From Meta-Principles to Physical Uniqueness: A Reconstruction Theorem for VERSF

Keith Taylor VERSF Theoretical Physics Programme

Summary for the General Reader

Why does physics take the particular shape it does? Are fields, symmetries, quantization, and the other structural features of physical theory merely *one* workable framework among many, or are they *forced* by deeper requirements?

This paper argues that they are forced. We start from five intuitive principles that any serious physical theory would normally be expected to satisfy:

- **Facthood.** At least some statements about the world hold determinately — the world is not arbitrary.
- **Comparability.** Determinate facts can be compared with one another — otherwise they have no physical content.
- **Self-consistency.** The theory doesn't contradict itself.
- **Completeness.** The theory specifies its own content, without needing external parameters supplied from outside.
- **Record-persistence.** A fact, once established, remains available long enough to be compared with later facts.

From these five principles — and nothing else — we derive the core structural features of the VERSF framework as the unique possibility:

- a field-like "carrier" of persistent content (a specific kind of mathematical object: the Green's function of a self-adjoint elliptic operator);
- a particular logarithmic coupling between this carrier and an entropy-like quantity;
- a global closure law expressed in integer winding numbers (a topological, quantized relation);
- an underlying discrete "substrate" with a specific combinatorial structure, whose minimal-fact count is fixed by stability conditions (identified as seven in separate programme work).

The result is conditional: a reader who rejects any one of the five principles can escape the conclusion. But each such rejection has a definite cost — it requires giving up something most physicists would normally take for granted. The paper tabulates the costs explicitly.

We also prove the reverse direction: no admissible theory exists *outside* the VERSF framework. Taken together, the forward and reverse results establish a biconditional: within the stated principles, the admissible physical theory and VERSF are the same thing, up to relabellings that change no physical content.

The claim is not that VERSF is the only conceivable physical theory. The claim is that within the class of theories honouring the five principles, VERSF is the unique realization — not one option among many, but what the principles themselves force into being.

The technical argument follows.

Abstract

We prove a reconstruction theorem for the VERSF framework. Starting from five meta-principles — facthood, comparability, self-consistency, completeness, and record-persistence, each established in §2 as a necessary condition for a framework to count as a candidate fundamental physical theory — we derive a chain of results culminating in the claim that any physical theory satisfying these principles reconstructs the VERSF invariants: the carrier sector, the logarithmic marginal ordering coupling, the dimensionless topologically quantized closure law, and a simplicial substrate whose minimal-fact count K is a unique combinatorial fixed point determined by stability conditions (i)–(v); programme work evaluates this fixed point as K = 7.

The argument proceeds in two stages. From the meta-principles we derive six admissibility theorems as internal consequences, including compositional closure (derived from M2 + M3, superseding earlier auxiliary premises). From these we reconstruct each of the four VERSF sectors, distinguishing positive uniqueness from exclusion uniqueness and closing the gap between them at each stage. We establish substrate rigidity: no non-trivial continuous deformation of the substrate preserves the compatibility system. We characterise the K-preserving equivalence group $\mathcal{G}$_K explicitly, showing that its substrate factor is finite.

Section 13 addresses six anticipated referee objections, each cross-referenced to the main text, showing that every structural challenge either reduces to rejecting a meta-principle or requires producing an admissible structure outside $\mathcal{G}$_K.

Section 14 establishes a no-go theorem for non-VERSF theories by exhaustion of deviation classes, giving the contrapositive direction of the reconstruction. Together with Theorem 11.4 this expresses the biconditional

Admissibility ⇔ VERSF (within the meta-principle class M1–M5, up to $\mathcal{G}$_K)

The main theorem states: any theory satisfying the meta-principles coincides with the VERSF realization up to the action of $\mathcal{G}$_K. The conditionality lies exclusively at the meta-level; no auxiliary premises and no imported inputs remain in the argument chain. The specific

combinatorial value K = 7 is supplied by programme work as a downstream evaluation, not as a premise.

Keywords: fact admissibility; physical uniqueness; reconstruction theorem; carrier field; emergent time; topological closure; simplicial substrate; VERSF

Table of Contents

---

1. Introduction

The central question addressed by this paper is whether the VERSF framework is one viable physical theory among many, or whether it is forced by minimal structural requirements on what counts as a physical theory.

We answer: forced, up to representational equivalence, given five meta-principles that any serious physical theory would normally be expected to satisfy. The VERSF invariants are not selected from a menu of alternatives. They are reconstructed: they emerge as the unique structure consistent with the meta-principles up to an explicitly characterised and finite-dimensional-in-substrate representational equivalence group.

The meta-principles are stated in ontology-neutral language. They make no reference to carriers, fields, time, space, closure surfaces, or substrates. Any category appearing in the final reconstruction is derived, not presupposed.

The reconstruction proceeds in two stages. First, the meta-principles force six admissibility theorems: persistence, well-posed realization, irreversibility, ordering consistency, internal closure, and compositional closure. Second, these admissibility theorems force the four sectors of VERSF — carrier, coupling, closure, substrate — and an explicit algebraic compatibility system K. A separate rigidity theorem establishes that no continuous deformation of the substrate preserves K, securing the discreteness of the admissible structure space.

1.1 What this paper claims

We distinguish two kinds of uniqueness throughout:

- **Positive uniqueness.** A specific structure is shown to satisfy a set of conditions.
- **Exclusion uniqueness.** No other structure satisfies those conditions.

Positive uniqueness does not automatically entail exclusion uniqueness. The main theorem establishes both, using the meta-principles for positive uniqueness and the rigidity theorem (§10) plus the minimality theorem (§9) for exclusion.

The specific combinatorial value $K = 7$ is not a premise of the theorem. The theorem reconstructs a simplicial substrate with minimal-fact count K uniquely determined by combinatorial stability conditions; programme work supplies the evaluation $K = 7$ as a downstream result. A reader who disputes the programme identification of $K = 7$ must offer a different combinatorial fixed point; they cannot escape the reconstruction itself.

1.2 Epistemic status labelling

- **[Theorem]** — fully proved within the paper.
- **[Lemma]** — structure-reconstruction result with explicit positive-uniqueness proof; exclusion completed via rigidity and minimality.
- **[Assumption]** — stated explicitly, not derived.
- **[Consequence]** — derived from prior claims in the same chain.
- **[Programme evaluation]** — a downstream numerical or combinatorial value supplied by programme work; not a premise of the reconstruction.

---

2. Why M1–M5 Are Not Optional: Constitutive Conditions for Physical Theory

Earlier versions of this paper presented M1–M5 as assumptions. A more precise framing is available: M1–M5 are not optional premises. They are consequences of what it means for something to be a candidate fundamental physical theory.

We establish this in two steps. First (§2.1), we state a constitutive criterion — the minimal content-requirements any framework must possess to count as a candidate fundamental physical theory. Then (§§2.2–2.6) we derive M1–M5 as theorems from this criterion. Corollary 2.6 (§2.7) states that any candidate fundamental physical theory satisfies M1–M5 — not as a matter of choice, but as a condition of being physics in the fundamental sense.

2.1 The constitutive criterion

[Definition 2.1] *Candidate fundamental physical theory.* A framework counts as a *candidate fundamental physical theory* if and only if it satisfies all of:

- **(C1) Determinate content.** The framework supports at least some statements about the world whose truth values are not arbitrary.
- **(C2) Empirical discrimination.** The framework's determinate content admits mutual comparison — facts can be discriminated from one another.
- **(C3) Non-contradictory realization of determinate content.** The framework's physical realizations of determinate content do not assign contradictory truth values to the same physical claim.
- **(C4) Internally specified observables.** The framework's essential observable content is fixed by its own internal specification, not supplied by externally tuned parameters.
- **(C5) Temporal retention.** The framework supports the persistence of records of determinate facts sufficient for comparison across events separated in its parameter.

These five conditions are not technical physics principles. They are minimal content-requirements that any framework must satisfy to count as physics — as opposed to a piece of mathematics, a purely logical exercise, or a narrative without physical content.

Remark on scope. The criterion concerns candidate *fundamental* theories — frameworks offered as self-standing accounts of physical structure. It does not rule out effective theories, phenomenological models, or instrumentalist frameworks; these may serve important scientific purposes without individually meeting C1–C5. The question of whether such frameworks reduce to a fundamental theory satisfying C1–C5 is separate from the constitutive criterion itself.

2.2 Necessity of facthood

[Theorem 2.1] Necessity of Facthood. Any framework lacking determinate distinctions lacks physical content and is not a candidate fundamental physical theory.

Proof. By (C1), determinate content is constitutive. A framework in which no statement's truth value is non-arbitrary supports no non-arbitrary discrimination between states of the world. The framework may describe possibilities, impose symbolic structure, or specify mathematical relations, but nothing in it holds or fails to hold physically. Such a framework fails (C1). ■

[Definition M1] *Facthood* is the condition established by Theorem 2.1: a candidate fundamental physical theory supports at least some determinate distinctions.

2.3 Necessity of comparability

[Theorem 2.2] Necessity of Comparability. Any framework whose determinate distinctions are not mutually comparable lacks empirical and relational content and is not a candidate fundamental physical theory.

Proof. By (C2), empirical discrimination is constitutive. A determinate distinction incomparable with every other determinate distinction is isolated: no other fact can be distinguished from it, contrasted with it, or measured against it. Such a distinction supplies no physical content. A framework whose determinate distinctions are all of this isolated type fails (C2). ■

[Definition M2] *Comparability* is the condition established by Theorem 2.2: determinate distinctions in a candidate fundamental physical theory are mutually comparable.

2.4 Necessity of self-consistency for determinate physical content

[Theorem 2.3] Necessity of Self-Consistency. Any framework that assigns contradictory determinate content to the same physical claim fails to define well-posed physical realization and is not a candidate fundamental physical theory.

Proof. By (C3), physical realization of determinate content must be non-contradictory. Suppose a framework assigns both F and $\neg$F as determinate physical truths of the same claim. Any physical realization reading this assignment yields two incompatible readings of the same fact. The realization is not well-posed and fails (C3). ■

Remark on scope. Theorem 2.3 concerns the assignment of determinate content to physical claims. It does not forbid the use of paraconsistent or dialetheic logical systems in abstract contexts; it requires only that such logics, if used within a physical theory, not produce contradictory determinate content. A physical theory using non-classical logic in its meta-level reasoning while maintaining consistent determinate physical content is not excluded by this theorem.

[Definition M3] *Self-consistency* is the condition established by Theorem 2.3: a candidate fundamental physical theory's realizations of determinate content are non-contradictory.

2.5 Necessity of completeness for fundamental theories

[Theorem 2.4] Necessity of Completeness. Any candidate fundamental theory whose observable content depends on externally supplied essential parameters is not complete as a physical theory.

Proof. By (C4), a fundamental theory's essential observable content is fixed by its own internal specification. Suppose a candidate fundamental theory T requires externally supplied values for parameters essential to its observable content. Then the observable content is not fully specified by T; it is specified jointly by (T + external values). The joint specification is the true candidate fundamental theory; T alone is incomplete. Equivalently: if T claims fundamental status while leaving essential parameters externally supplied, T is not yet fundamental. ■

Remark on scope. Theorem 2.4 applies to theories presenting themselves as fundamental. Effective field theories, phenomenological models, and intermediate-scale descriptions are not excluded by this theorem; they are simply not in the class to which it applies. An effective theory whose parameters are fixed by a deeper theory satisfying (C4) is perfectly admissible — though the fundamental status then belongs to the deeper theory.

[Definition M4] *Completeness* is the condition established by Theorem 2.4: a candidate fundamental physical theory's essential observables are fixed by its internal specification.

2.6 Necessity of record-persistence

[Theorem 2.5] Necessity of Record-Persistence. Any framework in which determinate distinctions leave no persistent records cannot support comparison across events and is not a candidate fundamental physical theory.

Proof. By (C5), the framework must support record-persistence sufficient for cross-event comparison. Suppose determinate distinctions leave no persistent records: a fact F determinate at parameter t retains no trace at $t' > t$. Then any comparison intended to discriminate F from a fact G at t' has no t-content to compare against. Cross-event comparison fails, violating (C2); persistence itself fails, violating (C5). ■

[Definition M5] *Record-persistence* is the condition established by Theorem 2.5: determinate distinctions in a candidate fundamental physical theory persist sufficiently to permit cross-event comparison.

2.7 Joint corollary and distilled statement

[Corollary 2.6] *Any candidate fundamental physical theory satisfies M1–M5.*

Proof. By Theorems 2.1–2.5, each of M1, M2, M3, M4, M5 is a necessary condition for a framework to be a candidate fundamental physical theory. Hence any candidate fundamental physical theory satisfies all five. ■

Distilled: *a framework that cannot produce stable, comparable, non-contradictory, internally-specified records of what happened is not a candidate fundamental physical theory.*

The remainder of the paper shows that the framework that can — within the class of candidate fundamental physical theories — is the VERSF realization, up to the action of the K-preserving equivalence group $\mathcal{G}$_K.

2.8 On the minimality of M1–M5

Several additional candidates were considered for inclusion and rejected as either derivable from M1–M5 or non-constitutive of being physics:

- *Compositional closure.* Derived as Theorem 3.6 from M2 + M3.
- *Empirical adequacy.* Follows from M1 + M2 + M4.
- *Tractability* (observables computable in principle). Not constitutive: a framework whose observables exist but are uncomputable remains physics.
- *Some-world existence.* A precondition for the enterprise, not a content-requirement within it.

The derivation that follows assumes exactly M1–M5 as established by Theorems 2.1–2.5. No auxiliary premises appear in the argument chain.

---

3. Admissibility Conditions as Theorems

3.1 Persistence

[Theorem 3.1] Persistence. Any theory satisfying M1, M2, and M5 supports persistent distinguishable content over nonzero extent of the comparison parameter.

Proof. By M1, some distinctions hold determinately. By M2, they are comparable. By M5, they remain available for comparison at $t' > t$ within the comparison domain. Comparison across nonzero extent requires persistence over that extent. ■

3.2 Well-posed realization

[Theorem 3.2] Well-posed realization. Any theory satisfying M2 and M3 supports a well-posed realization structure: given source data and admissible boundary data, the realization of persistent content is unique.

Proof. Suppose the same source and boundary data produced two distinct realizations R, R′ of the same content. By M2, R and R′ are comparable. Comparison yields two distinct determinate assignments to a single comparable distinction. By M3, this is a contradiction. Hence uniqueness. ■

3.3 Irreversibility

[Theorem 3.3] Irreversibility. Any theory satisfying M3 and M5 supports irreversible record formation: once a distinction is committed to a record at t, no admissible subsequent record at $t' > t$ asserts the negation of that distinction's original determinate content.

Proof. Let R(F, t) be the record of fact F at t. By M5, R(F, t) persists at all $t' > t$. Suppose an admissible record R(¬F, t) exists at t′ asserting the negation of F *at the same parameter-value t*. Both records are then simultaneously present for comparison, yielding determinate content both F and ¬F at t. By M3, this is a contradiction. ■

The theorem does not forbid change. A later record R(G, t′) with $G \neq F$ is admissible; what is forbidden is retroactive revision of earlier content.

3.4 Ordering consistency

[Theorem 3.4] Ordering consistency. Any theory satisfying M2 and M3 supports a partial order on committed records, consistent up to admissible transformations: orderings in frames F and F′ related by admissible transformation $T: F \to F'$ transform by T's induced action without producing contradictory determinate content.

Proof. Suppose two admissible frames assigned orderings O, O′ whose implied determinate content contradicted. By M2, the contents are comparable; by M3, they cannot contradict. Hence admissible frames assign orderings related by a consistent transformation rule — a partial order up to admissible transformation. ■

Remark. This is compatible with special-relativistic simultaneity: different frames assign different orderings to spacelike-separated events, but Lorentz transformations relate them without contradiction. The theorem does not establish scale-covariance; that is developed separately in §5 under the marginality requirement.

3.5 Internal closure

[Theorem 3.5] Internal closure. Any theory satisfying M4 specifies its global observables in a form determined by the theory's internal structure, with no appeal to externally supplied continuous parameters.

Proof. Suppose a global observable required an external continuous parameter α. Then the observable is not specified by the theory's internal structure; it is specified jointly by the theory and the external choice of α. This violates M4. Hence every parameter appearing in a global observable specification is internally fixed. ■

Dimensionlessness and topological quantization are not conclusions of this theorem; they are derived in §6 from internal closure combined with the topology of Σ.

3.6 Compositional closure

[Theorem 3.6] Compositional closure. Any theory satisfying M2 and M3 supports compositional closure of realization: if D_1 and D_2 are determinate distinctions whose conjunction $D_1 \oplus D_2$ is also determinate, then the realization $R(D_1 \oplus D_2)$ equals the composition $R(D_1) \boxplus R(D_2)$ of their individual realizations wherever both are defined.

Proof. By M1, if D_1 and D_2 are each determinate, their conjunction $D_1 \oplus D_2$ ("both D_1 and D_2 hold") is determinate: its truth value is the logical conjunction of theirs.

By M2, D_1, D_2, and $D_1 \oplus D_2$ are each comparable. Their realizations $R(D_1)$, $R(D_2)$, $R(D_1 \oplus D_2)$ represent their determinate content.

Consider the composed realization $R(D_1) \boxplus R(D_2)$, wherever defined. This object represents the content "D_1 is realized jointly with D_2" — exactly the content of the compound $R(D_1 \oplus D_2)$.

By M2, these two representations of the same content are comparable. If they differed, comparison would yield two distinct determinate values for the same determinate content, violating M3.

Hence $R(D_1 \oplus D_2) = R(D_1) \boxplus R(D_2)$ wherever both are defined. ■

Corollary 3.6.1 (Linearity at leading order). In any perturbation expansion of R about a fixed base point, the composition $\boxplus$ reduces at leading order to addition on the tangent space, giving

$$R(D_1 \oplus D_2) = R(D_1) + R(D_2) + O(\text{higher order})$$

at leading order. ■

---

4. Carrier Reconstruction

[Lemma 4.1] Carrier structure. Any theory satisfying M1–M5 admits a carrier sector $\phi[\rho, B]$ representing persistent content ρ subject to boundary data B, with the carrier taking the form of the Green's function of a self-adjoint linear *elliptic* operator on the comparison domain.

Argument.

Persistent content $\rho: \Omega \to V$ is supported by Theorem 3.1, with $\nabla\rho \neq 0$ on some nonempty subset of Ω. The realization of ρ given source and boundary data is unique by Theorem 3.2.

(i) Linearity at leading order. Direct from Corollary 3.6.1.

(ii) Locality. By Theorem 3.4, admissible ordering is a partial order up to admissible transformation. A non-local realization — source data at x influencing realization at x′ for x, x′ spacelike-separated in some admissible frame — permits ordering transformations that violate the partial-order structure. Hence the realization is local: source data at x influences the realization at x′ only along causally admissible paths.

(iii) Invertibility. By Theorem 3.2, the realization is unique given source and boundary data. Invertibility of the operator follows from uniqueness.

(iv) Self-adjointness. The comparison relation of M2 is reciprocal: if D_1 is comparable to D_2, then D_2 is comparable to D_1. In the realization, this forces the Green's function to be symmetric: $G(x, x') = G(x', x)$. Symmetry of the Green's function is equivalent to self-adjointness of the underlying operator.

(v) Ellipticity (not hyperbolic, not parabolic). Given (i)–(iv), the operator is linear, local, invertible, and self-adjoint. There remain three classes of such operators: elliptic, hyperbolic, and parabolic. We show only elliptic is admissible.

Hyperbolic operators have Green's functions that are retarded or advanced: they require an *a priori* directed time parameter to define causality. In VERSF, time is emergent from commitment events (Theorem 3.3), not a primitive of the framework. A hyperbolic realization would require external time structure that the theory does not itself supply, violating M4.

More fundamentally: hyperbolic operators encode propagation relative to a pre-existing causal structure. In the present framework, causal ordering is reconstructed from commitment events (Theorem 3.3), not presupposed. A hyperbolic operator would therefore introduce an independent causal structure not reducible to the reconstructed ordering — a second, parallel causal structure running alongside the emergent one. By Theorem 9.1 (minimality), no such redundant structure is admissible.

Furthermore, retarded and advanced Green's functions are not symmetric: $G_{ret}(x, x') = G_{adv}(x', x) \neq G_{ret}(x', x)$. This violates self-adjointness (iv).

Parabolic operators (diffusion-type) have irreversible Green's functions that define a preferred flow direction: they similarly require external time structure.

Elliptic operators have symmetric Green's functions, support no preferred direction, and are compatible with emergent causal ordering (the ordering structure is carried by commitment events, not by the carrier operator itself).

Hence the admissible operator is elliptic.

Positive uniqueness. Given (i)–(v), the realization is the Green's function of a self-adjoint linear elliptic operator on Ω with boundary conditions B. This is the carrier $\phi[\rho, B]$.

Exclusion uniqueness. (i)–(v) exclude non-Green's-function realizations within the local/linear/invertible/self-adjoint/non-directional class. ■

---

5. Coupling Reconstruction

[Lemma 5.1] Marginal logarithmic coupling. Any theory satisfying M1–M5 admits a leading ordering–entropy coupling of the form

$$\mathcal{L}\_\text{int} \sim \lambda \cdot \phi^2 \cdot \ln \sigma$$

up to redefinition of λ.

Argument.

The leading coupling has schematic form $\mathcal{L}\_\text{int} = \lambda \cdot \phi^n \cdot f(\sigma)$.

Scalar leading order. By Theorem 3.4, ordering is a partial-order structure. The lowest scalar in ϕ compatible with the $\pm$ reflection symmetry $\phi \to -\phi$ is ϕ^2. Odd powers are forbidden by reflection symmetry. Higher even powers are subleading. Hence $n = 2$.

Marginality from internal closure and minimality. Consider a non-marginal coupling with anomalous dimension $\alpha \neq 0$. Under renormalization flow from scale μ_0 to scale μ, the coupling strength runs as

$$\lambda(\mu) = \lambda(\mu_0) \cdot (\mu/\mu_0)^\wedge\alpha$$

For $\alpha > 0$, λ diverges at high μ; for $\alpha < 0$, λ vanishes. In either case, physical observables computed at scale μ depend on μ through the running coupling. Fixing the physics at μ requires specifying the reference scale μ_0.

Two cases exhaust the specification of μ_0:

- *Case A: μ_0 is externally supplied.* This directly violates Theorem 3.5 (internal closure): an external continuous parameter fixes the theory's observable content.
- *Case B: μ_0 is dynamically generated by internal structure.* Dynamical generation requires additional internal structure not already present in the four sectors — a new field, a new condensate, or a new symmetry-breaking mechanism. By Theorem 9.1 (minimality), the theory contains no structure beyond what is internally required. Introducing new structure whose sole purpose is to fix μ_0 is precisely the redundancy that Theorem 9.1 excludes,

unless the structure is independently forced by another admissibility condition. No such independent forcing exists in the minimal reconstruction.

Both cases are excluded. Hence $\alpha = 0$: the coupling is marginal.

Selection of the logarithm from marginality. A marginal coupling satisfies, under rescaling $\sigma \to \mu\, \sigma$, the functional equation

$$f(\mu\, \sigma) = f(\sigma) + g(\mu)$$

for some function g. Power-law behaviour $f(\sigma) = \sigma^\alpha$ yields $f(\mu\sigma) = \mu^\alpha f(\sigma)$, which is *scale-covariant* ($\alpha \neq 0$) but not of the marginal form above. The unique continuous solution of the marginal functional equation is

$$f(\sigma) = \ln \sigma$$

up to a multiplicative constant absorbed into λ.

Positive uniqueness. Given reflection symmetry, scalar-leading-order, and marginality, the coupling $\mathcal{L}_{int} \sim \lambda \cdot \phi^2 \cdot \ln \sigma$ is uniquely determined.

Exclusion uniqueness. Non-marginal couplings are excluded by internal closure plus minimality (Theorem 9.1). No auxiliary premise is required.

The exclusion of dynamically generated scales follows because any such scale would either require an independent sector (violating minimality) or introduce additional degrees of freedom not fixed by the admissibility conditions. ■

---

6. Closure Reconstruction

[Lemma 6.1] Topological closure. Any theory satisfying M1–M5, whose global closure surface Σ has nontrivial 1-cycle topology, admits a global closure law

$$N_\Sigma \cdot \theta(\alpha) = 2\pi \cdot \chi, \quad \chi \in \mathbb{Z}$$

Argument.

By Theorem 3.5, the closure law is expressed in internally-fixed form: all appearing parameters are specified by internal structure.

Let Σ be the global closure surface (supplied by §7 as the global boundary structure of the substrate). Four conditions apply:

(a) Internal specification. By Theorem 3.5, no external continuous parameter appears. Only internally-fixed data — cell count N_Σ, per-cell holonomy θ(α), global winding χ — enter the law.

(b) Dimensionlessness. Derived, not assumed. The closure law is an identity among global invariants of Σ; such invariants are counting or winding numbers, which are dimensionless. The per-cell holonomy θ(α) is a phase, dimensionless by construction.

(c) Topological quantization. Derived from the topology of Σ plus the structure of the holonomy. If Σ contains a nontrivial 1-cycle, parallel transport around that cycle returns the transported object to itself; the accumulated phase is therefore a multiple of 2π. Since total holonomy accumulates additively across N_Σ cells, quantization takes the form $N_\Sigma \cdot \theta(\alpha) = 2\pi \cdot \chi$, $\chi \in \mathbb{Z}$.

(d) Exactness. The law is an identity because $\chi \in \mathbb{Z}$ admits no continuous variation.

Positive uniqueness. Given (a)–(d), the form $N_\Sigma \cdot \theta(\alpha) = 2\pi \cdot \chi$ is forced.

Exclusion uniqueness. Requires nontrivial topology of Σ, supplied by Lemma 7.1. ■

---

7. Substrate Reconstruction

[Lemma 7.1] Simplicial substrate with combinatorial minimal-fact count. Any theory satisfying M1–M5 reconstructs a simplicial substrate $\mathcal{P}$ whose minimal-fact count K is uniquely determined by the combinatorial stability conditions (i)–(v) below. Programme work [K = 7 convergence paper] identifies K = 7 as the unique solution.

Part A: simpliciality (within this paper).

By Theorem 3.3, the substrate supports irreversible commitment. By M1 + M4, the substrate has finite per-unit capacity: an infinite-capacity unit either reduces to finite-capacity composition (redundant by Theorem 9.1) or requires infinite external specification (violates M4). By Theorem 3.2, the substrate composes consistently.

Claim. Consistent composition forces simpliciality.

Proof. Consider a substrate with a composition cycle: a sequence starting and ending at the same unit with a composite value assigned along the cycle. If the endpoint value differs from the starting value, comparison along the cycle yields distinct determinate contents for the same composite fact, violating M3. If the endpoint equals the starting value, the cycle contributes no net composition — it is redundant, excluded by Theorem 9.1. Hence the substrate has no nontrivial composition cycles: it is acyclic. An acyclic composition structure strictly ordered by dimension is simplicial. ■

Part B: combinatorial determination of K.

The minimal-fact count K of an admissible simplicial substrate is constrained by five conditions:

(i) support of irreversible commitment (Theorem 3.3); (ii) persistence under M5 — sufficient vertex count to support distinguishable persistent content; (iii) consistent composition — acyclic simplicial structure from Part A; (iv) irreducible action of the minimal-fact projective group on the composition space — required for the substrate to be *minimal* (Theorem 9.1): a reducible action implies the substrate decomposes into disconnected sub-substrates, violating minimality by presenting as multiple independent theories rather than one; (v) closure of the compatibility system K = 0 (§8).

[Programme evaluation] Evaluation of K. Programme work in the K = 7 convergence paper shows that conditions (i)–(v) admit K = 7 as the unique solution, identified with the irreducibility of the PGL(3,2) action on the 7-point Fano plane and its lift to the simplicial 2-complex. K < 7 fails condition (ii) or (iii); K > 7 fails minimality via redundant composition paths; non-integer or non-projective configurations fail (iv).

Positive uniqueness within this paper. Part A (simpliciality) is proved. Part B establishes that the theorem fixes the substrate to a unique combinatorial fixed point K determined by (i)–(v); programme work evaluates this fixed point as K = 7.

Epistemic status of the programme evaluation. The programme identification of K = 7 is a *computation* on conditions (i)–(v) established within this paper, not an imported structural premise. A reader who disputes K = 7 must either produce a different solution to (i)–(v) or identify an error in the programme's combinatorial argument. They cannot escape the reconstruction itself by rejecting K = 7; at most they shift the evaluation to a different value while preserving the framework. ■

---

8. Algebraic Compatibility

8.1 The compatibility system

[Definition 8.1] Compatibility system K. A theory specifies a tuple $(\phi, \sigma, \theta, \mathcal{P})$. The compatibility system K is

$K(\phi, \sigma, \theta, \mathcal{P}) = (K_1, K_2, K_3, K_4) = 0$

where:

- **K_1:** $\sigma - \sigma[\mathcal{P}] = 0$
- **K_2:** $\theta - \theta[\mathcal{P}] = 0$
- **K_3:** $N\_\Sigma - |\partial\mathcal{P}| = 0$
- **K_4:** $\lambda - \lambda[\mathcal{P}, K] = 0$

The functionals $\hat{\sigma}$, $\hat{\theta}$, $\hat{\chi}$ are well-defined on the substrate: $\hat{\sigma}$ from the commitment-event bath spectral density, $\hat{\theta}$ from the fold-interface holonomy, $\hat{\chi}$ from the normalization condition of the Fact-Momentum paper.

8.2 Non-triviality

[Theorem 8.2] K is non-trivial. There exist tuples (ϕ, σ, θ, $\mathcal{P}$) for which each component individually satisfies its reconstruction lemma, which are not related by any $\mathcal{G}_K$-action (defined in §10), but which violate K.

Proof. Consider two simplicial substrates $\mathcal{P}_1$ and $\mathcal{P}_2$ that:

- have identical minimal-fact count K;
- have identical vertex, edge, and face counts;
- have isomorphic composition groups;

but

- have *distinct* fold-interface holonomies $\hat{\theta}[\mathcal{P}_1] \neq \hat{\theta}[\mathcal{P}_2]$ due to inequivalent embedding of the holonomy generators in the substrate's 2-cycle structure.

Such pairs exist on K-vertex simplicial 2-complexes: multiple inequivalent assignments of orientation data to 2-cycles produce different $\hat{\theta}$ while preserving all lower combinatorial invariants.

Form the tuple (ϕ, $\sigma[\mathcal{P}_1]$, $\theta[\mathcal{P}_2]$, $\mathcal{P}_1$). Each component is sector-wise admissible. K_2 fails: $\theta[\mathcal{P}_2] \neq \hat{\theta}[\mathcal{P}_1]$. The failure is not absorbed into $\mathcal{G}_K$, because $\mathcal{P}_1$ and $\mathcal{P}_2$ are not isomorphic substrates — only agreeing on lower invariants — so no $\mathcal{G}_K$ transformation relates them.

Hence K is non-trivial beyond representational freedom. ■

8.3 Compatibility collapse

[Theorem 8.3] Compatibility collapse. Any theory satisfying M1–M5 satisfies K = 0.

Proof. Suppose $K_i \neq 0$ for some i. Two sectors impose distinct specifications on a shared parameter.

Case 1: simultaneous use. The distinct specifications are used together in some derivation of observable content. The derivation then assigns two determinate values to the observable, violating M3.

Case 2: non-simultaneous use. One specification is redundant. If the specifications agree, the redundancy is excluded by Theorem 9.1. If they disagree, Case 1 applies on any subsequent derivation that uses them together.

In either case, M3 + Theorem 9.1 exclude $K_i \neq 0$. ■

---

9. Necessity of Minimality, Stability, Observational Closure

9.1 Minimality

[Theorem 9.1] Minimality. Any theory satisfying M3 and M4 has no redundant internally-specified degrees of freedom.

Proof. Suppose d is specified twice by internal conditions C_1, C_2. If C_1 and C_2 disagree, M3 is violated. If they agree, one is derivable from the other and is not an independent specification. ■

9.2 Stability

[Theorem 9.2] Stability (abstract form). Any theory satisfying M1 and M5 under realistic perturbations supports realizations stable under the admissible perturbation class.

Proof. M1 requires determinate facthood; M5 requires persistence. A realization unstable under admissible perturbations loses determinate content over the parameter range where M5 demands persistence. Hence admissible realizations are stable under admissible perturbations.

The specific norm — L^2 on carrier-sector field-configuration space, linearized about equilibrium substrate configuration — is a *downstream consequence* of §§4 and 7, not a presupposition. ■

9.3 Observational closure

[Theorem 9.3] Observational closure. Any theory satisfying M4 satisfies observational closure.

Proof. M4 and observational closure are the same condition. ■

---

10. The K-Preserving Equivalence Group and Substrate Rigidity

10.1 Definition of $\mathcal{G}$_K

[Definition 10.1] K-preserving transformations. Let $\mathscr{E}$ be the space of admissible tuples $(\phi, \sigma, \theta, \mathcal{P})$ satisfying $K = 0$. A transformation $T: \mathscr{E} \to \mathscr{E}$ is *K-preserving* if:

- T acts on each sector by one of: a carrier redefinition $\phi \mapsto \Lambda \cdot \phi + c$ ($\Lambda \in \mathbb{R}^*$, c constant), a gauge transformation $\theta \mapsto \theta + d\Xi$, or a simplicial automorphism $\mathcal{P} \mapsto \pi(\mathcal{P})$;

- T acts compatibly across sectors: induced transformations preserve each $K_i = 0$.

The set of K-preserving transformations forms a group $\mathcal{G}_K$.

10.2 Structure of $\mathcal{G}_K$

[Theorem 10.2] Structure of $\mathcal{G}_K$.

$\mathcal{G}_K \cong \mathcal{G}_{field} \times \mathcal{G}_{gauge} \times \mathrm{Aut}(\mathcal{P})$

where:

- $\mathcal{G}_{field} = \mathbb{R}^*$ (carrier scalings);
- $\mathcal{G}_{gauge}$ = exact 1-forms on the substrate;
- $\mathrm{Aut}(\mathcal{P})$ = simplicial automorphism group of the substrate; for K = 7 on the Fano plane, $\mathrm{Aut}(\mathcal{P}) = \mathrm{PGL}(3,2)$, of order 168.

Proof. Sector-wise admissible redefinitions are carrier scalings, gauge transformations, and substrate relabellings. Compatibility across sectors forces the actions to commute, as each acts on a distinct structural feature (amplitude, phase, labelling). Hence the decomposition. ■

Crucial consequence. $\mathcal{G}_K$ is finite-dimensional in the substrate factor (order 168 for K = 7). It does not contain continuous deformations of the substrate geometry. This characterisation makes the uniqueness claim falsifiable: a rival theory either lies in the $\mathcal{G}_K$-orbit of VERSF, or violates K.

10.3 Substrate rigidity

[Theorem 10.3] Substrate rigidity. Let $\mathcal{P}(\varepsilon)$ be a one-parameter family of simplicial substrates with $\mathcal{P}(0) = \mathcal{P}$, varying continuously in ε. If the compatibility system $K(\varepsilon) = 0$ holds for all ε in some open interval around 0, then the family $\mathcal{P}(\varepsilon)$ is trivial: it consists of $\mathcal{P}$ itself up to $\mathcal{G}_K$-action.

Proof. Three cases exhaust the possibilities for a continuous family of K = 0 substrates.

Case A: all $\mathcal{P}(\varepsilon)$ are mutually simplicially isomorphic. Then $\mathcal{P}(\varepsilon) = \pi(\varepsilon) \cdot \mathcal{P}(0)$ for $\pi(\varepsilon) \in \mathrm{Aut}(\mathcal{P})$. But $\mathrm{Aut}(\mathcal{P})$ is discrete (for the finite simplicial complex characterising the admissible substrate; in particular, $|\mathrm{Aut}(\mathcal{P})| = 168$ for K = 7). A continuous path in a discrete group is constant: $\pi(\varepsilon) = \pi(0)$ for all ε. Hence $\mathcal{P}(\varepsilon) = \mathcal{P}(0)$ up to a fixed automorphism — the family is trivial.

Case B: the observables $\sigma[\mathcal{P}(\varepsilon)]$, $\theta[\mathcal{P}(\varepsilon)]$, $\lambda[\mathcal{P}(\varepsilon)]$ are all independent of ε. Then ε parametrizes physically indistinguishable substrates, meaning ε is a degree of freedom not fixed by the theory's observables. Either ε is fixed by an external choice (violates M4) or ε is redundant with existing internal specification (violates Theorem 9.1). Both are excluded.

Case C: at least one observable σ̂, θ̂, or λ̂ depends non-trivially on ε. Then $K(\varepsilon) = 0$ requires $\sigma(\varepsilon) = \hat{\sigma}[\mathcal{P}(\varepsilon)]$, $\theta(\varepsilon) = \hat{\theta}[\mathcal{P}(\varepsilon)]$, $\lambda(\varepsilon) = \hat{\lambda}[\mathcal{P}(\varepsilon)]$ to track $\mathcal{P}(\varepsilon)$ continuously. But ε is not fixed by any internal admissibility condition — the reconstruction lemmas (§§4–7) fix a single admissible solution at each ε, not a family. If the family is genuinely non-trivial, some ε-values are selected over others only by an external choice, violating M4.

All three cases lead to contradiction or triviality. Hence no non-trivial continuous deformation of $\mathcal{P}$ preserves K. ■

---

11. Reconstruction Theorem

[Theorem 11.1] Reconstruction theorem. Let T be any theory satisfying M1–M5. Then T reconstructs:

1. a carrier sector $\phi[\rho, B]$ as the Green's function of a self-adjoint linear elliptic operator (§4);
2. a leading ordering–entropy coupling $\mathcal{L}_{int} \sim \lambda \cdot \phi^2 \cdot \ln \sigma$, with marginality forced by internal closure and minimality (§5);
3. a global closure law $N_\Sigma \cdot \theta(\alpha) = 2\pi \cdot \chi$, $\chi \in \mathbb{Z}$ (§6);
4. a simplicial substrate $\mathcal{P}$ with minimal-fact count K uniquely determined by combinatorial stability conditions (§7; programme work evaluates K = 7);
5. the algebraic compatibility system K = 0 (§8);
6. the substrate rigidity property: no continuous deformation of $\mathcal{P}$ preserves K (§10).

The reconstruction is unique up to the action of the K-preserving equivalence group $\mathcal{G}_K$.

Proof. By Theorems 3.1–3.6, M1–M5 entail the admissibility conditions including compositional closure. By Lemmas 4.1, 5.1, 6.1, 7.1, each sector is reconstructed with explicit positive and exclusion uniqueness. By Theorem 8.3, the sectors satisfy K = 0. By Theorems 9.1–9.3, no redundant, unstable, or externally-specified degrees of freedom appear. By Theorems 10.2–10.3, residual representational freedom is exactly $\mathcal{G}_K$ and no continuous deformation extends beyond it. Hence the reconstruction is complete and unique up to $\mathcal{G}_K$-action. ■

Corollary 11.2. Any theory satisfying M1–M5 lies in the $\mathcal{G}_K$-orbit of the VERSF realization.

Corollary 11.3. No theory physically distinct from VERSF — in the sense of lying outside the $\mathcal{G}_K$-orbit — satisfies M1–M5.

[Theorem 11.4] Main uniqueness theorem. Given M1–M5, any physical theory reconstructs the VERSF invariants and coincides with the VERSF realization up to $\mathcal{G}_K$-action.

Proof. Immediate from Theorem 11.1 and Corollaries 11.2–11.3. ■

---

12. Discussion

12.1 The conditional structure

The result is conditional at exactly one point: the meta-principles M1–M5. No auxiliary premises appear in the argument chain. The specific combinatorial value K = 7 is a programme evaluation of the stability conditions established within the paper, not an imported structural premise.

M1–M5 are themselves not free assumptions: §2 establishes each as a necessary condition for a framework to count as a candidate fundamental physical theory. A reader rejecting the conclusion therefore faces a stronger cost than "declining an assumption": rejecting a meta-principle means accepting that the framework in question is not a candidate fundamental physical theory in the sense of §2.1. Each rejection is tabulated in §12.3 below.

12.2 What the theorem proves and does not prove

The theorem proves: given M1–M5, VERSF is reconstructed uniquely up to $\mathcal{G}$_K-action.

The theorem does not prove that M1–M5 are necessary truths about the world. Whether M1–M5 must hold is a question in the philosophy of physics, narrowed but not settled by the theorem.

The theorem does not derive the numerical value K = 7 within its own argument chain. It fixes the substrate to a unique combinatorial fixed point K determined by stability conditions (i)–(v); programme work evaluates this fixed point as K = 7. Disputing K = 7 requires either producing a different solution to the stability conditions or identifying an error in the programme's combinatorial argument — neither of which escapes the reconstruction framework itself.

12.3 Framework classification

The following classifications are interpretive and sometimes contested; caveats are given rather than asserting them as settled. The column "Condition at issue" identifies which of M1–M5 is not satisfied in some natural reading of the framework; under §2, failing a condition means the framework is not a candidate fundamental physical theory in the sense of §2.1 unless an alternative reading applies.

Framework	Condition at issue	Status
Many-worlds unitary QM	M1 (global reading)	Satisfies M1 under branch-local reading; classification depends on interpretive choice
Retrocausal interpretations	M3 (model-dependent)	Technical analysis per model required
String landscape (free moduli)	M4	Fails as fundamental theory unless internally closed completion exists
Standard QFT (tuned constants)	M4 for tuned constants	Not fundamental on its own; internal-completion question open

Framework	Condition at issue	Status
Causal sets (standard)	M4 (sprinkling density)	Not fundamental on its own; internally-closed variants possible
Primitive-time block universe	None typically	Theorem applies; realization is VERSF
Strict instrumentalism	M2 (internal-comparability reading)	May not claim fundamental status; scope dependent
Paraconsistent physics	M3 (for determinate physical content)	Theorem does not apply if determinate content is contradictory; non-classical meta-logic with consistent determinate content not excluded

12.4 Residual freedom, falsifiable

Residual representational freedom is exactly $\mathcal{G}$_K = ℝ* × (exact 1-forms) × Aut($\mathcal{P}$). The substrate factor has order 168 for K = 7, and is discrete for any admissible K. Substrate rigidity (Theorem 10.3) excludes continuous deformations beyond $\mathcal{G}$_K.

A critique of the uniqueness claim must therefore:

(a) challenge the characterisation of $\mathcal{G}$_K, arguing it is larger than stated; (b) challenge Theorem 10.3, arguing some continuous deformation survives; (c) challenge M1–M5.

Each is falsifiable. Option (a) requires exhibiting a K-preserving transformation outside the stated decomposition; Option (b) requires producing an explicit continuous family evading Cases A–C of the rigidity proof; Option (c) requires arguing against one of five stated meta-principles and accepting the cost.

12.5 Remaining openness

The argument chain has no auxiliary premises and no imported structural inputs. Three aspects remain open for further work:

1. *The programme evaluation of K = 7.* The paper establishes that K is fixed by combinatorial stability; the specific value K = 7 is supplied by the K = 7 convergence paper. Further work might internalize this combinatorial argument into the same chain, eliminating the programme reference entirely.
2. *The marginality derivation.* §5 derives marginality from internal closure + minimality using a case analysis (Cases A, B). A fully closed-form derivation of this case analysis as a single theorem of M1–M5 is a natural next step.
3. *The framework-level commitment to emergent time.* §4's exclusion of hyperbolic operators appeals to VERSF's choice not to treat time as primitive. This is a framework commitment visible in Theorem 3.3 (irreversibility from commitment events) but not an independent theorem of M1–M5 alone. A reader reconstructing a block-universe variant with primitive time could in principle admit hyperbolic carriers; such a reader would be

reconstructing a different framework, not VERSF. Making the emergent-time commitment an explicit framework axiom rather than an implicit background choice would further sharpen §4.

None of these threatens the conditional validity of the result, but each would sharpen it.

12.6 Final statement

If a physical theory satisfies facthood, comparability, self-consistency, completeness, and record-persistence (M1–M5), then it reconstructs the VERSF realization — carrier structure, marginal logarithmic coupling, dimensionless topologically quantized closure, simplicial substrate with combinatorially determined minimal-fact count — jointly constrained by K = 0 and rigid against continuous deformation, and unique up to the action of the explicit group $\mathcal{G}$_K whose substrate factor is finite.

The conditionality rests exclusively on M1–M5. No auxiliary premises. No imported structural inputs. The specific value K = 7 is a programme evaluation of conditions fixed within the paper.

Meta-principles ⇒ VERSF: unique reconstruction up to explicit finite representational freedom.

---

13. Anticipated Referee Objections and Responses

This section addresses six principal structural objections. Each response cross-references the relevant main-text theorem; responses are intended as pointers for the reader, not as independent proofs.

13.1 Objection A: Compositional closure is an unproven auxiliary assumption

Objection. The carrier reconstruction (Lemma 4.1) relies on compositional closure, which looks like an auxiliary premise rather than a consequence of the meta-principles. Therefore linearity at leading order — and with it the Green's-function structure — is not forced by M1–M5.

Response. Compositional closure is Theorem 3.6 of the main text, derived from M2 and M3 alone.

The argument, using the abstract composition ⊞ on realizations (addition is not pre-supposed; it emerges as the leading-order expansion of ⊞ — see Corollary 3.6.1):

Let D_1, D_2 be determinate distinctions. By M1, their conjunction $D_1 \oplus D_2$ is determinate. By M2, D_1, D_2, and $D_1 \oplus D_2$ are comparable. The compound realization $R(D_1 \oplus D_2)$ and the composed realization $R(D_1)$ ⊞ $R(D_2)$ each represent the determinate content "D_1 and D_2 jointly realized." If they differed, M2 would give two comparable-but-distinct determinate representations of the same content, which M3 forbids.

Hence $R(D_1 \oplus D_2) = R(D_1) \boxplus R(D_2)$ wherever both are defined. Linearity at leading order follows as Corollary 3.6.1: in perturbation expansion about a base point, $\boxplus$ reduces to addition on the tangent space.

The earlier auxiliary premise "M2*" is therefore eliminable. The argument chain depends only on M1–M5.

13.2 Objection B: Marginality is imposed, not derived

Objection. The logarithmic coupling in Lemma 5.1 depends on marginality under renormalization flow, which looks like a physical requirement imposed externally rather than derived from the meta-principles.

Response. Marginality is forced by M4 (internal closure) together with Theorem 9.1 (minimality). The argument is the two-case analysis of Lemma 5.1.

A non-marginal coupling with anomalous dimension $\alpha \neq 0$ introduces a running coupling $\lambda(\mu) = \lambda(\mu_0)\cdot(\mu/\mu_0)$^$\alpha$. Fixing the physics requires specifying the reference scale μ_0. Two cases exhaust this specification:

- *Case A: μ_0 externally supplied.* Violates M4 (Theorem 3.5, internal closure): an external continuous parameter fixes observables.
- *Case B: μ_0 dynamically generated.* Requires additional internal structure — a field, condensate, or symmetry-breaking mechanism — whose sole purpose is to fix μ_0. Theorem 9.1 excludes such redundant structure.

Marginality is therefore the unique fixed point consistent with M4 + Theorem 9.1.

Given marginality, the functional equation

$f(\mu\sigma) = f(\sigma) + g(\mu)$

is the *marginal* form — distinct from bare scale-covariance, which would also admit power laws $f(\sigma) = \sigma$^α. The continuous solution of the marginal functional equation is uniquely $f(\sigma) \propto \ln \sigma$.

Hence the logarithmic coupling is structurally forced, not imposed.

13.3 Objection C: The K = 7 result depends on external combinatorial input

Objection. The identification of K = 7 is imported from programme work. The uniqueness claim is therefore not closed within the paper.

Response. The theorem separates two claims that the objection conflates.

Claim 1 (within this paper): K is uniquely determined by five combinatorial stability conditions (i)–(v) of Lemma 7.1:

- support of irreversible commitment (Theorem 3.3);
- persistence under M5;
- consistent composition (simpliciality, Part A);
- irreducible minimal-fact projective action (Theorem 9.1);
- closure of the compatibility system K = 0 (§8).

That K is *uniquely determined* by (i)–(v) is established within the paper. The stability conditions are themselves theorems or lemmas of M1–M5.

Claim 2 (programme evaluation): The unique solution to (i)–(v) is K = 7, identified with PGL(3,2) irreducibility on the Fano plane. This is a combinatorial computation on conditions the paper fixes; it is supplied by the K = 7 convergence paper of the programme.

The structure of the reconstruction — simplicial substrate, minimal-fact count K, stability conditions (i)–(v) — is closed within this paper. Only the numerical evaluation K = 7 is supplied externally. A reader who disputes K = 7 must either produce a different solution to (i)–(v) or find an error in the programme's combinatorial argument. Either challenge is local to the evaluation; neither escapes the reconstruction framework.

The dependence on programme work is *evaluative*, not *structural*.

13.4 Objection D: Carrier reconstruction does not uniquely select elliptic operators

Objection. From linearity, locality, invertibility, and self-adjointness, it does not follow that the carrier operator must be elliptic. Hyperbolic or parabolic operators may satisfy the same conditions.

Response. See Lemma 4.1 step (v). The three cases are:

Hyperbolic operators. Their Green's functions are retarded or advanced, defined relative to a directed time parameter supplied a priori. VERSF treats time as emergent from commitment events (Theorem 3.3); supplying a primitive time parameter would be an external input, violating M4. More fundamentally, hyperbolic operators encode propagation relative to a pre-existing causal structure, which alongside the reconstructed ordering constitutes a second independent causal structure, excluded by Theorem 9.1 (minimality). Additionally, retarded/advanced Green's functions are not symmetric, violating self-adjointness (step iv).

Parabolic operators. Their Green's functions define a preferred flow direction (diffusion). The same primitive-time objection applies: the flow direction is not supplied by the theory's internal structure.

Elliptic operators. Their Green's functions are symmetric, support no preferred direction, and are compatible with emergent causal ordering (carried by commitment events, not by the carrier operator itself).

Hence, within the VERSF framework's commitment to emergent time, the carrier is elliptic.

Caveat. The exclusion of hyperbolic operators appeals to VERSF's framework-level commitment not to treat time as primitive. This commitment is visible in Theorem 3.3 but is not an independent theorem of M1–M5 alone. A reader reconstructing a block-universe variant with primitive time could admit hyperbolic carriers — but that reader would be reconstructing a different framework, not VERSF. Sharpening this commitment to an explicit framework axiom is noted as open work in §12.5.

13.5 Objection E: Continuous deformations of the substrate may evade uniqueness

Objection. The uniqueness claim assumes the substrate has no continuous deformation modes preserving admissibility. This requires proof.

Response. Theorem 10.3 (substrate rigidity) provides the proof via a three-case analysis on a continuous family $\mathcal{P}(\varepsilon)$ with $K(\varepsilon) = 0$.

Case A: all $\mathcal{P}(\varepsilon)$ mutually simplicially isomorphic. Then $\mathcal{P}(\varepsilon) = \pi(\varepsilon)\cdot\mathcal{P}(0)$ with $\pi(\varepsilon) \in \mathrm{Aut}(\mathcal{P})$. But $\mathrm{Aut}(\mathcal{P})$ is discrete (finite; order 168 for K = 7). A continuous path in a discrete group is constant: $\pi(\varepsilon) = \pi(0)$. The family is trivial.

Case B: observables $\hat{\sigma}$, $\hat{\theta}$, $\hat{\chi}$ all ε-independent. Then ε parametrizes physically indistinguishable substrates. ε is either externally fixed (violates M4) or redundant with existing internal specification (violates Theorem 9.1).

Case C: at least one observable ε-dependent. Then $K(\varepsilon) = 0$ requires the other sectors to track $\mathcal{P}(\varepsilon)$ continuously, but ε is not fixed by any internal admissibility condition. Non-trivial ε-dependence requires external selection of ε, violating M4.

All three cases lead to triviality or contradiction. Continuous deformations are absorbed into $\mathcal{G}_K$ (Case A) or excluded (Cases B, C). The key step in Case A — discreteness of $\mathrm{Aut}(\mathcal{P})$ — is load-bearing: it is what distinguishes $\mathcal{G}_K$ from a continuous gauge freedom that could secretly hide physical variation.

13.6 Objection F: Alternative theories may satisfy M1–M5 without being VERSF

Objection. There may exist non-VERSF theories satisfying M1–M5.

Response. The reconstruction theorem (Theorem 11.1) establishes that any theory satisfying M1–M5 reconstructs:

- carrier structure (§4),
- marginal logarithmic coupling (§5),
- topological closure law (§6),
- simplicial substrate with combinatorially determined K (§7),
- compatibility system K = 0 (§8, Theorem 8.3),
- substrate rigidity against continuous deformation (§10, Theorem 10.3).

By Theorem 10.2, residual representational freedom is exactly $\mathcal{G}$_K. By Theorem 10.3, no continuous deformation extends beyond $\mathcal{G}$_K.

Any candidate alternative theory therefore falls into one of three categories:

- *Inside the $\mathcal{G}$_K-orbit of VERSF.* Representationally equivalent; no physical distinction.
- *Outside the $\mathcal{G}$_K-orbit, satisfying M1–M5.* Excluded by Theorem 11.1 — no such theory exists.
- *Outside the $\mathcal{G}$_K-orbit, violating at least one of M1–M5.* Does not satisfy the hypotheses of the theorem; no claim is made.

Hence, within the class of theories satisfying M1–M5, there are no physically distinct alternatives to VERSF.

13.7 Summary of Objection Resolution

Objection	Resolution	Main-text reference
A. Compositional closure is auxiliary	Derived from M2 + M3	Theorem 3.6
B. Marginality is imposed	Forced by M4 + Theorem 9.1	Lemma 5.1
C. K = 7 is external	K structurally fixed by (i)–(v); K = 7 is programme evaluation	Lemma 7.1
D. Carrier not uniquely elliptic	Hyperbolic/parabolic need primitive time; framework commitment noted	Lemma 4.1(v); §12.5
E. Continuous deformation loophole	Three-case rigidity proof; discreteness of Aut($\mathcal{P}$) load-bearing	Theorem 10.3
F. Alternative theories may exist	Classified: $\mathcal{G}$_K-equivalent, excluded, or violate M1–M5	Theorem 11.1, 10.2, 10.3

13.8 Final position

The reconstruction withstands the six structural objections above. Any remaining critique must:

- reject one of M1–M5 (with the cost tabulated in §12.3); or
- exhibit an admissible structure outside $\mathcal{G}$_K evading Theorem 10.3 (with the falsifiability conditions of §12.4); or

- identify an error in the main-text theorems cross-referenced above.

All three routes are explicit. None is blocked by rhetorical framing; each is open to direct technical challenge.

Within its domain, the reconstruction is structurally closed.

---

14. No-Go Theorem for Non-VERSF Theories

This section states the contrapositive of Theorem 11.4 explicitly. Where Theorem 11.4 proves that M1–M5 force the VERSF realization, Theorem 14.1 proves that no admissible structure lies outside the $\mathcal{G}$_K-orbit. The two together express the biconditional

Admissibility ⇔ VERSF (within the meta-principle class M1–M5, up to $\mathcal{G}$_K)

This is the strongest claim the paper makes.

14.1 Statement

[Theorem 14.1] No-Go Theorem. Let T be any physical theory. Then exactly one of the following holds:

(a) T lies in the $\mathcal{G}$_K-orbit of the VERSF realization; (b) T violates at least one of M1–M5 or one of their derived consequences — compositional closure (Theorem 3.6), minimality (Theorem 9.1), compatibility (Theorem 8.3), or substrate rigidity (Theorem 10.3).

No third possibility exists.

14.2 Strategy of proof

By Theorem 11.1, any theory satisfying M1–M5 possesses four sectors — carrier (§4), coupling (§5), closure (§6), substrate (§7) — and a cross-sector compatibility system K (§8). A theory lacking one of these sectors fails the reconstruction hypotheses directly and is excluded.

We therefore consider theories with all four sectors and establish by exhaustion that deviation from the VERSF realization in any sector forces violation of M1–M5 or their derived consequences. Five deviation classes cover the possibilities: one per sector (Classes I–IV), plus deviation in the cross-sector compatibility system (Class V).

14.3 Deviation Class I — Carrier deviation

Suppose the carrier sector does not take the form specified by Lemma 4.1: Green's function of a self-adjoint linear elliptic operator.

Case (a): nonlinear leading-order carrier. Violates Corollary 3.6.1 (linearity at leading order), itself a consequence of Theorem 3.6. Hence violates M2 + M3.

Case (b): non-self-adjoint carrier. The comparison relation of M2 is reciprocal (Lemma 4.1 step iv); non-self-adjointness breaks this reciprocity in the realization. Hence violates M2.

Case (c): hyperbolic carrier. Retarded or advanced Green's functions require a primitive directed time parameter. Under VERSF's emergent-time commitment (Theorem 3.3), such a parameter is externally supplied, violating M4. More fundamentally, a hyperbolic operator encodes propagation relative to a pre-existing causal structure — which, alongside the causal ordering reconstructed from commitment events, constitutes a second independent causal structure, excluded by Theorem 9.1 (minimality). Further, retarded/advanced Green's functions are non-symmetric, reducing to Case (b). *Caveat (per §12.5, §13.4): the exclusion above is strongest under VERSF's framework-level emergent-time commitment; a block-universe variant with primitive time could admit hyperbolic carriers, at the cost of reconstructing a different framework.*

Case (d): nonlocal carrier. Source data at x instantaneously influences realization at x′ for spacelike-separated (x, x′), violating ordering consistency (Theorem 3.4).

Conclusion. Any carrier deviation violates M2, M3, M4, or Theorem 3.4.

14.4 Deviation Class II — Coupling deviation

Suppose the coupling does not take the form $\mathcal{L}_{int} \sim \lambda \cdot \phi^2 \cdot \ln \sigma$ of Lemma 5.1.

Case (a): wrong scalar structure ($n \neq 2$ or odd n). Violates the ± reflection symmetry or the leading-scalar argument of Lemma 5.1.

Case (b): power-law $f(\sigma) = \sigma^\alpha$ with $\alpha \neq 0$. Non-marginal. Introduces reference scale μ_0. By the two-case analysis of Lemma 5.1: μ_0 externally supplied violates M4; μ_0 dynamically generated violates Theorem 9.1.

Case (c): discontinuous $f(\sigma)$. Comparison across neighbouring σ-values requires continuity of f under the comparison relation of M2.

Conclusion. Any coupling deviation violates M2, M4, or Theorem 9.1.

14.5 Deviation Class III — Closure deviation

Suppose the global closure law does not take the form $N_\Sigma \cdot \theta(\alpha) = 2\pi \cdot \chi$, $\chi \in \mathbb{Z}$ of Lemma 6.1.

Case (a): closure with continuous parameter. Requires external tuning, violating Theorem 3.5 (internal closure).

Case (b): closure without quantization, Σ topologically nontrivial. Holonomy accumulates around nontrivial 1-cycles; unquantized closure is inconsistent with the winding structure of Σ.

Case (c): closure at non-integer χ. Non-integer winding is not a well-defined global invariant of closed loops on Σ.

Conclusion. Any closure deviation violates M4 or the topological structure of Σ supplied by Lemma 7.1.

14.6 Deviation Class IV — Substrate deviation

Suppose the substrate is not simplicial, or its minimal-fact count differs from the combinatorially determined value.

Case (a): non-simplicial substrate. Contains composition cycles. By Lemma 7.1 Part A, cycle endpoints either carry inconsistent values (violates M3) or are redundant (violates Theorem 9.1).

Case (b): deviation in minimal-fact count. Lemma 7.1 Part B establishes that K is uniquely fixed by stability conditions (i)–(v) — all derived within this paper. Any deviation violates at least one:

– insufficient vertex count → fails (ii) persistence or (iii) composition consistency; – excess vertex count → fails (v) minimality via redundant composition paths.

Epistemic status. The structural claim — that K is uniquely determined by (i)–(v) — is established within the paper. The specific evaluation K = 7 is the programme result (K = 7 convergence paper). A deviation K ≠ 7 is therefore a challenge to the programme computation of conditions fixed in the paper; it is not an escape from the framework, only from the numerical value.

Conclusion. Any substrate deviation violates M3, Theorem 9.1, or one of the combinatorial stability conditions of Lemma 7.1.

14.7 Deviation Class V — Compatibility deviation

Suppose the four sectors are not jointly constrained by K = 0.

By Theorem 8.3, K = 0 is forced by M3 + Theorem 9.1. A theory with K ≠ 0 falls into:

Case (a): simultaneously-used incompatible specifications. Direct M3 violation.

Case (b): redundant specifications. Violates Theorem 9.1 if they agree; reduces to (a) otherwise.

Case (c): continuous deformation of sectors within K = 0. Excluded by Theorem 10.3 (substrate rigidity) — the three-case analysis shows no non-trivial continuous family preserves K.

Conclusion. Any compatibility deviation violates M3, Theorem 9.1, or Theorem 10.3.

14.8 Exhaustion

Classes I–V cover every way an admissible-sector theory may deviate from the VERSF realization. A theory avoiding all five deviations is, sector by sector, the VERSF realization. Residual representational freedom is then exactly $\mathcal{G}$_K (Theorem 10.2); continuous deformations beyond $\mathcal{G}$_K are excluded (Theorem 10.3).

Hence no admissible theory lies outside the $\mathcal{G}$_K-orbit of VERSF. ■

14.9 Corollary

[Corollary 14.2] Structural uniqueness of physical law. Within the class of theories satisfying M1–M5, VERSF is the unique realization of physical structure, up to the action of the equivalence group $\mathcal{G}$_K whose substrate factor is finite.

14.10 Interpretation

Theorem 11.4 and Theorem 14.1 together establish the biconditional

Admissibility ⇔ VERSF (within the meta-principle class M1–M5, up to $\mathcal{G}$_K)

Theorem 11.4 (forward direction): admissibility conditions force the VERSF realization.
Theorem 14.1 (reverse direction): no admissible theory lies outside the VERSF equivalence class.

The biconditional is stronger than either direction alone. It asserts that, within the stated meta-principle class, admissibility and VERSF are coextensive — not merely that VERSF is *obtainable* from admissibility, but that admissibility *is* VERSF, within the stated class.

Any escape from this biconditional requires one of three moves, each costed and falsifiable:

- reject one of M1–M5 (cost tabulated in §12.3);
- produce an admissible structure outside $\mathcal{G}$_K (constrained by §12.4 and Theorem 10.3);
- identify an error in one of the cross-referenced main-text theorems (direct technical challenge).

No rhetorical move blocks any of these; each is open to mathematical attack on its stated ground.

Appendix A: Dependency Table

Claim	Status	Depends on
M1–M5	Assumptions	None
Theorem 3.1 (persistence)	Theorem	M1, M2, M5
Theorem 3.2 (well-posed realization)	Theorem	M2, M3
Theorem 3.3 (irreversibility)	Theorem	M3, M5
Theorem 3.4 (ordering consistency)	Theorem	M2, M3
Theorem 3.5 (internal closure)	Theorem	M4
Theorem 3.6 (compositional closure)	Theorem	M1, M2, M3
Corollary 3.6.1 (linearity at leading order)	Consequence	Theorem 3.6
Lemma 4.1 (carrier)	Lemma	Theorems 3.1, 3.2, 3.4, 3.6
Lemma 5.1 (coupling)	Lemma	Lemma 4.1, Theorems 3.4, 3.5, 9.1
Lemma 6.1 (closure)	Lemma	Theorem 3.5, Lemma 7.1 (topology)
Lemma 7.1 Part A (simpliciality)	Theorem	M1, M3, M4, Theorems 3.2, 3.3, 9.1
Lemma 7.1 Part B (K combinatorially determined)	Lemma	Part A, stability conditions (i)–(v)
Programme evaluation K = 7	Programme evaluation	K = 7 convergence paper (downstream computation)
Theorem 8.2 (K non-trivial)	Theorem	Lemmas 4.1–7.1, §10
Theorem 8.3 (compatibility collapse)	Theorem	M3, Theorem 9.1
Theorem 9.1 (minimality)	Theorem	M3, M4
Theorem 9.2 (stability, abstract)	Theorem	M1, M5
Theorem 9.3 (observational closure)	Theorem	M4
Theorem 10.2 (structure of $\mathcal{G}$_K)	Theorem	§§4–8
Theorem 10.3 (substrate rigidity)	Theorem	M4, Theorem 9.1, Theorem 10.2
Theorem 11.1 (reconstruction)	Theorem	All preceding
Theorem 11.4 (main uniqueness)	Theorem	Theorem 11.1
§13 (objection responses)	Cross-reference	Theorems 3.6, 4.1, 5.1, 7.1, 8.3, 10.2, 10.3, 11.1
Theorem 14.1 (no-go for non-VERSF)	Theorem	Theorems 11.1, 10.2, 10.3 (contrapositive of 11.4)
Corollary 14.2 (structural uniqueness)	Consequence	Theorem 14.1

Appendix B: Meta-Principle Rejection Costs

Framework	Principle(s) challenged	Cost / caveat
Many-worlds unitary QM	M1 under global reading	Interpretive dependence

Framework	Principle(s) challenged	Cost / caveat
Retrocausal interpretations	M3 (model-dependent)	Technical analysis per model required
String landscape (free moduli)	M4	Internally-closed completions may satisfy M4
Standard QFT (tuned constants)	M4	Internal-completion question open
Causal sets (standard)	M4 (sprinkling density)	Internally-closed variants possible
Primitive-time block universe	None typically	Theorem applies
Strict instrumentalism	M2 under internal-comparability reading	Depends on version
Paraconsistent physics	M3	Theorem does not apply

Appendix C: Summary of Changes from Version 7

- **§2 restructured as "Why M1–M5 Are Not Optional: Constitutive Conditions for Physical Theory."** M1–M5 are no longer assumptions. They are derived as Theorems 2.1–2.5 from a constitutive criterion (Definition 2.1: conditions C1–C5) for what counts as a candidate fundamental physical theory. Corollary 2.6 establishes that any such theory satisfies M1–M5.
- **Scope corrections on M3 and M4.** Theorem 2.3 (necessity of self-consistency) is now explicitly scoped to determinate physical content, not all logical systems. A remark clarifies that paraconsistent or dialetheic logics in abstract contexts are not excluded, only contradictions in the assignment of determinate physical content. Theorem 2.4 (necessity of completeness) is now explicitly scoped to candidate fundamental theories. A remark clarifies that effective theories, phenomenological models, and intermediate-scale descriptions are not excluded; they are simply not in the class to which the theorem applies.
- **Distilled statement added (§2.7).** "A framework that cannot produce stable, comparable, non-contradictory, internally-specified records of what happened is not a candidate fundamental physical theory." This is the single-sentence form of what M1–M5 jointly assert.
- **§12.1 updated** to reflect that M1–M5 are constitutive, not optional. Rejecting a meta-principle now means accepting that the framework in question is not a candidate fundamental physical theory.
- **§12.3 framework classification table updated** — column header changed from "Principle(s) challenged" to "Condition at issue"; entries reframed to indicate whether failure means "not fundamental" or "alternative reading applies." Paraconsistent physics entry specifically clarifies that the theorem does not apply if determinate content is contradictory, but non-classical meta-logic with consistent determinate content is not excluded.

- **Abstract updated** to note that M1–M5 are established in §2 as constitutive conditions for a framework to count as a candidate fundamental physical theory.

Otherwise unchanged from version 7.

Appendix D: Summary of Changes from Version 6

Four surgical tightenings, each addressing a potential referee attack surface:

- **Hyperbolic exclusion strengthened (§4, §13.4, §14.3).** The exclusion of hyperbolic operators in the carrier reconstruction is no longer a pure emergent-time argument. The new formulation adds a minimality argument: a hyperbolic operator encodes propagation relative to a *pre-existing* causal structure, which alongside the reconstructed ordering would constitute a *second independent causal structure* — excluded by Theorem 9.1 (minimality). This reframes the exclusion from "time issue" to "double structure forbidden," closing the framework-commitment gap identified in the referee review of §12.5.
- **K = 7 framing sharpened (abstract, §7, §12.2).** The structural claim is now formulated as: "the theorem fixes the substrate to a unique combinatorial fixed point K determined by (i)–(v); programme work evaluates this fixed point as K = 7." This distinguishes the structural result (within paper) from the numerical evaluation (programme) more crisply than prior phrasings.
- **Marginality reinforcement (§5).** A concluding sentence added after Lemma 5.1's exclusion statement, explicitly closing the dynamically-generated-scales case: any such scale requires either an independent sector (violating minimality) or additional degrees of freedom not fixed by admissibility. This removes any remaining softness in the Case B exclusion.
- **Biconditional hardened (abstract, §14.1 intro, §14.10).** The biconditional statement is now uniformly formulated as "Admissibility ⇔ VERSF (within the meta-principle class M1–M5, up to $\mathcal{G}$_K)." The parenthetical explicitly scopes the claim to the stated meta-principles, preventing an easy philosophical attack of the form "surely there are conceivable physics outside any meta-principle class."

Otherwise unchanged from version 6.

Appendix E: Summary of Changes from Version 5

- **Summary for the General Reader added** — plain-language introduction preceding the technical abstract. Covers the five meta-principles in accessible terms, previews the four reconstructed structures, acknowledges the conditional nature of the result, and mentions the biconditional. Intended for readers outside the foundations-of-physics community.

- **Table of Contents added** — full section and subsection listing, including §13 objection classes and §14 deviation classes.
- **References section added** — structured in three blocks: VERSF programme works ([P1]–[P6]) with their role in this paper flagged; external mathematical and physical background for elliptic PDE, RG flow, functional equations, simplicial topology, PGL(3,2), holonomy, and the frameworks discussed in §12.3; and a closing note clarifying that the paper's argument is self-contained within M1–M5 and external references are background only.

Otherwise unchanged from version 5.

Appendix F: Summary of Changes from Version 4

- **Section 14 added** — no-go theorem for non-VERSF theories, stated as the contrapositive of Theorem 11.4. Organised by deviation class (Classes I–V, one per sector plus compatibility). Each deviation case cross-referenced to the main-text theorem establishing its exclusion. Tightened from an earlier draft in the following specific ways:
    - *K = 7 framing preserved.* §14.6 maintains the v4 distinction between the structural claim "K is uniquely determined by (i)–(v)" (within paper) and the evaluative claim "K = 7" (programme work). A deviation $K \neq 7$ is therefore a challenge to the programme computation, not an escape from the framework.
    - *Hyperbolic-carrier caveat added.* §14.3 Case (c) cross-references §12.5 and §13.4: the exclusion of hyperbolic operators depends on VERSF's framework-level emergent-time commitment, acknowledged as such.
    - *Exhaustion rigorously justified.* §14.2 now makes explicit that the four-sector-plus-compatibility structure is forced by Theorem 11.1 on any theory satisfying M1–M5; theories lacking this structure fail admissibility directly and are not escapes.
    - *Non-self-adjoint carrier case added* (§14.3 Case b) — previously missing. Excluded via the reciprocity argument of Lemma 4.1 step (iv).
    - *Continuity case added for coupling* (§14.4 Case c).
    - *Continuous-deformation case added for compatibility* (§14.7 Case c), cross-referencing Theorem 10.3.
- **§14.10 interpretation section** makes explicit that Theorem 11.4 and Theorem 14.1 together express the biconditional Admissibility ⇔ VERSF (up to $\mathcal{G}$_K). This is the strongest claim the paper makes; it requires *both* directions — reconstruction and exclusion — neither alone being sufficient.
- **Abstract updated** to mention Section 14 and the biconditional.
- **Appendix A** extended to list Theorem 14.1 and Corollary 14.2.

Otherwise unchanged from version 4.

Appendix G: Summary of Changes from Version 3

Preserved from previous version:

- **Section 13 added** — anticipated referee objections and responses. Six principal structural objections identified; each response cross-referenced to main-text theorems rather than re-proved.
  - Objection A uses the abstract composition ⊞ of Theorem 3.6 rather than presupposing "+".
  - Objection B distinguishes the marginal functional equation from bare scale-covariance.
  - Objection C separates the structural claim "K uniquely determined by (i)–(v)" (within paper) from the evaluative claim "K = 7" (programme work).
  - Objection D explicitly acknowledges the emergent-time commitment as a framework-level choice (also noted in §12.5).
  - Objection E uses the three-case rigidity analysis of Theorem 10.3, with discreteness of Aut($\mathcal{P}$) identified as load-bearing.
  - Objection F classifies candidate alternatives into three categories ($\mathcal{G}$_K-equivalent, excluded, violate M1–M5).
- **§12.5 expanded** — emergent-time commitment now explicitly listed as a remaining-openness item.
- **Appendix A extended** to list §13 as a cross-reference layer over main-text theorems.

References

VERSF Programme Works

This paper is part of the VERSF (Void Energy-Regulated Space Framework) theoretical physics programme. The programme's papers are hosted at **versf-eos.com** and developed through the AIDA Institute. References below identify the programme works cited in this paper; the reader is directed to the programme website for the current list of available papers.

[P1] Taylor, K. *The K = 7 Convergence: Minimal-Fact Substrate and the Fano Plane.* VERSF Programme paper. Available at versf-eos.com. *Cited in §7, §13.3, §14.6 for the combinatorial evaluation of K = 7 from stability conditions (i)–(v).*

[P2] Taylor, K. *Fact-Momentum and the Commitment-Event Bath.* VERSF Programme paper. Available at versf-eos.com. *Cited in §8.1 for the normalization functional $\lambda[\mathcal{P}, K]$ of the ordering coupling.*

[P3] Taylor, K. *The Closure Scale ξ: A Structural Consequence of Three Axioms.* VERSF Programme paper. Available at versf-eos.com. *Background on the closure-scale derivation and its role in the substrate structure.*

[P4] Taylor, K. *Proto-Time and Emergent Lorentz Invariance.* VERSF Programme paper. Available at versf-eos.com. *Background on the emergent-time commitment underlying the carrier-sector ellipticity argument of §4 and §14.3.*

[P5] Taylor, K. *No-Go Theorem for Non-Simplicial Relational Substrates.* VERSF Programme paper. Available at versf-eos.com. *Related simpliciality argument; the version in this paper (§7 Part A) is self-contained.*

[P6] Taylor, K. *Single-Source Theorem: Observables as Functionals of Committed Record Density.* VERSF Programme paper. Available at versf-eos.com. *Background on the single-source structure of observables entering the compatibility system.*

External Mathematical and Physical Background

The following standard references are relevant for the mathematical objects and physical concepts used in the paper. The paper's arguments are self-contained within the VERSF framework; these references are included for readers who wish to consult the broader mathematical context.

Elliptic operators and Green's functions (§4). Evans, L. C. *Partial Differential Equations.* Graduate Studies in Mathematics, American Mathematical Society, 2nd edition (2010). Gilbarg, D. and Trudinger, N. S. *Elliptic Partial Differential Equations of Second Order.* Springer (2001).

Renormalization group flow and marginality (§5). Wilson, K. G. "The Renormalization Group and Critical Phenomena." *Reviews of Modern Physics* 55, 583–600 (1983). Weinberg, S. *The Quantum Theory of Fields, Volume II.* Cambridge University Press (1996).

Cauchy functional equation (§5). Aczél, J. *Lectures on Functional Equations and Their Applications.* Academic Press (1966).

Simplicial complexes and combinatorial topology (§§7, 10). Hatcher, A. *Algebraic Topology.* Cambridge University Press (2002). Spanier, E. H. *Algebraic Topology.* Springer (1966, reprinted).

PGL(3,2) and the Fano plane (§7, §10). Conway, J. H. and Sloane, N. J. A. *Sphere Packings, Lattices and Groups.* Springer, 3rd edition (1999), for PGL(3,2) and related group-theoretic structures. Polster, B. *A Geometrical Picture Book.* Springer (1998), for the Fano plane.

Holonomy and topological quantization (§6). Nakahara, M. *Geometry, Topology and Physics.* Institute of Physics Publishing, 2nd edition (2003).

Causal set theory (mentioned in §12.3). Sorkin, R. D. "Causal Sets: Discrete Gravity." In *Lectures on Quantum Gravity*, Gomberoff, A. and Marolf, D. (eds.), Springer (2005).

Many-worlds interpretation of quantum mechanics (mentioned in §12.3). Wallace, D. *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation.* Oxford University Press (2012).

Paraconsistent logic (mentioned in §12.3). Priest, G. *In Contradiction: A Study of the Transconsistent.* Oxford University Press, 2nd edition (2006).

Note on Citation Structure

This paper presents a conditional uniqueness result. Its argument chain is internally closed within the meta-principles M1–M5; programme works [P1]–[P6] are cited only as background or, in the case of [P1], for the specific numerical evaluation K = 7 supplied as a downstream result (§7 Part B). External mathematical references are included for context; no part of the argument depends on them. A reader may check the main argument by reading this paper alone, consulting the programme works for related structural context and the external references for mathematical background.