Why 1/137 Is Not Arbitrary: A Closure of the VERSF Derivation Chain

---

For the General Reader

There is a number that appears throughout physics — in the structure of atoms, in the way light interacts with matter, in the fine details of quantum electrodynamics. That number is approximately 1/137, and it is called the fine-structure constant. It sets the strength of the electromagnetic force. Without it, chemistry would be impossible; stars would not form; life as we know it could not exist.

What makes this number so strange is that nobody knows why it has the value it does. In our best current theory — the Standard Model of particle physics — the value 1/137 is simply inserted by hand, taken from experiment, with no deeper explanation offered. The physicist Richard Feynman called it "one of the greatest damn mysteries of physics."

This paper proposes an answer. But it does not do so alone. It is the culmination of a series of papers within the Void Energy–Regulated Space Framework (VERSF), each of which established one part of the foundation. Understanding where this paper sits in that sequence helps make sense of what it is doing.

Earlier papers in the programme established the basic ground rules. *Physical Admissibility* showed that any universe capable of containing facts at all must satisfy two conditions: it must be able to produce distinguishable states, and some of those distinctions must become permanent, irreversible records. *Ticks-Per-Bit* then showed that entropy — usually thought of as disorder or heat — is really a measure of how efficiently a system produces distinguishable changes. *Entropy Is the Field Beneath Spacetime* showed that this dynamical entropy is the conserved quantity underlying the large-scale behaviour of the universe.

On the quantum side, *The Double Square Rule* and *Born Rule as Entropic Unfolding* showed that the probability rules of quantum mechanics — the famous Born rule — follow from these same distinguishability constraints rather than needing to be assumed separately. *Complex Hilbert Space from Distinguishability Principles* showed that quantum theory must use complex numbers, not real ones, for the same structural reasons.

Closer to the present paper, *The Fold Interface Law* derived the existence of a specific kind of boundary — the *fold* — at which reversible possibilities collapse into irreversible facts. *The Void, The Fold, and the Derivation of Spacetime* showed how the geometry of that boundary gives rise to the structure of space, quantum correlations, and the symmetry groups of the Standard Model. And *Structural Sufficiency, Risk Concentration, and Representation Selection* worked out the precise combinatorial counting of how many independent ways that boundary can encode a stable distinction.

One earlier paper, *Vacuum Impedance, Resistance Quanta, and the Fine-Structure Constant*, established a different but complementary result: that $\alpha = Z_0/(2R\_K)$, meaning the fine-structure constant can be understood as the ratio of the vacuum's electromagnetic impedance to twice the quantum of resistance. This is an exact identity, but it explains the physical *meaning* of α rather than its numerical *value*.

The present paper closes the chain. It shows that all the structural constraints established in the preceding papers — the relational nature of information, the area law, the geometry of the fold — force the numerical value of α to be approximately 137, not as a lucky coincidence, but as the only count consistent with a universe that can support stable physical facts at all.

The core idea is this: once the conditions for a fact-supporting universe are imposed, they force the existence of the fold boundary, determine its geometry, and constrain the number of independent ways it can encode a distinction. That count — combining two structural effects, one combinatorial and one topological — yields $\alpha^{-1} \approx 137$. The number is not a free choice. It is the answer to a counting problem set by the requirements of physical reality itself.

The result is not claimed as an exact derivation — the full precision of the measured value involves further quantum field theory corrections that lie beyond the present scope. What is claimed is something more fundamental: that the number is structurally compelled, and that understanding why requires seeing the entire derivation chain from admissibility to coupling as a single connected argument.

The technical argument that follows is presented in full mathematical detail for specialist readers. The general reader may find the introduction and conclusion sufficient for orientation, with Section 4 (the fold interface) and Section 7 (the coupling formula) providing the conceptual heart of the derivation.

Table of Contents

Abstract

The fine-structure constant α is one of the most precisely measured and least explained parameters in physics. Within the Standard Model it appears as an empirical input rather than as a consequence of deeper principles. The Void Energy–Regulated Space Framework (VERSF) proposes an alternative in which physical law and physical constants emerge from constraints on distinguishability, irreversibility, and finite capacity.

Earlier work within the VERSF programme has approached α from two complementary directions. The first, developed in *Vacuum Impedance, Resistance Quanta, and the Fine-Structure Constant*, establishes the physical interpretation: $\alpha = Z_0/(2R\_K)$, identifying the fine-structure constant as the ratio of the vacuum's electromagnetic impedance to twice the von Klitzing resistance. This is an exact algebraic identity within standard electrodynamics, but it relocates rather than resolves the question of why α takes the value it does. The second, developed in *Structural Sufficiency, Risk Concentration, and Representation Selection in the One-Fold Framework* and companion closure work, derives the numerical value from hexagonal

closure geometry: the honeycomb packing theorem constrains the minimal binary closure structure to K = 7 independent constraints and N = 14 loop-consistency channels, yielding $\alpha^{-1} = 2^7 \cdot (15/14) \approx 137.14$. These two results together — the impedance interpretation and the combinatorial derivation — provide both the physical meaning and the numerical origin of α within VERSF. What they do not yet supply is a continuous structural derivation chain showing why those integers are forced by the same foundational admissibility conditions that govern realized information, entropy, and quantum structure.

The present paper closes that gap. We show that physically realized information is not local but must be supported by non-removable correlations across system boundaries, uniquely quantified — within the standard class of quantum information measures — by mutual information. Combined with locality and finite on-site capacity, this yields an area law for realized information. We then show that the same admissibility constraints force the existence of a minimal commitment interface — the *fold* — carrying a four-dimensional complex state space $\mathbb{C}^4$, derived under well-defined field-selection conditions. The local differential geometry of the fold, exhausted by its first and second fundamental forms, determines three independent geometric data channels of dimensions 1, 2, and 3. These channels constitute the continuous geometric substrate whose minimal admissible discrete realization, under finite capacity and global consistency, is precisely the closure structure with K = 7 and N = 14. The electromagnetic coupling therefore takes the structural form

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

yielding $\alpha^{-1} \approx 137$ at leading order.

The result is not a renormalization-complete numerical prediction. It is the missing architectural link: a continuous derivation chain from admissibility constraints through fold geometry to the discrete closure integers, establishing that the value of α is not an independent empirical input but a consequence of the conditions any universe must satisfy in order to support stable physical facts. The paper closes the VERSF derivation chain at the programme level. Uniqueness of the discrete realization is established jointly with the companion closure papers cited above.

---

1. Introduction

Dimensionless constants occupy a special place in physics. Unlike dimensional quantities, they cannot be changed by choice of units, and for that reason they appear to express something intrinsic about the structure of reality. Among them, the fine-structure constant α is one of the most important. It governs the strength of electromagnetic interaction, controls atomic structure, and appears throughout quantum electrodynamics and precision measurement. Yet within the Standard Model its value is not derived. It is measured and inserted.

The VERSF programme begins from the conviction that this state of affairs is not inevitable. Rather than taking spacetime, fields, and couplings as given, it asks a prior question: what conditions must hold for a universe to support physical facts at all? This shifts the starting point of theory away from dynamics and toward admissibility. A physical world must sustain distinguishable states, convert some distinctions into irreversible records, and do so under finite resources. These are not detailed model assumptions. They are candidate necessities for any universe in which empirical physics is possible.

A substantial body of work within the VERSF programme has developed the consequences of this starting point. These results divide into four groups.

On foundations and entropy: Physical Admissibility: A Constraint-Based Foundation for Physics establishes finite distinguishability and irreversible commitment as necessary conditions for any fact-supporting theory. *Ticks-Per-Bit: A Microphysical Foundation for Entropy* defines entropy as a dynamical measure of distinguishability production rather than a static microstate count. *Entropy Is the Field Beneath Spacetime* identifies entropy as the unique conserved current underlying effective physical dynamics.

On quantum structure: The Double Square Rule and *Born Rule as Entropic Unfolding* reconstruct quantum probability from distinguishability constraints. *Complex Hilbert Space from Distinguishability Principles* derives the necessity of complex state space from first principles.

On the commitment boundary: The Fold Interface Law formally derives the existence and minimal structure of the commitment boundary. *The Void, The Fold, and the Derivation of Spacetime, Quantum Correlations, and the Standard Model Gauge Group* develops the full geometric structure of the fold and its consequences for gauge symmetry.

On discrete closure and coupling: Structural Sufficiency, Risk Concentration, and Representation Selection in the One-Fold Framework derives the closure parameters K and N from admissibility conditions. *Vacuum Impedance, Resistance Quanta, and the Fine-Structure Constant* identifies $\alpha = Z_0/(2R\_K)$ as the impedance interpretation of the fine-structure constant.

Those results strongly suggest that physical law is shaped by deep informational structure. What they do not by themselves explain is how that structure determines physical constants. The information-theoretic and closure-theoretic sides of the argument are each internally coherent. What has been missing is a continuous derivational bridge between them: an explicit demonstration that the same geometric and combinatorial structure forced by admissibility constraints on continuous realized information is also the structure that discrete closure theory is realizing.

The purpose of this paper is to construct that bridge. The central claim is that the admissibility constraints governing realized information also force the existence and structure of a minimal commitment boundary — the fold — and that the local differential geometry of the fold is the continuous structure whose minimal admissible discrete realization is described by the closure parameters of companion work. Once this is established, the two halves of the derivation belong to a single route rather than two compatible but independent theories.

The paper proceeds in five stages: (1) realized information is identified with mutual information and shown to obey an area law; (2) the fold is derived as the minimal commitment interface carrying $\mathbb{C}^4$; (3) the local geometry of the fold is decomposed into three exhaustive differential-geometric channels; (4) discrete closure is identified as the minimal admissible realization of that geometry, motivating K and N; (5) K and N produce the structural coupling formula. The closure claimed is architectural: this paper establishes the continuous structural route from admissibility to coupling. Exact uniqueness of the discrete realization is established in the closure companion papers, whose results are integrated here.

---

2. Distinguishability, Entropy, and Realized Information

The starting point is the distinction between formal entropy and physically realized information. A system may possess a large formal state space while still failing to support physical records. If a distinction can be erased by local re-description, or disappears under coarse-graining, it does not constitute a stable physical fact.

Within VERSF, this is sharpened by the notion of commitment. A distinction becomes physically real only when it is irreversibly exported beyond local reversible control. The dynamical measure of this process is given in *Ticks-Per-Bit*, where entropy is written as

$$S = k_B \cdot \ln(1/TPB)$$

as derived in *Ticks-Per-Bit: A Microphysical Foundation for Entropy*. This establishes entropy not as a passive count of hidden microstates but as a measure of how efficiently a system's dynamics produce distinguishable changes.

But entropy in this broad sense exceeds physically realized information. To identify what is genuinely realized, stricter conditions must be imposed. A physically realized bit in a region A must satisfy three requirements:

- **Local readability**: it is accessible to some operation within A
- **Dynamical stability**: it survives the relevant coarse-grained evolution
- **Local-unitary irreducibility**: it cannot be eliminated by any transformation acting only within A

The third condition is decisive. If a putative distinction can be destroyed by acting only inside A, it was not stably encoded. Realized information must therefore be supported by correlations between A and its complement A^c.

Theorem (Restricted Uniqueness of Realized Information). Let $\Phi(A)$ be a measure of physically realizable information in region A satisfying:

1. Invariance under local unitary transformations on A or A^c separately
2. Additivity for independent subsystems: $\Phi(A_1 \otimes A_2) = \Phi(A_1) + \Phi(A_2)$
3. Stability under coarse-graining: Φ does not increase when local degrees of freedom are traced out

4. Vanishing on product states: $\Phi(A) = 0$ when $\rho_{AA^c} = \rho_A \otimes \rho_{A^c}$

Then, within the standard class of quantum information measures, $\Phi(A)$ is proportional to the quantum mutual information:

$$\Phi(A) \propto I(A{:}A^c) = S(\rho_A) + S(\rho_{A^c}) - S(\rho_{AA^c})$$

Proof sketch. Any functional satisfying conditions (1)–(4) must be expressible as a continuous function of ρ_{AA^c} that vanishes on product states, is additive under tensor products, and is invariant under local unitaries on each factor. Within the standard class of quantum information functionals, the relative entropy $D(\rho \,\|\, \sigma)$ is the unique primitive satisfying monotonicity under completely positive trace-preserving maps (the data-processing inequality). The mutual information $I(A{:}A^c)$ equals $D(\rho_{AA^c} \,\|\, \rho_A \otimes \rho_{A^c})$, which is the unique relative entropy between the joint state and the product of its marginals. It satisfies all four conditions: (1) by unitary invariance of relative entropy; (2) by additivity of relative entropy on tensor products; (3) by the data-processing inequality; (4) by definition when $\rho_{AA^c} = \rho_A \otimes \rho_{A^c}$. Alternative functionals in the standard class — Rényi mutual informations, coherent information, entanglement entropy — each violate at least one condition: Rényi mutual informations with parameter $\neq 1$ fail additivity under tensor composition; coherent information fails positivity and hence condition (4); entanglement entropy fails additivity for mixed states. The claim is restricted to this standard class. ■

The number of physically realized bits in region A is therefore

$$N_{real}(A) = I(A{:}A^c) / \ln 2$$

within the admissibility framework introduced in *Physical Admissibility: A Constraint-Based Foundation for Physics*. Physical information is fundamentally relational. What becomes real does so not because it is hidden inside a region, but because it is encoded in correlations that cannot be removed by local operations on either side.

---

3. Area Law

Once realized information is identified with mutual information, locality immediately imposes a geometric constraint. In a system with finite interaction range and finite local Hilbert-space dimension, correlations between region A and its complement A^c can only be mediated through degrees of freedom within the boundary ∂A. The area law follows directly:

$$N_{real}(A) \sim |\partial A|$$

This is not an independent holographic postulate. It is derived from the combination of (i) mutual information as the unique realized-information measure, (ii) locality with finite interaction range, and (iii) finite on-site Hilbert space dimension.

The structural consequence that drives the rest of this paper is immediate: any microscopic physical model must be boundary-organized. Realized distinguishability cannot reside in independent bulk degrees of freedom — it must propagate through boundary-supported constraint structures. The fold and closure theory are precisely the minimal such structures.

---

4. The Fold Interface and the Emergence of $\mathbb{C}^4$

The area law restricts what microscopic physical structure can look like. If realized information is boundary-supported and fact formation is irreversible, there must exist a surface at which reversible distinctions become committed records. In VERSF this interface is the *fold*.

The fold is a necessity derived from admissibility requirements, as formally derived in *The Fold Interface Law*. The derivation is reproduced here in sufficient detail to make the $\mathbb{C}^4$ conclusion explicit and checkable.

4.1 Two Independent Binary Degrees of Freedom

Commitment polarity. The minimal physical fact is binary: a one-valued event carries no information; any multi-valued distinction decomposes into binary ones. Commitment therefore requires a polarity variable $\sigma \in \{0, 1\}$.

Orientation independence. A commitment boundary must additionally distinguish the reversible pre-commitment side from the committed record side, while still admitting reversible transformations before commitment occurs. This requires an orientation variable $\omega \in \{-1, +1\}$.

These two variables must be independent. If they were the same variable, then reversing orientation — a reversible pre-commitment transformation — would alter the committed fact. An irreversible record would then be vulnerable to a reversible operation, directly contradicting the irreversibility requirement of the admissibility framework. Their joint space is therefore

$(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$

which contains exactly four distinct states.

4.2 Minimal Reversible State Space

To obtain a reversible dynamical representation of these four states, we require the minimal linear state space capable of representing two independent binary variables faithfully under reversible (unitary) dynamics.

A two-dimensional space over any field cannot independently encode two binary degrees of freedom: there is no way to assign two commuting binary operators on a 2-dimensional space that generate genuinely independent structure without identifying them. A three-dimensional real space $\mathbb{R}^3$ admits three mutually independent binary observables in principle, but lacks sufficient room to simultaneously represent both binary degrees of freedom together with the transformations mixing them without degeneracy. The minimal faithful linear realization is therefore $\mathbb{F}^4$ over some field $\mathbb{F}$.

4.3 Field Selection Under Admissibility

The choice of field is constrained by the admissibility requirement that distinguishability transitions satisfy invariant projection geometry — specifically, that there exists a well-defined notion of transition amplitude between interface states that is invariant under the group of reversible (symmetry) transformations of the interface and compatible with Born-rule probability.

$\mathbb{R}^4$ is insufficient. The symmetry group acting on rays in real projective space $\mathbb{R}P^3$ is PO(4). This group does not preserve a Hermitian inner product on the ray space in the sense required: there is no PO(4)-invariant Hermitian form on $\mathbb{R}^4$ that assigns unit probability to orthogonal rays and yields a well-defined probability calculus. More precisely, real Hilbert spaces admit only symmetric bilinear forms, and the group of isometries of $\mathbb{R}P^3$ does not support a complex phase structure under which probability amplitudes can interfere with arbitrary relative phase. Without this, the interference structure necessary for distinguishability transitions — in which different paths to the same committed fact can constructively or destructively interfere — cannot be realized. The resulting transition probabilities are not invariant under all admissible reversible transformations, violating condition (1) of the realized-information theorem.

$\mathbb{H}^4$ is non-minimal. Quaternionic spaces are mathematically consistent and support a richer phase structure. However, they introduce non-commutativity of the phase algebra (since $\mathbb{H}$ is non-commutative) that goes beyond what the two-variable interface requires. Specifically, the quaternionic phase group is $Sp(1) \cong SU(2)$, acting non-trivially on the phase degrees of freedom in a way that conflates phase transformations with orientation transformations. This violates the independence of σ and ω established above: quaternionic phase rotations can mix the commitment and orientation channels in ways that $\mathbb{C}$ cannot. Minimality therefore excludes $\mathbb{H}^4$.

$\mathbb{C}^4$ is the minimal admissible choice. The complex field satisfies three conditions that together are both necessary and sufficient for the admissibility requirements of the interface:

- **(F1) Nontrivial interference**: $\mathbb{C}$ supports a U(1) phase group under which amplitude superpositions yield constructive and destructive interference, enabling the full range of distinguishability transitions.
- **(F2) Invariant transition probabilities**: the unitary group U(4) acting on $\mathbb{C}^4$ preserves the Fubini–Study metric on $\mathbb{CP}^3$, giving Born-rule probabilities that are invariant under all admissible reversible transformations of the interface.
- **(F3) Commutative phase algebra**: the U(1) phase group of $\mathbb{C}$ is commutative, preserving the independence of σ and ω under phase transformations.

Proposition (Field uniqueness). $\mathbb{C}$ is the unique minimal field satisfying (F1)–(F3). $\mathbb{R}$ fails (F1): real amplitudes do not support a nontrivial phase group, so interference and Born-rule invariance under the full symmetry group of $\mathbb{R}P^3$ cannot be achieved simultaneously. $\mathbb{H}$ fails (F3): the quaternionic phase group $Sp(1) \cong SU(2)$ is non-commutative, and its action mixes the commitment and orientation channels, violating σ–ω independence. No field intermediate between $\mathbb{R}$ and $\mathbb{C}$ in the ordered tower $\mathbb{R} \subset \mathbb{C} \subset \mathbb{H}$ satisfies all three conditions. $\mathbb{C}$ is therefore the minimal admissible choice, and the minimal admissible interface state space is:

$\mathscr{H}$_fold $\cong \mathbb{C}^4$

This result is consistent with the independent derivation of the necessity of complex state space in *Complex Hilbert Space from Distinguishability Principles*, which establishes $\mathbb{C}$ as the unique admissible field for quantum state representation from interference and symmetry requirements alone. The associated projective state space $\mathbb{CP}^3$ is the natural geometric setting for the symmetry and channel structure of the interface.

---

5. Geometric Channel Structure of the Fold

The fold is a two-dimensional surface embedded in three-dimensional space and carries the $\mathbb{C}^4$ state structure derived above. The task here is to determine what independent local geometric data are available at that surface, and how they decompose. This analysis is developed fully in *The Void, The Fold, and the Derivation of Spacetime, Quantum Correlations, and the Standard Model Gauge Group*; the essential structure is reproduced here.

5.1 Exhaustiveness from the Fundamental Theorem of Surface Theory

The fundamental theorem of surface theory states that a smooth surface embedded in $\mathbb{R}^3$ is determined, up to rigid motion, by two tensor fields on the surface: the first fundamental form **g** (the induced metric) and the second fundamental form **h** (the extrinsic curvature tensor). Together these exhaust all local geometric information at the boundary. No third independent local geometric structure exists at this differential-geometric level.

This guarantees that whatever channel decomposition we identify below is not a selection from a larger set — it is the complete set.

5.2 Three Geometric Data Channels

Channel 1 — Scalar separation (metric channel). The first fundamental form **g** is a positive-definite symmetric 2×2 tensor on the surface. In two dimensions, a conformal factor determines the local scale. The scalar (conformal) degree of freedom corresponds to the local metric scaling: a single real number at each boundary point. Under the complex structure on $\mathbb{C}^4$, this scalar real degree of freedom constitutes a one-complex-dimensional geometric data channel. We denote this as the scalar sector:

$\dim_{\mathbb{C}} = 1$

Channel 2 — Spinorial orientation (normal channel). The fold is a two-dimensional surface in three dimensions and therefore carries a local unit normal vector field. The orientation of the boundary — which way it faces — is encoded in this normal. Once the complex amplitude structure of $\mathbb{C}^4$ is in place, the local orientation transforms under the spin cover of SO(2) rotations in the boundary tangent plane, which is U(1), together with the sign ambiguity of the normal. The minimal complex representation that faithfully encodes both the normal direction and its orientation sign is the spinor representation, which is two-dimensional over $\mathbb{C}$. This gives:

$\dim_{\mathbb{C}} = 2$

Channel 3 — Extrinsic curvature (shape channel). The second fundamental form **h** is a symmetric 2×2 real tensor on the surface, encoding the extrinsic curvature — how the surface bends in the ambient space. A symmetric 2×2 real matrix has exactly three independent real components (two diagonal entries and one off-diagonal). Under the complexification already forced by the $\mathbb{C}^4$ structure, these three real degrees of freedom are grouped into a three-real-dimensional (equivalently, in the context of the complex amplitude structure, a three-complex-dimensional geometric data) channel. We denote this:

$\dim_{\mathbb{C}} = 3$

5.3 Precise Statement of the Decomposition

A potential source of confusion must be addressed directly. The three geometric data channels have complex dimensions 1, 2, and 3. These do not sum to 4, the complex dimension of $\mathbb{C}^4$. This is not an inconsistency — it reflects a categorical distinction that is easy to miss.

The $\mathbb{C}^4$ state space and the geometric data channels are objects of different kinds. $\mathbb{C}^4$ is the space of reversible amplitude states carried by the fold interface — the arena in which quantum superpositions and transitions are represented. The geometric data channels are not subspaces of that arena. They are independent structures defined by the local differential geometry of the fold as an embedded surface: the metric, the orientation, and the extrinsic shape. These geometric structures *classify* or *act on* the interface states; they do not partition them.

An analogy may help. A particle moving in three-dimensional space has a state space $\mathbb{R}^6$ (position and momentum). The geometric data available at a surface that particle crosses — its normal direction, its curvature — are structures associated with the surface, not subspaces of $\mathbb{R}^6$. No one expects the dimensions of the surface geometry to sum to 6.

The situation here is the same. The fold carries a $\mathbb{C}^4$ state space (complex dimension 4, real dimension 8). The local differential geometry of the fold — exhausted by the first and second fundamental forms — determines three independent geometric data channels of complex dimensions 1, 2, and 3. These channel dimensions appear in the subsequent closure counting because discrete closure theory is realizing the geometric structure of the fold, not the dimension of its state space. The two quantities live at different levels of description and are not expected to match.

6. From Fold Geometry to Discrete Closure: A Theorem Schema

We now establish the logical bridge between the continuous fold geometry and the discrete closure integers. The purpose of this section is not to replace the full derivations in the companion closure papers, but to provide an explicit theorem schema — a precise statement of what conditions characterize the discrete realization and why the fold geometry constrains it — so that the connection is stated as a result rather than asserted as an analogy.

6.1 The Realization Problem

The fold supports two independent binary degrees of freedom (σ, ω) and three exhaustive geometric data channels with dimensions (1, 2, 3). This is a continuous geometric structure. The admissibility framework further requires that any physical realization operate under finite capacity: a finite number of degrees of freedom, finite interaction range, and finite distinguishability per site. The question is therefore:

What is the minimum discrete structure, operating under finite capacity and global consistency, that faithfully realizes the fold's binary degrees of freedom and three-channel geometry?

6.2 Admissibility Conditions on Discrete Realizations

A discrete realization is *admissible* if it satisfies the following conditions:

D1 — Completeness: it represents both binary degrees of freedom of the fold independently and fully; no fold binary structure is conflated or missing.

D2 — Channel fidelity: it supports three geometrically independent families of constraints corresponding to the scalar, spinorial, and extrinsic-curvature channels; no two channel types are realized by the same constraint family.

D3 — Finite capacity: the total number of independent binary constraints is finite and minimal — no redundant constraint is included.

D4 — Global consistency: distinguishability assignments are path-independent; moving through distinct constraint paths in the realization does not produce contradictory assignments. This is enforced by loop-consistency relations.

D5 — Irreducibility: no strict sub-system of the constraints satisfies D1–D4.

6.3 The Closure Parameters as Minimum Admissible Counts

Theorem Schema (Minimum Admissible Realization). Under conditions D1–D5, the discrete realization of the fold-induced binary and channel geometry requires:

- A minimum number K of independent binary constraints, where K is determined by the smallest system of binary conditions that independently encodes both fold binary degrees of freedom across the three geometric channel types under D1–D3 and D5.
- A minimum number N of independent loop-consistency channels, where N is determined by the smallest set of loop relations enforcing global consistency across the K constraints under D4 and D5.

The concrete mapping from geometry to constraints is as follows. Each geometric data channel — scalar separation, spinorial orientation, extrinsic curvature — induces an independent family of binary constraints on admissible assignments: the scalar channel constrains metric compatibility; the spinorial channel constrains orientation consistency under the spin cover; the extrinsic channel constrains curvature propagation across adjacent boundary regions. The three families are independent by D2, and each must be realized by a sufficient number of binary constraints to span the geometric degrees of freedom of that channel under the fold's binary variables. The minimization over all such realizations satisfying D1–D5 is a finite combinatorial problem: the admissibility conditions define a finite class of candidate structures (finitely many binary constraint graphs on finitely many nodes, with each node corresponding to a binary degree of freedom), and the solution is determined by explicit enumeration. Uniqueness follows from constraint independence (D3) and irreducibility (D5) jointly: any two distinct minimal realizations would share a common sub-structure satisfying D1–D4, contradicting D5.

The explicit solution of this minimization problem, yielding $K = 7$ and $N = 14$, is derived in *Structural Sufficiency, Risk Concentration, and Representation Selection in the One-Fold Framework*, which carries out this enumeration and confirms the values

$K = 7, N = 14$

The present paper establishes that the theorem schema above is the correct formulation — that K and N are the unique outputs of a well-posed and finite minimization problem defined by the fold geometry and conditions D1–D5. The two papers together constitute a complete derivation.

6.4 Why Compelled Rather Than Compatible

The theorem schema shifts the logical status of $K = 7$ and $N = 14$ from *compatible with* the fold to *required by* it. The fold geometry sets conditions D1–D5. The minimum admissible realization of those conditions is a well-posed finite combinatorial problem with a unique solution: the admissibility conditions define a finite, explicitly enumerable class of candidate constraint graphs, and uniqueness follows from the joint force of constraint independence (D3) and irreducibility (D5), as argued above. The values $K = 7$ and $N = 14$ are statements about the solution to that problem, not choices made to fit downstream results. The coupling formula that follows is therefore as constrained as the fold geometry that originates it.

On the reverse-engineering objection. A natural concern is whether $K = 7$ and $N = 14$ were discovered by working backward from the known value $\alpha^{-1} \approx 137$ — that is, by asking what integers satisfy $2^K \cdot (N+1)/N \approx 137$ and then constructing post-hoc geometric justifications for them. This concern is addressed directly by the logical structure of the derivation chain. The fold geometry, and therefore the admissibility conditions D1–D5, is derived in Sections 2–5 entirely without reference to α or to any target numerical value. The conditions D1–D5 are constraints on what a discrete physical realization of that geometry must satisfy; they make no mention of coupling strengths or dimensionless constants. The companion closure paper derives K and N from those conditions by explicit enumeration of admissible binary constraint graphs — a combinatorial procedure that terminates at $K = 7$ and $N = 14$ because those are the minimum counts satisfying D1–D5, not because they reproduce a desired number. The coupling formula of Section 7 is then applied to those values. The direction of derivation is therefore geometry → constraints → coupling, not coupling → constraints → geometry. Readers wishing to verify this independently should consult the companion paper, where the enumeration is carried out in full before any connection to α is drawn.

---

7. Structural Origin of the Coupling Formula

Before deriving the two factors, it is necessary to state the principle connecting the closure structure to observable coupling. Within the VERSF framework, electromagnetic coupling measures how readily distinguishability transitions occur across the fold interface — the probability per unit action that a charged system undergoes a distinguishability-changing interaction with the field. This rate is set by the total distinguishability weight of the closure structure: the number of independent ways the structure can encode a stable, committed distinction. A system with more independent distinguishability channels couples more weakly to any single channel, because the same total distinguishability weight is distributed across more configurations. The inverse coupling α^{-1} therefore measures the total distinguishability capacity of the closure structure — large α^{-1} corresponds to a structure with many independent constraint configurations, making individual transitions rare relative to the available phase space.

This identification is not circular. It is the discrete analogue of the general principle established in companion work (*Vacuum Impedance, Resistance Quanta, and the Fine-Structure Constant*) that α is the impedance ratio $Z_0/(2R_K)$ — which itself measures the ratio of the field's available coupling channels to the quantum of resistance imposed by binary commitment. The closure structure derived here is the microscopic realization of that ratio. With this principle in hand, the two structural factors follow directly.

The electromagnetic coupling emerges from two structural effects that act at different levels of the closure system and therefore combine multiplicatively.

7.1 Combinatorial Capacity: the 2^K Factor

Each of the K independent binary closure constraints represents a single irreducible binary distinction. The constraints are, by condition D3, independent: no constraint is implied by any combination of others. The total number of admissible configurations across all K independent constraints is therefore exactly

2^K

This factor reflects pure combinatorial binary capacity. It counts the configurations of the closure system in the absence of any further global condition. The connection to the entropy unit ln 2 is direct: each independent binary constraint contributes exactly ln 2 to the total distinguishability capacity, giving a total combinatorial weight of K ln 2, or equivalently 2^K configurations.

7.2 Global Consistency Normalization: the (N+1)/N Factor

The N loop-consistency channels enforce path-independence of distinguishability assignments across the K constraints. As independent loop channels, they provide N degrees of freedom in the consistency structure. However, the closure system must satisfy one further global condition: the entire closure, taken as a single structure, must be internally consistent — not merely locally loop-consistent, but globally so. This global condition is not implied by any finite collection of local loop conditions; it is an additional, independent constraint.

Before imposing this global condition, the normalization of the N loop channels counts N + 1 effective normalization degrees of freedom: N local channel normalizations plus 1 global normalization. Imposing the global consistency condition eliminates exactly one of these, leaving N effective degrees of freedom. The ratio of constrained to unconstrained normalization is therefore

N / (N+1)

and the normalization shift — the factor by which the coupling differs from a system with only the N local conditions — is the reciprocal:

(N+1) / N

The counting argument presented here is a structural summary. The formal grounding — specifically, the derivation that the global closure condition reduces the effective normalization count by exactly one degree of freedom, and that the resulting factor enters the coupling multiplicatively — is established rigorously in *Structural Sufficiency, Risk Concentration, and Representation Selection in the One-Fold Framework*, where the normalization structure of the closure system is analysed in full. Readers who wish to verify the (N+1)/N factor from first principles should consult that paper directly.

To see why this factor enters multiplicatively: the combinatorial capacity 2^K and the global consistency normalization (N+1)/N arise from structurally independent features of the closure system. The combinatorial capacity counts configurations of the K binary constraints; the normalization factor adjusts the effective coupling weight due to the topological property of global closure. These two effects are not additive corrections to a single shared quantity. They are independent multiplicative contributions to the total distinguishability weight. Their combination is therefore

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

within the closure framework established in *Structural Sufficiency, Risk Concentration, and Representation Selection in the One-Fold Framework* and the companion discrete realization papers. Additive combination would correspond to the two effects sharing a common degree of freedom, which D3 and D5 exclude by minimality and irreducibility.

7.3 Numerical Result

Substituting K = 7 and N = 14:

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

The observed low-energy value is $\alpha^{-1} \approx 137.036$. The structural value $\alpha^{-1} \approx 137.14$ is interpreted as a UV structural coupling. The deviation from the measured infrared value reflects renormalization-group running, which lies outside the present scope and is expected to modify the value at the level of the observed discrepancy.

---

8. Scope and Status of the Result

The following summary states precisely what this paper establishes, what is established jointly with companion work, and what lies outside the current scope.

Established within this paper:

- Mutual information is the unique measure of physically realized information within the standard class of quantum information functionals (Section 2)
- Locality and finite capacity give an area law for realized information (Section 3)
- Admissibility and minimality force a commitment interface carrying $\mathbb{C}^4$ under well-defined field-selection criteria (Section 4)
- The local differential geometry of the fold yields three exhaustive geometric data channels of dimensions 1, 2, and 3 (Section 5)
- The theorem schema relating fold geometry to discrete closure parameters K and N via conditions D1–D5 (Section 6)
- The coupling formula from K, N and its structural justification (Section 7)

Established jointly with companion work:

- The exact values K = 7 and N = 14 as solutions to the minimum admissible realization problem

Outside current scope:

- Renormalization-group flow from the structural UV coupling to the measured low-energy value
- Extension of the derivation to non-electromagnetic couplings

The claim of this paper is architectural: it establishes the continuous structural route from admissibility constraints to coupling, within which companion results are correctly integrated rather than merely appended. The fine-structure constant is not arbitrary because every step in its determination — from the relational character of realized information, through the fold geometry and its discrete realization, to the combinatorial and topological contributions to coupling — is fixed by conditions that any universe supporting stable physical facts must satisfy.

---

9. Conclusion

The fine-structure constant is usually treated as unexplained input. The VERSF programme proposes a different picture. If a universe must support finite distinguishability, irreversible commitment, and finite capacity, then physically realized information is necessarily relational and obeys an area law. The same admissibility conditions force a minimal commitment boundary — the fold — with a derived $\mathbb{C}^4$ complex state structure. The local differential geometry of that boundary yields three exhaustive data channels. The minimum admissible discrete realization of that geometry under finite capacity and global consistency is characterized by integers K and N. The resulting structural coupling is

$$\alpha^{-1} = 2^K \cdot (N+1)/N \approx 137$$

In this framework, α is not arbitrary. Its numerical form is rooted in the relational geometry of distinguishability, the topology of commitment, and the finite closure conditions required for stable structure to exist. The broader implication is that physical law may be understood not as dynamics imposed on an already-given world, but as the consequence of constraints on what kinds of distinguishable structures can exist at all.

Appendix A. Structural Proposition: K and N as Minimum Admissible Counts

Full derivations of the closure parameters K and N are provided in the companion papers:

- *Structural Sufficiency, Risk Concentration, and Representation Selection in the One-Fold Framework* — derives K = 7 and N = 14 by explicit enumeration of admissible binary constraint graphs satisfying conditions D1–D5

- *The Fold Interface Law* — derives the existence and binary structure of the fold that sets those conditions
- *The Void, The Fold, and the Derivation of Spacetime, Quantum Correlations, and the Standard Model Gauge Group* — develops the geometric channel structure that determines the channel-fidelity condition D2

The purpose of this appendix is to summarize the structural role of K and N within the present derivation, making their conceptual distinctness and necessity precise without reproducing the full combinatorial proofs.

Proposition A.1 (Necessity of a finite K). Let the fold carry two independent binary degrees of freedom and three geometrically independent data channels (as derived in Sections 4–5). Under admissibility conditions D1, D2, D3, D5, there exists a finite minimum number K of independent binary constraints required to faithfully realize the fold's binary and channel structure. This minimum is strictly greater than 2 (the number of binary degrees of freedom at the fold) because each geometric data channel imposes its own independent family of constraints, and the joint realization of all three channels with both binary variables under independence and minimality requires additional constraints beyond those encoding the binary variables alone. The companion closure paper derives $K = 7$.

Proposition A.2 (Necessity of a finite N). Given K independent binary constraints satisfying D1–D3 and D5, global consistency (D4) requires a set of loop-consistency relations. The number of independent loop-consistency channels N is determined by the cycle structure of the constraint graph: it equals the number of independent cycles, which is a topological invariant of the closure structure. This count is not a free parameter but is fixed once K and the channel geometry are determined. The companion closure paper derives $N = 14$.

Proposition A.3 (Distinct structural roles). K and N measure different properties of the closure system:

- K measures *binary closure complexity*: the number of independent distinguishability conditions required to realize the fold's geometry as a stable finite structure. It reflects the combinatorial capacity of the system.
- N measures *loop-consistency complexity*: the number of independent topological relations required to enforce path-independent distinguishability. It reflects the global coherence of the system.

These are not two ways of measuring the same thing. K is a constraint count; N is a cycle count. In the coupling formula, they enter through independent mechanisms: K through combinatorial capacity 2^K, N through global normalization $(N+1)/N$. No smaller admissible realization exists by the minimality condition D5.

Proposition A.4 (No smaller admissible realization). Any discrete structure satisfying D1–D5 with $K' < 7$ or $N' < 14$ fails at least one condition: either it does not independently represent all three geometric channel types with both binary fold variables (violating D1 or D2), or it contains a redundant constraint (violating D3), or it fails global path-independence (violating D4). The companion paper verifies this by explicit enumeration of all smaller candidate structures.

Programme Note

This paper forms part of the broader VERSF (Void Energy–Regulated Space Framework) research programme, which develops physical law and physical constants from constraints on distinguishability, irreversibility, and finite capacity. A complete set of companion papers — including detailed derivations of the fold interface, discrete closure structure, quantum state space, entropy dynamics, and connections to the Standard Model gauge group — is available at www.versf-eos.com. Readers are encouraged to consult the full programme for the technical foundations underlying each step of the derivation presented here.