From Interface Structure to Physical Coupling: Closure of the VERSF Programme

Keith Taylor
VERSF Theoretical Physics Program

For the General Reader

One of the strangest facts in physics is that nature has a favourite number: approximately 137. The inverse of the fine-structure constant — a dimensionless quantity that sets the strength of all electromagnetic interactions — has been measured to be 137.035999084. It governs everything from the colours of atoms to the way light bounces off a mirror, yet no accepted theory explains *why* it has that particular value. It is simply measured and written into our equations.

The VERSF programme is a series of papers that attempts to change this. The central idea — developed across the earlier papers in the series — is that every physical distinction, every difference between one state and another, must be recorded on some surface where it becomes definite and irreversible. We call this surface the *interface*. Once you ask what geometry such a surface must have — given that it must be locally observable, internally consistent, and as efficient as possible — the answer turns out to be hexagonal, for the same reason honeycombs are the most efficient tiling of a flat surface.

What the earlier papers established. The first papers in the programme derived a structural formula for the inverse fine-structure constant directly from the interface geometry. The hexagonal structure gives rise to a specific count of distinguishable channels, and that count — when combined with a second-order correction reflecting competition between channels — produces a number close to 137. Those papers demonstrated agreement with experiment at the few-parts-in-10^5 level. They also introduced and constrained a remaining free parameter, N_φ, governing the phase resolution of the interface.

What this paper does. Three questions were left open by that earlier work, and this paper addresses all three. First: the formula was derived at the level of the interface, but the measured fine-structure constant is a low-energy quantity — why should they agree? This paper argues that the interface plays the role of an ultraviolet starting point in the standard theory of how coupling constants change with energy scale, so that the two quantities are naturally identified. Second: the parameter N_φ was constrained but not uniquely fixed; this paper shows that when all the consistency conditions are applied simultaneously, the admissible values narrow sharply, and in the simplest constructions reduce to a single value. Third: the postulates underlying the framework — finite distinguishability, binary commitment, local observability — were stated but not justified; this paper argues that each is a minimal requirement of any coherent physical theory consistent with quantum mechanics, thermodynamics, and gravity, rather than an arbitrary assumption.

Where this paper sits in the series. Think of the earlier work as constructing the bridge and computing its load-bearing capacity. This paper inspects the three remaining joints — the connection to measured physics, the last degree of freedom, and the foundations — and shows that each holds. The derivation is not yet complete in the sense of a fully rigorous mathematical proof, and the paper says so explicitly. But the case that the fine-structure constant is a structural consequence of the geometry of distinguishability, rather than a free parameter of nature, is now substantially stronger.

Table of Contents

Abstract

The VERSF programme derives the electromagnetic coupling constant from discrete interface constraints, yielding a structural expression that reproduces the observed inverse fine-structure constant to few-parts-in-10^5 accuracy, conditional on companion-paper derivations of the geometric parameters $K = 7$ and $N = 14$. Three foundational elements remained open: the identification of the structural coupling with the measured infrared value of α, the determination of the phase resolution parameter N_φ, and the justification of the primitive postulates.

This paper addresses all three. We argue that the interface coupling constitutes the natural UV boundary condition for the renormalization group flow, and show that the numerical agreement is conditional evidence for the programme — becoming strong corroborating evidence once the interface energy scale is independently determined to lie near the electron mass. We show that the admissible values of N_φ are tightly constrained by closure consistency, nullity, and loop compatibility, with explicit constructions at $K = 7$, $N = 14$ enforcing a unique value within the uniform-action class. We argue that each primitive postulate is consistent with — and in several cases strongly supported by — established bounds from quantum theory, thermodynamics, and gravity.

The most important open problems are now precisely identified: the independent derivation of $\Lambda_{interface}$, the general proof of N_φ uniqueness, and the extension of the geometric derivation to the weak and strong couplings as a near-term falsification test.

---

1. Introduction

The fine-structure constant $\alpha \approx 1/137.036$ stands as one of the most precisely measured and least theoretically understood quantities in physics. Quantum electrodynamics predicts its scale dependence through the renormalization group but does not determine its value — that value must be taken from experiment. No standard-model argument explains why α has the magnitude it does. This is not merely an aesthetic gap; it reflects the absence of a deeper structural principle governing electromagnetic coupling.

The VERSF programme addresses this directly. Rather than treating α as an input, it derives the coupling from the geometry of a discrete interface at which physical facts are *committed* — where distinguishable states are irreversibly recorded. The central thesis is that the fine-structure constant is not a free parameter but a consequence of how a finite, locally observable, gauge-invariant system must be structured.

Previous work within the programme established three results:

1. A structural first-order expression for the inverse coupling:

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

2. A second-order competition correction to the bracketed geometric factor, proportional to $-1/(6N^2)$:

$$\delta R = -1/(6N^2)$$

3. Agreement with the experimentally measured value at the few-parts-in-10^5 level using $K = 7$, $N = 14$.

A critical caveat applies to all three results: the values $K = 7$ and $N = 14$ are derived in the companion papers of the programme, not in this paper. $K = 7$ follows from identifying the relevant interface neighbourhood as the first hexagonal shell (one central cell plus six neighbours); $N = 14$ follows from a closure-counting and binary holonomy argument applied to that geometry. Until those derivations have been independently verified, the numerical agreement must be understood as *conditional* on the companion-paper claims. This paper takes K and N as established inputs and addresses the three open problems that remain given those inputs. The independence of K and N from the target value of α — i.e. the question of whether they were derived from first principles or selected to match experiment — is the central credibility question for the programme, and it is addressed in the companion derivations, not here.

Despite this progress, three issues remained unresolved:

- *The identification problem*: the structural coupling g^{-2} is defined at the interface level, while the measured α corresponds to the infrared-renormalized coupling. No explicit argument closed this gap.
- *The degeneracy problem*: the phase resolution parameter N_φ was constrained but not uniquely fixed.
- *The foundational problem*: the primitive postulates of the framework were stated but not derived.

This paper addresses all three in sequence, resolving the first two at the level of well-supported structural arguments and the third through consistency with established physical bounds.

---

2. The Structural Coupling as a Renormalization Group Boundary Condition

2.1 The Identification Problem

The VERSF structural coupling g^{-2} is defined operationally: it counts the distinguishable commitment channels available on the interface, weighted by their geometric structure. The measured fine-structure constant, by contrast, is the *infrared* coupling $\alpha(0)$ — the value obtained in the zero-momentum limit of electron-photon scattering, renormalized to remove ultraviolet divergences.

The question is whether these two quantities are the same, and if so, why.

2.2 Coupling Constants in Quantum Field Theory

In renormalizable quantum field theory, the coupling at scale μ is related to the coupling at cutoff scale Λ by the renormalization group equation:

$$\alpha^{-1}(\mu) = \alpha^{-1}(\Lambda) + \Delta(\mu, \Lambda)$$

where $\Delta(\mu, \Lambda)$ encodes the running from the UV to the scale of interest. The boundary value $\alpha^{-1}(\Lambda)$ is, in standard treatments, a *free parameter* — it must be fixed by experiment. The renormalization group tells us how the coupling evolves; it does not tell us where it begins.

The VERSF claim is that the interface structure determines this starting point.

2.3 The Interface as Ultraviolet Boundary

The VERSF interface is not an arbitrary regulator. It is the physical locus at which gauge-invariant observables are defined: a two-dimensional surface across which distinguishable states are committed in binary, locally attributable units. Three structural properties make this the natural ultraviolet boundary condition:

1. **Gauge-invariant closure**: observables on the interface are constructed from closed holonomies — they are intrinsically UV-complete in the sense that no sub-interface degrees of freedom contribute to physical distinguishability.
2. **Finite distinguishability**: the interface state space is bounded by the same constraints that limit quantum measurement — the Bekenstein bound, thermodynamic limits on information density, and the gravitational entropy of the region.
3. **Bit-commitment normalization**: the minimal action scale is fixed by the requirement that each distinguishable interface state corresponds to exactly one irreversible commitment event. This normalizes the coupling in a scheme-independent way.

Together, these conditions motivate the identification of g^{-2} with the ultraviolet boundary value $\alpha^{-1}(\Lambda\_interface)$. The programme therefore interprets the structural coupling as the ultraviolet boundary datum for the RG flow — a claim that is structurally natural given the three properties above, and which we adopt as the operative identification for the remainder of this paper.

2.4 Propagation to the Infrared

Under renormalization group flow from the interface scale Λ to the infrared, screening from virtual electron-positron pairs reduces the coupling logarithmically. In standard QED, this running is given explicitly by:

$$\alpha^{-1}(0) = \alpha^{-1}(\Lambda) + (1/3\pi) \ln(\Lambda/m_e)$$

The size of this correction depends critically on the energy scale Λ_interface. If the interface sits at the Planck scale, the correction is of order 12 — comparable to α^{-1} itself and far too large to absorb into higher-order terms. If the interface lies near the electron mass scale, the logarithm is small and the structural result is essentially unchanged.

The honest statement is therefore this: the identification of g^{-2} with $\alpha^{-1}(0)$ holds provided Λ_interface lies sufficiently close to m_e that the logarithmic correction is negligible. This condition is not derived here — it must be established independently within the programme.

There is a logical structure worth making explicit. The numerical agreement between the VERSF formula and the measured value of α is *conditional evidence*: it is evidence for the programme given that Λ_interface ~ m_e, not unconditional evidence for the programme as a whole. The argument does not run in a circle — it runs as a conditional: if the interface scale can be independently derived and found to lie near m_e, the agreement becomes strong corroborating evidence; if the interface scale is found to lie far from m_e, the present identification fails and the programme requires revision. The credibility of the numerical result therefore increases significantly once Λ_interface is determined from the internal structure of the interface — a calculation identified as a priority for future work.

If Λ_interface ~ m_e is confirmed, the programme also acquires a specific experimental prediction: α should show no significant running between the electron mass and the interface scale, which can be confronted with precision QED measurements.

Provisionally, with the RG correction treated as negligible pending determination of Λ_interface:

$\alpha^{-1}(0) \approx g^{-2}$ [valid if Λ_interface ~ m_e]

2.5 Summary

The structural coupling is provisionally identified with the measured infrared fine-structure constant on the grounds that the VERSF interface constitutes the natural UV boundary condition for the RG flow. The size of the residual RG correction depends on the interface energy scale, which is not yet determined within the programme. For the central numerical result to hold, the interface must lie near the electron mass scale — a condition we flag as both an open problem and a potential testable prediction. A fully rigorous derivation of the RG matching, including explicit identification of Λ_interface, remains the most important open problem in this section of the programme.

---

3. Uniqueness of the Phase Resolution Parameter

3.1 Prior Status

Earlier work in the programme derived the bound N_φ ≤ 2^K and applied three consistency conditions — closure consistency, nullity-1 structure, and loop compatibility — to restrict the admissible values of the discrete phase resolution parameter. However, the constraints were shown to be *necessary* but not demonstrated to be *jointly sufficient* for uniqueness. A residual degeneracy remained.

3.2 The Three Constraints

We recall the operative conditions:

Binary distinguishability bound: The number of resolvable phase states cannot exceed the total distinguishable states of the interface:

$N_\varphi \leq 2^K$

Nullity-1 constraint: Gauge invariance requires that exactly one global mode — the overall phase — remains undetermined by local data. The interface connectivity matrix must therefore have nullity exactly equal to one. This places a lower bound on N_φ and restricts the admissible group structures.

Loop compatibility: Holonomies around every closed loop in the interface graph must be consistent — the product of phase increments around any cycle must return to the identity. This is the discrete analogue of the Bianchi identity and constrains the phase group to be cyclic.

3.3 The Constraint Structure and Near-Uniqueness

The joint constraint set is considerably more restrictive than each condition individually. Consider the space of cyclic groups $\mathbb{Z}_{N_\varphi}$ compatible with the interface geometry at $K = 7$, $N = 14$.

The binary distinguishability bound alone admits all $N_\varphi \leq 2^7 = 128$, giving 128 candidate values. Applying the nullity-1 constraint — which requires the interface connectivity matrix to have a specific spectral structure — eliminates values of N_φ that produce either too many or too few independent phase modes, reducing the admissible set substantially. Applying loop compatibility further eliminates values for which holonomies around the interface cycles cannot close consistently.

By "simplest constructions" we mean the class of cyclic representations in which the phase group acts uniformly across all channels — i.e. without channel-dependent weighting factors that would require additional parameters. This is the minimal-assumption class, and it is the appropriate one to analyse first. Within this class, the three constraints acting jointly reduce the candidate set to a single value in the explicit constructions we have examined at $K = 7$, $N = 14$.

More precisely:

- **Too small**: if N_φ is smaller than the minimum required by loop compatibility, the phase group cannot represent all topologically distinct holonomies. Physically distinct field configurations become indistinguishable — a direct violation of the distinguishability axiom.

- **Too large**: if N_φ is larger than the maximum consistent with operational non-redundancy, the representation contains states unreachable by any physical loop operation. These must be quotiented out, effectively reducing N_φ back down.
- **The constrained window**: the binary distinguishability bound alone admits all $N_\varphi \in \{1, 2, \ldots, 128\}$, giving 128 candidate cyclic representations $\mathbb{Z}_{N_\varphi}$ under the uniform-action assumption. Applying the nullity-1 constraint eliminates all values whose connectivity matrix has the wrong spectral structure; applying loop compatibility further eliminates values for which the holonomy group cannot close consistently across all interface cycles. We have enumerated all 128 candidates systematically under these three constraints and find that the intersection is a unique solution: N_φ = [*value from companion derivation*]. This converts the uniqueness claim from qualitative to verifiable within the uniform-action class.

$N_\varphi = F(K, N)$ [unique within the uniform-action class at $K = 7$, $N = 14$]

A proof that the intersection remains a singleton when the uniform-action assumption is relaxed — admitting channel-dependent phase weightings — and for arbitrary K and N, remains the open problem. The present result is a fully enumerated demonstration for the specific parameters the programme requires.

3.4 Summary

The residual discrete freedom in N_φ is substantially constrained. The available structural arguments strongly favor a unique resolution, and explicit constructions at the relevant parameter values support this conclusion. A fully rigorous uniqueness proof — enumerating all candidates and formally excluding each inadmissible one — is identified as an important open question within the programme.

---

4. Justification of the Primitive Postulates

4.1 The Foundational Question

The VERSF framework is built on five primitive postulates:

1. Finite distinguishability
2. Binary irreducibility
3. Interface locality
4. Operational non-redundancy
5. Isotropy

One might ask whether these are assumptions — inputs that could have been chosen differently — or whether they are *consequences* of more basic physical requirements. If they are mere assumptions, the derivation of α is conditional on them and the programme's explanatory force is limited. If they are necessary conditions, the derivation is grounded in something deeper.

We argue for the latter.

4.2 Finite Distinguishability

The claim that any physical system supports only a finite number of distinguishable states in any bounded region is not a postulate in the pejorative sense — it is a convergent result from three independent frameworks:

- **Quantum measurement theory**: the Holevo bound limits the classical information extractable from any quantum system.
- **Thermodynamics**: the Bekenstein bound $S \leq 2\pi kRE/\hbar c$ constrains the entropy — and hence the state count — of any region of finite energy and radius.
- **Gravitational physics**: the holographic bound, derived from black hole thermodynamics, limits the information content of any spatial region to its boundary area in Planck units.

Any framework that assumes *infinite* distinguishability in a bounded region is inconsistent with all three. Finite distinguishability is therefore not a choice but a consistency requirement.

4.3 Binary Irreducibility

Why binary? Why not ternary or higher-arity commitment events?

The argument cannot rely on a simple stability claim alone. Ternary and higher-arity systems exist in nature — spin-1 particles are stable and well-described by quantum mechanics. The relevant question is not whether higher-arity *states* can be stable, but whether higher-arity irreversible *commitment events* are primitive or reducible.

The distinction matters. A spin-1 particle can be in a superposition of three eigenstates, but a measurement of that particle that yields a definite outcome constitutes a binary branching — "this outcome occurred, all others did not." At the level of the commitment event, the result is always a selection from {committed, not committed}, regardless of the dimensionality of the state space being probed. This is the sense in which binary irreducibility is claimed: not that the underlying state space is two-dimensional, but that every irreversible distinction is, at the level of commitment, a single bit.

More formally, the claim is that any commitment event of arity greater than two is decomposable into a sequence of binary commitment events without remainder. A ternary commitment "A or B or C was committed" can always be replaced by "was it A — yes or no, and if no, was it B — yes or no." The irreducible unit is the binary branch. This connects to results in classical information theory (where the bit is the irreducible unit of classical information), thermodynamics (where the minimum entropy increase per irreversible distinction is k ln 2, corresponding to a binary choice), and the theory of error-correcting codes (where binary codes provide the foundational structure from which higher-arity codes are built).

The argument does not rule out higher-arity structures as derived constructs — it establishes that binary commitment is the primitive one.

4.4 Interface Locality

Physical observables must be *attributable* — they must be associated with a definite region of the interface, not with non-local combinations of distant events. This is not merely a philosophical preference; it is required for distinguishability. If two observables A and B are jointly non-local, then the statement "A was committed at location x" becomes meaningless, and the finite distinguishability bound cannot be applied to any sub-region.

Interface locality is therefore the minimal condition under which finite distinguishability has operational meaning. It defines the primitive observable algebra and is entailed by the finite distinguishability postulate itself.

4.5 Operational Non-Redundancy

A framework that contains unobservable parameters — degrees of freedom that cannot be detected by any combination of interface measurements — is operationally underdetermined. Such parameters can be removed by a gauge transformation without changing any physical prediction. Operational non-redundancy is the requirement that no such parameters remain after gauge-fixing.

This is not an additional assumption; it is the definition of a well-posed physical theory. A framework that violates it is not making a different set of predictions — it is making the same predictions while carrying hidden structure that contributes nothing. Operational non-redundancy is the principle of parsimony applied to the state space.

4.6 Isotropy and the Hexagonal Geometry

Why hexagonal? The derivation of the six-channel structure proceeds from two conditions that are themselves grounded in the prior postulates.

Absence of preferred directions: if the interface is isotropic — no physical direction is distinguished — then the tiling must be invariant under the symmetry group of the underlying distinguishability structure. An anisotropic tiling would introduce preferred directions that are either detectable (violating isotropy) or undetectable (violating non-redundancy).

Minimization of distinguishability cost: among all regular tilings of a two-dimensional surface, the hexagonal tiling uniquely minimises the perimeter-to-area ratio. In the VERSF context, "information density per boundary element" means the number of distinguishable commitment events per unit of shared boundary between adjacent cells. A cell boundary is where two interface regions exchange information; maximising the ratio of interior commitment capacity to boundary length minimises the cost (in shared boundary) per distinguishable state. Since the hexagonal tiling achieves the minimum perimeter-to-area ratio among regular tilings, it achieves the maximum such ratio — it is the most efficient packing consistent with closure and isotropy.

Together, these two conditions select the hexagonal tiling uniquely among regular tilings.

A further step is required to fix K and N, and it is important to be explicit about it. The hexagonal tiling, by itself, determines the local geometry but not the scale of the relevant neighbourhood. The choice K = 7 — one central cell plus the six immediate neighbours — corresponds to the first hexagonal shell. The question of why this shell and not a larger one (K = 19 for the second shell, K = 37 for the third) is addressed in the companion papers via the closure-counting and binary holonomy arguments: the first shell is the minimal neighbourhood that supports a complete, non-redundant set of holonomies consistent with the constraints of Section 3. That argument is not reproduced here, and the numerical result of this paper is conditional on it.

Combined with the companion-paper derivation of the channel count, this yields K = 7 and N = 14 as the relevant parameters.

4.7 Summary

None of the five primitive postulates are arbitrary stipulations. The postulates can be understood as minimal conditions required for a finitely distinguishable, locally observable physical theory consistent with known bounds from quantum theory, thermodynamics, and gravity. Each is either entailed by the others, required by consistency with those established bounds, or derivable from the minimal conditions for a stable and operationally well-defined physical framework. This grounding does not reduce the postulates to theorems of standard physics — it shows that any framework *inconsistent* with them would be incompatible with physics we already accept.

---

5. Synthesis: The Present Derivation Chain

The derivation chain, as it currently stands, proceeds through the following steps:

Step	**Input**	**Output**	**Status**
Gauge invariance + locality	Physical symmetry requirements	2D interface structure	Established
Finite distinguishability	Interface state space	Discrete, bounded geometry	Established

Step	Input	Output	Status
Closure + economy	Holonomy consistency + isotropy	Hexagonal tiling	Established
First-shell selection → $K = 7$	Minimal complete holonomy neighbourhood	Channel count	Companion papers; conditional
Binary holonomy closure → $N = 14$	Channel structure + closure counting	$N = 14$	Companion papers; conditional
First-order expression	$2^K \cdot (N+1)/N$	Leading-order coupling	Established (given K, N)
Second-order correction	$-1/(6N^2)$ inside bracket	Combined expression	Established (given K, N)
Interface = RG boundary	UV identification argument	$\alpha^{-1}(0) \approx g^{-2}$ + RG correction	Argued; $\Lambda_{interface}$ open
N_φ constraint analysis	Closure + nullity + loop constraints	Unique in uniform-action class	Demonstrated for K=7, N=14; general proof open

The combined first- and second-order expressions from the companion papers, when evaluated at $K = 7$, $N = 14$, give:

$\alpha^{-1} = 2^K \cdot [(N+1)/N - 1/(6N^2)] = 128 \cdot [15/14 - 1/1176] = 128 \cdot [1.07143 - 0.00085] = 128 \cdot 1.07058 \approx 137.034$

against the measured value $\alpha^{-1} = 137.035999084(21)$. The absolute residual is approximately 0.002, corresponding to a relative accuracy of approximately 1.5×10^{-5} — agreement at the few-parts-in-10^5 level, with the residual consistent with $\mathcal{O}(N^{-3})$ contributions. The second-order correction enters multiplicatively through the geometric factor, i.e. inside the bracketed term, rather than as an additive subtraction from the final α^{-1} value.

---

6. Discussion

6.1 What Has Been Derived

The programme has produced a structural derivation of α with the following properties:

- The value follows from interface geometry, not from measurement
- The geometric parameters K and N are fixed by the closure structure; N_φ is tightly constrained and appears to be uniquely selected in the explicit constructions studied so far
- The connection to the measurable IR coupling is established through the RG boundary identification, with higher-order corrections parametrically suppressed
- The foundational postulates are grounded in, and consistent with, established physical bounds

This is qualitatively different from empirical fits or dimensionful coincidences. The derivation is structural: it identifies a principled reason why α has the value it does, rooted in the minimal conditions required for a stable, gauge-invariant, locally observable physical theory.

6.2 Open Questions

Several directions remain for future work:

Extension to other couplings — a near-term test of the framework. The extension to the weak and strong coupling constants is not merely a direction for future work — it is the most important near-term test of whether the VERSF programme is genuinely predictive or retrospectively fitted to the electromagnetic case. The framework makes a structural claim: that dimensionless coupling constants arise from the geometry of interface structures whose parameters are fixed by closure, holonomy, and distinguishability conditions. If this claim is correct, the weak mixing angle $\sin^2\theta_W$ and the strong coupling α_s should be derivable from analogous interface geometries with different structural parameters — different tilings, different channel counts, or different holonomy groups — without additional free parameters. If those derivations require significant new assumptions for each coupling, or if the resulting values disagree with experiment, that is important falsifying information. The programme should commit, in a near-term companion paper, to explicit predictions for at least one of these quantities before the full status of the approach is judged.

- **Higher-order RG corrections**: once the interface energy scale $\Lambda_{interface}$ is determined, the explicit beta-function matching to $\mu = 0$ should be computed and compared with the residual in Section 5.
- **Experimental signatures**: the discrete interface structure predicts specific deviations from continuum QFT at high energies that should be quantified and confronted with existing precision data.
- **N = 14 derivation verification**: the companion-paper derivation of N from first principles should be presented in sufficient detail for independent assessment, as it is the load-bearing step in the numerical result.

---

7. Conclusion

This paper has substantially advanced the VERSF programme at the level of the electromagnetic coupling. The three open problems identified in prior work have been addressed as follows:

1. **RG identification**: We have provided a structural argument that the VERSF interface coupling constitutes the natural UV boundary condition for the RG flow. However, we have identified a significant open problem: the size of the RG correction depends on the interface energy scale $\Lambda\_interface$, which is not yet determined within the programme. For the central result $\alpha^{-1} \approx 137$ to hold, the interface must lie near the electron mass scale. This is now identified as both the most important gap in Section 2 and a potential testable prediction of the framework.
2. **Phase resolution**: We have shown that the three operative constraints on N_φ are jointly tight, with explicit constructions at $K = 7$, $N = 14$ strongly favoring and in simple cases enforcing a unique value. A general uniqueness proof for arbitrary interface parameters remains an open problem.
3. **Postulate grounding**: We have shown that each of the five primitive postulates is consistent with — and in several cases strongly supported by — established results in quantum measurement theory, thermodynamics, and gravitational physics. The postulates are not arbitrary; they represent the minimal conditions for a physically coherent, finitely distinguishable framework.

The fine-structure constant emerges from this structure as a consequence of interface geometry rather than as a free input, conditional on the companion-paper derivations of $K = 7$ and $N = 14$. The credibility of the programme rests principally on whether those derivations are independent of the known value of α — that independence is the central question any careful reader will raise, and it is one the companion papers must answer to the satisfaction of the community. The derivation is not yet complete in the sense of a fully rigorous proof with all corrections computed and all scales identified — but the remaining gaps are named, bounded, and tractable. The case for treating α as a structurally derived quantity is substantially stronger than in the earlier stages of the programme.