$ ($L^2 M T^{-1}$), which enters through the Planck length:

$$\ell_P^2 = \hbar G / c^3$$

The combination $\Lambda \ell_P^2$ is dimensionless, and the quantity $N_\Lambda = 3/(\Lambda \ell_P^2)$ is therefore a pure number. The closure condition presented here constrains five constants — α, G, Λ, c, and $\hbar$ — not four.

---

2. Cosmological Closure Scale

The cosmological constant defines a global curvature scale:

$$r_\Lambda = \sqrt{3 / \Lambda}$$

This scale characterises global geometric closure — the scale at which the universe's geometry closes upon itself — rather than evolution in time. Combining this with the Planck area, we define a purely geometric, dimensionless quantity:

$$N_\Lambda = 3 / (\Lambda \ell_P^2)$$

N_Λ represents the ratio of the cosmological closure scale to the Planck area. The physical cell count on the de Sitter horizon 2-sphere (area $A_\Sigma = 4\pi r_\Lambda^2$) is:

$$N_\Sigma \equiv A_\Sigma / \ell_P^2 = 4\pi r_\Lambda^2 / \ell_P^2 = 12\pi / (\Lambda \ell_P^2) = 4\pi N_\Lambda$$

N_Σ is the total number of Planck-area cells tiling the closure surface — the global, timeless geometric invariant that enters the Chern-Weil derivation of the closure condition (Section 4).

2.1 Relation to de Sitter Horizon Entropy

The quantity N_Λ is directly related to the entropy of the de Sitter horizon. The Bekenstein-Hawking entropy of the cosmological horizon is:

$$S_{dS} = A / (4\ell_P^2) = 4\pi r_\Lambda^2 / (4\ell_P^2) = \pi \cdot 3 / (\Lambda \ell_P^2)$$

and therefore:

$N\_\Lambda = S\_dS / \pi$

and therefore:

$N\_\Sigma = 4\pi\, N\_\Lambda = 4\, S\_dS$

The closure condition therefore relates global horizon entropy to total gauge holonomy. The quantity $S\_dS \approx 3.3 \times 10^{122}$ is itself well established in the de Sitter entropy literature; the cell count $N\_\Sigma = 4S\_dS$ is not a new number. What is new is the interpretation: the proposal that this count of Planck-area cells participates in a topological closure condition linking it to gauge holonomy. The familiar cosmological hierarchy ($\sim 10^{120}$) appears not as a fine-tuning problem but as the entropy-scaled topological charge required for closure. This anchors the framework in established gravitational thermodynamics.

---

3. The Fine Structure Constant as U(1) Gauge Holonomy

The fine structure constant admits a direct geometric interpretation through gauge holonomy. In U(1) gauge theory, the physically meaningful object associated with a closed loop is the Wilson loop holonomy:

$W(C) = \exp(i\, (q/\hbar) \oint\_C A\_\mu\, dx\^\mu)$

where the physically meaningful quantity is the accumulated phase around a closed loop. The fundamental phase quantum is 2π; all gauge-invariant holonomies are defined modulo this value.

The fine structure constant,

$\alpha = e^2 / (4\pi\varepsilon_0\, \hbar c)$

is precisely the dimensionless quantity that sets the scale of phase accumulation per elementary electromagnetic interaction. It is therefore natural to interpret α as the fundamental phase strength — the holonomy per elementary interaction loop.

To eliminate the free-function ambiguity, we define θ operationally as the U(1) Wilson-loop phase on an elementary cell loop: $W(C) = \exp(i\theta)$ with $\theta = (q/\hbar) \oint A \cdot dx$. For weak coupling, gauge invariance and analyticity at $\alpha = 0$ (no interaction implies no phase) require:

$\theta(\alpha) = 2\pi(c_1\alpha + c_2\alpha^2 + c_3\alpha^3 + \cdots)$

The expansion is in powers of α because α is the unique dimensionless coupling governing U(1) phase strength, and the factor of 2π ensures θ is measured in units of the fundamental phase quantum. Under the principle that no additional dimensionless structure exists at the cell scale, the leading coefficient c_1 cannot be a complicated number — any value other than unity would constitute an unexplained new parameter. This fixes $c_1 = 1$, yielding:

$$\theta(\alpha) = 2\pi\alpha + O(\alpha^2)$$

In the present work we adopt the leading-order form $\theta(\alpha) = 2\pi\alpha$. This is not an arbitrary ansatz: linearity at leading order is forced by gauge invariance, analyticity, and the absence of additional cell-scale structure. The coefficient is fixed by minimality. Any deviation from linearity corresponds to higher-order geometric structure at the Planck scale and would shift the inferred integer χ, providing an explicit falsifiability channel: if the framework's predictions fail at leading order, the first diagnostic is whether $O(\alpha^2)$ corrections are required.

Normalisation and What Is (Not) Derived

The Chern-Weil and discretisation steps (Section 4) derive the form $N_\Sigma \cdot \theta = 2\pi\chi$ and the integrality of χ, but they do not fix the numerical mapping between the empirical U(1) coupling and the geometric per-cell holonomy scale. We therefore treat:

$$\theta(\alpha) = 2\pi\, c_1\, \alpha + O(\alpha^2)$$

as the most general leading-order relation. The choice $c_1 = 1$ is a normalisation convention consistent with minimality, not a theorem. The framework becomes more predictive to the extent that c_1 can be fixed by independent structure (e.g., representation or transport normalisation on the substrate) or constrained by scheme-invariance tests (Section 10.1). All numerical predictions in this paper assume $c_1 = 1$; departures from this value rescale χ linearly and are therefore directly testable.

Within the broader VERSF programme, this identification is further supported by the independent derivation of α as a dimensionless closure ratio controlling U(1) phase normalisation. If α is the per-loop closure fraction for U(1) phase — as derived in that context — then the per-cell phase increment entering the global winding closure is exactly $2\pi\alpha$. The mapping $\theta(\alpha) = 2\pi\alpha$ is therefore not an extra postulate but the translation of an independently derived interpretation of α into holonomy units.

3.1 Scale Dependence of α

The fine structure constant is known to run with energy scale under renormalisation. In the present framework, α is interpreted as the effective infrared U(1) coupling relevant to global closure — the renormalisation-group invariant value governing large-scale holonomy accumulation. Operationally, α is taken to be the low-energy (Thomson-limit) electromagnetic coupling relevant to macroscopic phase accumulation. Because the closure condition is formulated as a macroscopic global consistency constraint on the effective long-distance geometry, the low-energy α is the appropriate normalisation; any UV completion must reproduce this IR value while satisfying the closure constraint.

Alternatively, one may ask whether the closure condition is scheme-invariant: the topological class $\chi = (1/2\pi)\int_\Sigma F$ is a geometric invariant independent of renormalisation scale. What depends on μ is the mapping $\theta(\alpha(\mu))$. Scheme-invariance therefore requires that the inferred topological class be unchanged when α is evaluated at different scales — a condition that constrains the allowed $O(\alpha^2)$ terms in $\theta(\alpha)$. This is explored further in Section 10.1.

3.2 Bridge Hypothesis: Phase Bookkeeping on the Coherence Substrate

The closure condition requires that α — the electromagnetic coupling constant — appears in a geometric context. This subsection makes explicit the physical bridge between the two domains.

Bridge Hypothesis (Phase Bookkeeping Connection). The U(1) connection appearing in the closure condition is not the macroscopic electromagnetic field. It is an effective phase-holonomy bookkeeping connection defined on the pre-geometric coherence substrate that governs relational transport at the Planck scale.

In any relational or simplicial pre-geometry, the fundamental observable associated with transport is parallelism: the comparison of relational states across neighbouring cells. Parallel transport necessarily introduces a gauge redundancy. The minimal such redundancy is U(1) phase.

The normalisation of this phase response must be set by the unique dimensionless coupling controlling U(1) phase accumulation in nature, namely the fine structure constant α. The appearance of α in the geometric closure condition therefore reflects the empirical normalisation of phase holonomy inherited by the electromagnetic sector from the underlying geometric substrate.

In this sense, the closure condition does not assume that electromagnetic fields live on Planck cells. It assumes that the same U(1) phase normalisation observed in electromagnetism originates in, and constrains, the deeper geometric phase bookkeeping structure.

Minimal two-parameter bridge (geometric U(1) vs electromagnetic U(1)). A priori, the geometric U(1) connection could carry its own dimensionless normalisation α_g, distinct from the low-energy electromagnetic coupling α_{EM}. The most general minimal relation is:

$$\alpha_{EM} = \kappa \alpha_g$$

where κ is an embedding/renormalisation factor encoding how the emergent electromagnetic sector inherits the substrate phase normalisation. The closure condition constrains α_g through $\theta(\alpha_g)$, and therefore constrains κ indirectly. The identification $\alpha_g = \alpha_{EM}$ corresponds to $\kappa = 1$. This is the simplest bridge hypothesis, but the framework does not require it: allowing $\kappa \neq 1$ preserves the closure form while rescaling the inferred χ. In this way the bridge is not a mere assertion; it is a quantitatively constrainable embedding whose value can in principle be fixed by independent structure or by consistency with the inferred topological class.

---

4. Geometric Closure Condition

4.1 Derivation from Chern-Weil Integrality

The mathematical foundation is the Chern-Weil integrality theorem for U(1) principal bundles. For any U(1) connection on a principal bundle over a closed, oriented 2-manifold Σ:

$$c_1 = (1/2\pi) \int_\Sigma F \in \mathbb{Z}$$

This is a theorem — it follows from the classification of principal U(1) bundles over closed surfaces by their first Chern class. The total U(1) flux through any closed surface is necessarily quantised in units of 2π. This is the same mathematical content as Dirac's magnetic monopole quantisation condition, and it requires no physical assumptions beyond the existence of the bundle and the closedness of the surface.

Discretisation. Decompose Σ into N cells $\{C_i\}$. By additivity of the integral:

$$\int_\Sigma F = \Sigma_i \int_{C_i} F = \Sigma_i \varphi_i$$

where φ_i is the flux through cell i. By Stokes' theorem, φ_i equals the holonomy around the boundary of cell i: $\varphi_i = \oint_{\partial C_i} A = \theta_i$. Therefore:

$$\Sigma_i \theta_i = 2\pi c_1$$

This is exact — no approximation is involved in the discretisation step. The total holonomy over all cells equals 2π times an integer, regardless of how the cells are shaped or sized.

Homogeneity. If every cell carries the same holonomy $\theta_i = \theta$, the sum reduces to:

$$N \cdot \theta = 2\pi\chi$$

where $\chi = c_1$. The homogeneity assumption is not arbitrary — it follows from the symmetry of the background geometry. De Sitter space has maximal isometry group SO(4,1). Any spatial 2-sphere cross-section inherits an SO(3) rotational symmetry. If the geometric U(1) bundle respects this symmetry — i.e., if the geometric substrate does not spontaneously break spatial isotropy — then the connection is SO(3)-invariant on Σ, which forces the flux per unit area to be constant across the surface. Combined with a uniform cell decomposition, this gives $\theta_i = \theta$ for all i. Homogeneity is therefore not an additional assumption but a consequence of the cosmological principle applied to the geometric substrate.

Coupling identification. The per-cell holonomy is $\theta(\alpha) = 2\pi\alpha$, as argued in Section 3 from gauge invariance, analyticity, and the absence of additional cell-scale structure. Substituting:

$$N \cdot 2\pi\alpha = 2\pi\chi$$

The factor of 2π cancels:

$$N \cdot \alpha = \chi \in \mathbb{Z}$$

4.2 Geometric Identification

The natural closed 2-surface in de Sitter space is the cosmological horizon 2-sphere, with area $A_\Sigma = 4\pi r_\Lambda^2 = 12\pi/\Lambda$. Tiled by Planck-area cells (area ℓ_P^2 each), the cell count is:

$$N_\Sigma = A_\Sigma / \ell_P^2 = 12\pi / (\Lambda \ell_P^2) = 4\pi N_\Lambda$$

as defined in Section 2. The Chern-Weil derivation therefore gives the closure condition:

$$N_\Sigma \cdot \alpha = \chi \in \mathbb{Z}$$

where $\chi \equiv c_1$ is the first Chern number. This is the exact integrality condition with no unresolved normalisation factors — the cell count matches the physical area of the closure surface, and χ is the standard topological invariant.

4.3 What Is Derived vs. What Is Assumed

The Chern-Weil derivation substantially reduces the postulate content of the closure condition. Here is the accounting:

Now derived (previously postulated):

- The integrality of χ. This is a theorem — it follows from the topology of U(1) bundles over closed surfaces via Chern-Weil. It requires no physical assumption beyond the existence of the bundle and the closedness of Σ.
- The form of the equation, $N \cdot \theta = 2\pi\chi$. This follows from discretising the Chern-Weil integral. It is exact, not approximate.
- The homogeneity of per-cell holonomy. This follows from the maximal symmetry of de Sitter space, provided the geometric U(1) bundle respects spatial isotropy.

Still assumed (but now explicitly isolated):

- The existence of a geometric U(1) bundle over the closure surface. This is the VERSF substrate postulate — that a U(1) phase structure exists at the Planck scale, from which the electromagnetic interaction inherits its coupling (see Section 3.2).
- The coupling identification $\theta = 2\pi\alpha$. Already argued in Section 3; the derivation does not change its status.
- The geometric identification $N = N_\Sigma = 4\pi N_\Lambda$ (the physical Planck-area cell count on the de Sitter horizon sphere).

Given the existence of a geometric U(1) bundle over the closure surface, Chern-Weil integrality and exact discretisation imply the form $N_\Sigma \cdot \theta(\alpha) = 2\pi\chi$. The remaining physical content is the substrate-bundle postulate and the identification of θ with the α-normalised holonomy per cell.

Physical Content of the Substrate Postulate. The central physical assumption of the framework is the existence of a U(1) phase bookkeeping structure on the coherence substrate.

This is not an exotic addition: any relational theory of transport requires a notion of parallelism, and parallelism generically introduces a gauge redundancy. The minimal such redundancy is U(1). The present work explores the consequences of this assumption; its ultimate validity must be decided by whether the resulting closure condition survives empirical and theoretical consistency tests.

4.4 Minimality and Uniqueness

Proposition (Minimality of the Closure Form). Among dimensionless global constraints linking (i) the total number of Planck-area cells N_Σ, (ii) a local U(1) holonomy normalisation $\theta(\alpha)$, and (iii) a topological quantisation condition, the form $N_\Sigma \cdot \theta(\alpha) = 2\pi\chi$ is essentially unique under the following minimal requirements:

Dimensionlessness: the constraint involves no dimensionful scales beyond those already appearing in N_Σ.

Topological protection: the right-hand side is quantised and invariant under smooth deformations.

Coarse-graining stability: the condition is preserved under aggregation of microscopic cells.

No additional structure: no new dimensionless parameters are introduced.

Any alternative functional form of comparable simplicity either introduces extra dimensionless coefficients, breaks topological protection, or fails to be stable under coarse-graining. In this sense, the closure condition is not an arbitrary choice but the simplest admissibility criterion consistent with these requirements.

4.5 Properties of χ and the Nature of the U(1) Bundle

The U(1) bundle in question is not identified with the physical photon field. A literal electromagnetic Chern number $\chi \sim 10^{120}$ would require an enormous net magnetic flux through the cosmological horizon — equivalently, an enormous net magnetic charge inside — which is not observed. Here χ is computed for an effective U(1) connection describing phase-holonomy bookkeeping on the pre-geometric coherence substrate — the same structure from which the electromagnetic U(1) emerges in the VERSF programme — not for a macroscopic Maxwell field configuration. The fact that α appears in both the electromagnetic and geometric contexts is not a coincidence to be explained but a consequence of their common origin: the geometric U(1) structure of the substrate is the structure from which the electromagnetic interaction inherits its coupling. Operationally, this effective U(1) connection represents the phase-holonomy bookkeeping associated with relational transport on the coherence substrate; α enters because it is the empirically observed normalisation of U(1) phase coupling inherited by the electromagnetic sector.

χ counts the total number of U(1) holonomy quanta required for global closure. It is an integer by the Chern-Weil theorem (topological quantisation of the first Chern class) and is not identified with the Euler characteristic of a simple manifold. The value $\chi \sim 10^{120}$ is large by comparison with Chern numbers encountered in condensed matter or string compactifications, where the base manifolds have areas of order unity in natural units. Here the base surface has area $\sim 10^{122}$ Planck areas; a large total winding number is the expected consequence of integrating a small holonomy per cell over an enormous surface, just as a large total curvature results from integrating small local curvature over a large manifold.

This equation is a static consistency condition. It expresses the requirement that local interaction geometry and global spacetime geometry close coherently, without reference to temporal evolution.

Under the leading-order form $\theta(\alpha) = 2\pi\alpha$ derived in Section 3, the closure condition yields the definite prediction:

$N_\Sigma \cdot \alpha = \chi \in \mathbb{Z}$

This removes all free functions from the framework and yields a specific numerical value for the topological charge given measured constants.

---

5. Relationship Between Constants

Substituting the definition of N_Σ into the closure condition yields:

$(12\pi / (\Lambda \ell_P^2)) \cdot \theta(\alpha) = 2\pi\chi$

or equivalently,

$\Lambda = (6 / \chi) \cdot (\theta(\alpha) / \ell_P^2)$

Since $\ell_P^2 = \hbar G / c^3$, this relation explicitly links:

- **α** — local U(1) gauge holonomy strength,
- **G** — minimal geometric resolution (Planck area),
- **Λ** — global geometric closure scale,
- **c** — causal adjacency constant governing the structure of permitted causal connections,
- ***ħ*** — minimum quantum of geometric commitment.

The speed of light appears here not as a temporal velocity but as a structural constant governing causal connectivity in the geometric substrate. Similarly, *ħ* enters not as a dynamical action quantum but as the minimum scale of irreversible geometric commitment.

---

6. Numerical Consistency Check

The closure condition yields a concrete numerical prediction. Using current measured values:

Constant	Value
Λ	1.1056×10^{-52} m^{-2}
ℓ_P	1.616255×10^{-35} m
ℓ_P^2	2.611×10^{-70} m^2
α	1/137.036

The number of Planck cells tiling the closure surface is:

$$N_\Sigma = 12\pi / (\Lambda\, \ell_P^2) \approx 1.31 \times 10^{123}$$

Under the leading-order form $\theta(\alpha) = 2\pi\alpha$, the implied topological winding number is:

$$\chi = N_\Sigma \cdot \alpha \approx 9.52 \times 10^{120}$$

The magnitude of χ is not problematic — it reflects the cosmological hierarchy expressed as a topological count. The same order of magnitude ($\sim 10^{120}$) appears independently in the Bekenstein-Hawking entropy of the cosmological horizon (as shown in Section 2.1) and in the ratio of the Hubble scale to the Planck scale squared. The closure condition reinterprets this hierarchy: it is the total winding number required for the universe's local holonomy to close globally.

6.1 Uncertainty and Integrality Test

The closure condition does not require exact integrality of χ at finite experimental precision, but convergence toward integrality within observational uncertainty.

From

$$\chi = N_\Sigma \cdot \alpha = 12\pi\alpha / (\Lambda\, \ell_P^2)$$

the fractional uncertainty is dominated by the uncertainty in Λ:

$$\sigma_\chi / \chi \approx \sigma_\Lambda / \Lambda$$

Current measurements constrain Λ at the few-percent level, implying an uncertainty of order 10^{118} in χ — a window that admits an enormous number of candidate integers. Even heroic improvements in Λ only shrink the integer window linearly; therefore χ-integrality is primarily a structural requirement of the framework, while correlated-variation constraints (Section 10) are the practical observational handle.

Exact integrality is therefore an in-principle constraint, not a near-term experimental test. The correlated variation of constants (Section 10) is the framework's primary near-term falsifiable prediction.

---

7. The Status of Time

Time does not appear as a fundamental dimension in this framework. At the foundational level, geometry is static and the closure condition is timeless. What is conventionally described as time emerges only when geometric commitments become irreversible — producing an ordered, monotonically increasing record of distinguishable states.

Accordingly, cosmological expansion and interaction dynamics are interpreted as emergent, coarse-grained descriptions of underlying geometric commitment processes rather than as primitive features of reality. The cosmological constant, in particular, is not an energy density driving accelerated expansion; it is a geometric boundary condition specifying the closure scale. The appearance of acceleration is an artefact of projecting a static geometric constraint into a time-dependent cosmological framework.

A note on observational access: the measured value of Λ used in this paper is extracted from dynamical cosmological observations — supernova luminosity distances, baryon acoustic oscillations, and CMB power spectra — all of which are interpreted within the standard ΛCDM framework, which treats time as fundamental. This creates an apparent tension: the paper uses outputs of a dynamical theory while reinterpreting Λ as a static geometric quantity. The resolution is that the numerical value of Λ is an empirical datum independent of its interpretation. ΛCDM provides the measurement; the closure condition provides a different reading of what is measured. The same number is extracted regardless of whether one interprets it as a vacuum energy density or as a geometric closure scale, because the observable consequences (luminosity distances, angular diameter distances) are identical at the level of the background geometry. What changes is the ontology, not the phenomenology.

Within the Void Energy-Regulated Space Framework (VERSF), the closure condition presented here can be understood as arising from the requirement that entropy gradients on a zero-entropy substrate produce a globally consistent geometric structure. The zero-entropy substrate admits no temporal coordinate; the closure condition is the mathematical expression of global consistency on that substrate. Temporal emergence, and with it all apparent dynamics, follows from the irreversibility of local entropy production within the closed geometry.

---

8. Relation to Prior Work

The scale $1/(\Lambda\ \ell_P^2)$ and its relation to horizon entropy have been extensively studied. Bekenstein (1973) and Hawking (1975) established that black hole entropy scales as $A/(4\ell_P^2)$;

the extension to de Sitter horizons yields $S\_dS \sim 1/(\Lambda\, \ell\_P^2)$, the same quantity that appears as N_Λ in this paper. Banks and Fischler (2000, 2001) argued that the Hilbert space of de Sitter space is finite-dimensional, with dimension set by the horizon entropy — a conclusion structurally adjacent to the present framework's identification of N_Λ as the total count of geometric degrees of freedom. Bousso's covariant entropy bound (1999) provides a rigorous holographic constraint linking areas to entropy across general spacetimes. The Cohen-Kaplan-Nelson bound (1999) relates UV and IR cutoffs in a way that produces $\Lambda \sim \ell\_P^{-2} \times L^{-2}$, where L is a cosmological scale — remarkably similar in structure to the relation $N\_\Lambda = 3/(\Lambda\, \ell\_P^2)$ used here. Separately, large-number hypotheses dating to Dirac (1937, 1938) noted numerical coincidences among cosmological and microphysical scales, and gauge theories employ holonomy, winding numbers, and topological quantisation to enforce global consistency.

The quantity $N\_\Lambda \approx S\_dS/\pi \approx 10^{122}$ is therefore not new — it is well established in the de Sitter entropy literature. What is new is the closure interpretation: the proposal that N_Λ is not merely an entropy count but participates in a global topological constraint linking gauge holonomy to cosmological geometry.

The present work differs from these predecessors in that it imposes a single closure condition connecting gauge structure to cosmological geometry. Unlike large-number hypotheses or numerological relations, the closure condition is dimensionless, topological, and falsifiable. It does not attempt to compute constants from microphysics but constrains them through global admissibility. The identification of χ as a first Chern number, rather than an empirical ratio, gives the large number a standard topological interpretation.

The present work does not attempt to derive the numerical values of the fundamental constants. Independent derivations or structural constraints for α, Λ, c, and G exist within the broader VERSF programme. The role of this paper is different: it establishes a global geometric admissibility condition that must be satisfied if those independently constrained constants are to coexist within a single closed, self-consistent geometry. Failure of this condition would falsify the framework even if the individual derivations remained internally consistent.

The present paper should therefore be read neither as a derivation of constants nor as an alternative cosmological model, but as a global consistency test that any viable theory of fundamental constants must satisfy.

---

9. Implications and Extensions

If the closure condition is correct, several consequences follow.

The constants α, G, Λ, c, and $\hbar$ are not independent parameters requiring separate explanation. They are jointly fixed by a single geometric requirement. The number of truly free parameters in fundamental physics is reduced, and the question "why these values?" is replaced by "what does geometric closure require?"

The cosmological constant problem — the ~120 order-of-magnitude discrepancy between the quantum field theory vacuum energy prediction and the observed Λ — is reframed as a category error. Λ is not a vacuum energy to be calculated from particle physics; it is a geometric closure scale. The apparent fine-tuning is an artefact of treating a geometric boundary condition as a dynamical energy density.

The framework naturally suggests extension to non-abelian gauge holonomy. If α enters the closure condition as the U(1) phase coupling, then the strong and weak coupling constants (α_s and α_w) may enter through SU(3) and SU(2) holonomy contributions to a generalised closure condition. A brief exploratory non-abelian extension — including the derivation of Casimir weights $k_2 = 3/4$ (SU(2)) and $k_3 = 4/3$ (SU(3)) from small-loop Wilson holonomy expansion — is provided in Appendix A. This material is not used in the main argument; the core closure condition of this paper involves only α and the U(1) sector.

---

10. Falsifiability

10.1 Scheme-Invariance of the Closure Class

The closure condition is formulated using the low-energy (Thomson-limit) normalisation of α relevant to macroscopic phase accumulation. RG running of $\alpha(\mu)$ is not a violation of closure; it reflects the scale dependence of the effective coupling in a particular renormalisation scheme. A UV completion must reproduce the IR value while remaining compatible with closure.

However, the topological class $\chi = (1/2\pi)\int_\Sigma F$ is a geometric invariant that does not depend on renormalisation conventions. What depends on the scale μ is the mapping $\theta(\alpha(\mu))$. Scheme-invariance of the closure condition therefore requires that the inferred topological class be unchanged when α is evaluated at different renormalisation scales:

$$\chi_{inferred}(\mu_1) = \chi_{inferred}(\mu_2)$$

This constrains the allowed $O(\alpha^2)$ terms in $\theta(\alpha)$: if $\theta = 2\pi(\alpha + c_2\alpha^2 + \cdots)$, then changing μ shifts $\alpha \to \alpha(\mu)$, and the inferred χ must remain the same integer. At leading order ($\theta = 2\pi\alpha$), the constraint is automatically satisfied if one commits to a single reference scale (the Thomson limit). Beyond leading order, scheme-invariance provides a nontrivial test: do there exist $O(\alpha^2)$ corrections to θ that maintain a consistent topological class across scales?

This tests the geometry-phase bridge itself: if no choice of higher-order corrections yields scheme-invariant closure compatible with the known running of α, the closure hypothesis fails even if its low-energy consequences appear consistent.

10.2 Observational Predictions

The framework makes the following testable predictions, ordered by near-term observational accessibility:

Correlated variation of constants (primary near-term test). Holding χ fixed implies a correlated variation constraint:

$$\delta\Lambda/\Lambda = \delta\alpha/\alpha + 3(\delta c/c) - \delta\hbar/\hbar - \delta G/G$$

This assumes χ is topologically protected — as a first Chern number, it cannot vary continuously but can in principle jump between integers under sufficiently extreme topological transitions (analogous to phase transitions in condensed matter). If such jumps are possible, the constraint changes character: small continuous variations of constants must still satisfy the equation above, but discrete jumps in χ would permit discontinuous shifts in the constant relationships. The possibility of topological phase transitions in the cosmological geometry is not addressed here but would have observable signatures distinct from smooth variation.

Any observed smooth variation of one constant must be accompanied by compensating variation in the others. Independent variation of α alone — as suggested by some quasar absorption spectra — would violate closure. This prediction is testable with existing and near-future astrophysical data without requiring extreme precision.

Distinguishability from alternative varying-constant frameworks. Many theories that allow varying constants (e.g., Bekenstein-type varying-α models or dilaton/runaway scalar scenarios) predict correlations among variations, but the structure differs: typically a single scalar field controls multiple couplings with model-dependent relative coefficients. In the present framework, by contrast, the constraint is anchored to a topological class χ, which admits two distinctive features:

Quantised-class protection: small continuous variations must lie on a constrained manifold in constant-space; large changes may occur only via discrete class changes (topological transitions), which would produce non-smooth signatures.

Fixed coefficient structure: the relative coefficients (+1, +3, −1, −1) in the correlated variation equation are not fit parameters but follow from the dimensional structure of the closure relation once χ is held fixed.

Framework	Variation driver	Discrete transitions?	Coefficients predicted?
Bekenstein varying-α	Scalar field coupled to F^2	No	Model-dependent
Dilaton/runaway	Scalar field in string moduli	No	Moduli-dependent
Geometric closure (this work)	Topological class χ	Yes (class jumps)	Yes (from dimensional structure)

A detailed phenomenological comparison with specific scalar-field models is deferred; the key point is that the present framework predicts either strict constancy (if χ is rigid) or correlated variation punctuated by discrete transitions — a pattern not generic in smooth scalar-field models.

Geometric Λ vs substance-like dark energy. Within this framework Λ is interpreted as a geometric boundary condition, not as a dynamical fluid component. This does not by itself rule out additional light fields in nature. What it does imply is that late-time acceleration should be consistent with an exactly constant Λ ($w = -1$) except possibly for discrete transition-like departures if the topological class changes. Smooth quintessence-like evolution of $w(a)$ would therefore falsify the geometric closure interpretation. Detection of time-varying $w(a)$ consistent with smooth scalar-field dynamics would be direct evidence against this framework.

Λ is exactly constant. Because the closure condition is topological and static, it selects a constant Λ. A time-varying Λ would imply a changing N_Σ, and hence either a changing χ or compensating variation in α, G, or $\hbar$. Evidence for $w(a) \neq -1$ would falsify the framework unless correlated variation preserves integer χ. Independent dynamical dark energy is incompatible with geometric closure.

Integer winding number (structural constraint). The quantity N_Σ · α must be an exact integer. Current uncertainties on Λ permit a window of $\sim 10^{118}$ candidate integers, making this an in-principle rather than near-term test. Its primary role is structural: once Λ and α are constrained sufficiently, the framework requires N_Σ · α to lie within observational uncertainty of an integer. Persistent failure of that convergence would falsify the closure hypothesis.

10.3 Falsification Logic and Derivation Roadmap

On Falsification. As with other global consistency principles in physics, the closure condition is not tested by direct measurement of its quantised integer, but by the web of constraints it imposes. The framework is falsified if:

- Scheme-invariance of the closure class cannot be maintained (Section 10.1),
- correlated variation relations fail observationally,
- or the closure condition proves incompatible with independently derived values of α or Λ without introducing new parameters.

Roadmap to a Complete Derivation. The Chern-Weil derivation in Section 4.1 completes steps 2 and 3 below, conditional on the substrate postulate. Elevating the full framework to a derived theorem would require:

1. A microscopic model of the coherence substrate admitting a natural U(1) transport redundancy. **(Open — the central remaining gap.)**
2. A proof that the coarse-grained closure surface carries a quantised U(1) bundle class. **(Completed — Chern-Weil integrality theorem, Section 4.1.)**
3. A derivation showing that the total bundle class equals the sum of per-cell holonomies. **(Completed — exact discretisation via Stokes' theorem, Section 4.1.)**

4. A demonstration that alternative closure functionals either violate RG stability or introduce additional dimensionless structure. **(Partially addressed — Section 4.4 establishes minimality; full proof requires step 1.)**

The present work establishes the admissibility principle and its consequences; deriving it from a complete microscopic theory remains a goal for future work within the VERSF programme.

---

11. Conclusion

The constants α, G, Λ, c, and ℏ are best understood as projections of a single underlying geometric closure requirement. Their observed values reflect the unique configuration in which local U(1) holonomy, geometric resolution, causal structure, quantum commitment scale, and global closure are mutually compatible. The closure condition $N_\Sigma \cdot \theta(\alpha) = 2\pi\chi$ is derived from the Chern-Weil integrality theorem for U(1) bundles over closed surfaces, with homogeneity following from de Sitter symmetry. The holonomy function $\theta(\alpha) = 2\pi\alpha$ is derived at leading order from gauge invariance, analyticity, and the absence of additional cell-scale structure. The remaining physical assumptions — the existence of the geometric U(1) bundle and the coupling identification — are explicitly enumerated and individually falsifiable.

Under this framework, the closure condition yields a specific, falsifiable prediction: $\chi \approx 9.52 \times 10^{120}$, an integer constrained by improving measurements of Λ. The framework removes the need to treat fundamental constants as independent inputs, reframes the cosmological constant as a geometric boundary condition, dissolves the cosmological constant problem as a category error, and identifies the cosmological hierarchy ($\sim 10^{120}$) as a topological invariant rather than a fine-tuning coincidence.

The framework is falsifiable, connects to established gravitational thermodynamics through the de Sitter entropy relation, is extensible to non-abelian gauge structure, and is grounded in the timeless geometric consistency of the VERSF programme.

Appendix A: Exploratory Non-Abelian Extension

This appendix is an exploratory sensitivity check, not part of the main claim. The core closure condition of this paper involves only α and the U(1) sector.

To incorporate non-abelian sectors in a scalar closure functional, we define the effective per-cell holonomy contribution as the gauge-invariant size of small-loop Wilson holonomy in the matter representation. For a non-abelian group SU(N), the Wilson loop in representation R is:

$$W_R(C) = (1/d_R)\, \mathrm{Tr}_R\, [P \exp(ig \oint_C A_\mu^a T^a dx^\mu)]$$

For a small loop bounding area Σ, expanding to leading non-trivial order, the gauge-invariant deviation from identity scales as:

$$W_R(C) \approx 1 - (g^2 C_2(R) / 2) \Sigma^2 \langle F^a F^a \rangle + \cdots$$

where $C_2(R)$ is the quadratic Casimir. For the fundamental representations at low energies:

$$C_F(SU(2)) = 3/4, \quad C_F(SU(3)) = 4/3$$

The choice of fundamental representation is motivated by Standard Model matter content; adjoint Casimirs (2, 3) would yield different weights. A geometric derivation of the relevant representation remains open.

The generalised closure condition takes the form:

$$N_\Sigma \cdot [\alpha + (3/4)\alpha_w + (4/3)\alpha_s] = \chi$$

where $\alpha_w \equiv g_w^2/(4\pi)$ and $\alpha_s \equiv g_s^2/(4\pi)$. Couplings must be evaluated at a common renormalisation scale.

Mathematical caveat: The U(1) closure uses the phase directly ($\theta = 2\pi\alpha$), whereas the non-abelian Wilson loop gives deviation from identity scaling as $g^2C_2(R)$. The additive combination assumes linearity of the closure functional — a simplifying assumption, not a derived result. **Topological caveat:** U(1) winding uses the first Chern number; SU(2) and SU(3) topology naturally involves the second Chern class (instanton number). A full treatment incorporating appropriate Chern classes for each gauge sector is deferred to future work.