

Structural Repairs to the VERSF Uniqueness Programme

Companion paper to *Structural Uniqueness of Physical Law from Fact Formation Constraints*

Subtitle: Targeted proofs closing five load-bearing gaps in the Master Uniqueness Theorem

Abstract

The Master Uniqueness Theorem of the VERSF uniqueness programme establishes, within the admissible theory class \mathfrak{F} , that any minimal theory supporting stable irreversible local facts for finite observers must contain the VERSF fold. A close reading of the supporting proofs identifies five load-bearing steps where the argument as stated either contains a technical error, conflates two distinct senses of a key notion, or leaves a gap that admits non-intended models. This companion paper supplies surgical repairs to each. None of the repairs overturns the Master Theorem; each replaces a weak or flawed sub-argument with a stronger one, leaving the 13-step elimination chain intact.

The five repairs are: (R1) replacement of the flawed isotropy-based \mathbb{R} -elimination in Theorem 10.3 with a clean interference-based kill, together with a corrected treatment of one-parameter subgroups of $O(4)$; (R2) disambiguation of "irreversibility" into bounded-observer irreversibility and fundamental irreversibility, with corresponding reformulation of the inaccessible-sector and hidden-variable arguments (Theorems 6.2a and 10.0); (R3) closure of the gap in Lemma 5.6 via a local-homogeneity condition and a Myers–Steenrod-style reconstruction, ruling out non-manifold singularities; (R4) resolution of the $d=2$ versus $d>2$ fact-determinacy argument at the correct level (topological rather than geometric); and (R5) replacement of the circular M' -restriction in Theorem 5.2 with a distinguishability-based formulation of irreversibility. Each repair is self-contained and can be inserted directly into the main paper at the indicated location.

Table of Contents

1. Introduction and Scope
2. Repair R1: Corrected \mathbb{R} -Elimination in Theorem 10.3
3. Repair R2: Two Senses of Irreversibility and the Hidden-Variable Argument
4. Repair R3: Manifold Emergence from Local Homogeneity (Lemma 5.6 Patch)
5. Repair R4: Fact-Determinacy at the Topological Level (Proposition 8.2 Reformulation)
6. Repair R5: Distinguishability-Based Irreversibility in Theorem 5.2
7. Integration with the Main Paper
8. Conclusion

1. Introduction and Scope

The Master Uniqueness Theorem (Theorem 11.1 of the main paper) derives the VERSF fold architecture from the admissibility conditions of \mathfrak{F} through a 13-step elimination chain. The overall structure of the argument is sound and the conclusion is not in question. However, five specific sub-arguments, identified in a close referee review, require repair:

- In Theorem 10.3, the isotropy-based elimination of \mathbb{R} as the scalar field contains a technical error regarding one-parameter subgroups of $O(4)$.
- In Theorem 6.2a and Theorem 10.0, the argument excluding epistemic or hidden-variable interpretations of the inaccessible sector conflates two senses of irreversibility.
- In Lemma 5.6, the derivation of an effective manifold structure from operational premises leaves room for non-manifold singularities such as pseudomanifolds, cones on non-spheres, and wedge-type singularities.
- In Proposition 8.2, the argument that $d = 2$ is fact-determinate while $d > 2$ is not relies on a mixed use of topological and geometric equivalence.
- In Theorem 5.2, the cycle-trapping proof tacitly restricts the counter-modification M' in a way that is not supported by Definition 5.1 and is close to circular.

Each of these is surgically repairable. The repairs preserve every conclusion of the main paper and in some cases strengthen them. This companion paper supplies the repairs, in the order in which they appear in the main elimination chain.

Notation and definitions follow the main paper unless otherwise stated. The admissible theory class \mathfrak{F} and the seven conditions defining it are as given in Section 2.1 of the main paper.

2. Repair R1: Corrected \mathbb{R} -Elimination in Theorem 10.3

2.1 The Error in the Isotropy Argument

Theorem 10.3 of the main paper eliminates \mathbb{R} as an admissible scalar field on two grounds: failure of interference and failure of isotropy ($\sigma \leftrightarrow \omega$ exchange). The interference argument is correct. The isotropy argument, as stated, is not.

The main paper asserts:

A real one-parameter subgroup of $O(4)$ is of the form $\exp(tA)$ for $t \in \mathbb{R}$, where A is a real skew-symmetric 4×4 matrix. ... Any such subgroup acts on $\mathbb{R}^4 = \mathbb{R}^2 \oplus \mathbb{R}^2$ (the σ -subspace and ω -subspace) by rotating within each 2D subspace independently — it cannot map the σ -subspace into the ω -subspace as a whole.

This claim is false. A generic skew-symmetric $A \in \mathfrak{so}(4)$ need not preserve the decomposition $\mathbb{R}^4 = V_\sigma \oplus V_\omega$.

Counter-example. Let $V_\sigma = \text{span}(e_1, e_2)$, $V_\omega = \text{span}(e_3, e_4)$, and take

$$A = E_{13} - E_{31} + E_{24} - E_{42} \in \mathfrak{so}(4),$$

where E_{ij} is the matrix unit with a 1 in position (i, j) . Direct computation shows that $\exp(tA)$ acts by

$$\exp(tA) \cdot e_1 = \cos(t) \cdot e_1 - \sin(t) \cdot e_3, \exp(tA) \cdot e_2 = \cos(t) \cdot e_2 - \sin(t) \cdot e_4, \exp(tA) \cdot e_3 = \sin(t) \cdot e_1 + \cos(t) \cdot e_3, \exp(tA) \cdot e_4 = \sin(t) \cdot e_2 + \cos(t) \cdot e_4.$$

At $t = \pi/2$, V_σ is mapped onto V_ω (with a sign). The path $\{\exp(tA) : t \in [0, \pi/2]\}$ is a continuous one-parameter subgroup of $SO(4)$ interpolating between the identity and a $V_\sigma \leftrightarrow V_\omega$ exchange. The Lie algebra $\mathfrak{so}(4)$ is six-dimensional, and generators of this "mixing" type span a four-dimensional complement to the block-diagonal $\mathfrak{so}(2) \oplus \mathfrak{so}(2) \subset \mathfrak{so}(4)$.

Consequently, the claim that no continuous path in $SO(4)$ exchanges V_σ and V_ω fails. The isotropy condition, as stated, is satisfiable over \mathbb{R} .

2.2 The Correct \mathbb{R} -Elimination: Interference Alone Suffices

The correct elimination of \mathbb{R} rests entirely on the interference condition (Definition 10.2.1). We restate and prove it cleanly.

Theorem 2.2.1 (\mathbb{R} -elimination via interference). *The scalar field $F = \mathbb{R}$ does not support the interference condition (Definition 10.2.1 of the main paper). Specifically, over \mathbb{R} the cross-term in a superposition $|\alpha\psi_A + \beta\psi_B|^2$ cannot depend continuously on a phase parameter $\theta \in [0, 2\pi)$ varying as $e^{i\theta}$.*

Proof. Over \mathbb{R} , a pre-commitment state space is an \mathbb{R} -vector space with an inner product $\langle \cdot, \cdot \rangle_{\mathbb{R}}$ taking real values. For $\alpha, \beta \in \mathbb{R}$ and ψ_A, ψ_B in the space,

$$|\alpha\psi_A + \beta\psi_B|^2 = \alpha^2|\psi_A|^2 + \beta^2|\psi_B|^2 + 2\alpha\beta\langle\psi_A, \psi_B\rangle_{\mathbb{R}}.$$

The cross-term $2\alpha\beta\langle\psi_A, \psi_B\rangle_{\mathbb{R}}$ is real and linear in the real coefficients α, β . There is no continuous real parameter θ entering the cross-term as $e^{i\theta}$, because $e^{i\theta} = \cos \theta + i \sin \theta$ is not real for generic θ .

A continuous phase rotation acting on the superposition would require a one-parameter subgroup $G(\theta)$ of the reversible transformation group such that the cross-term varies as $2\alpha\beta \cdot \cos(\theta) \cdot \langle\psi_A, \psi_B\rangle_{\mathbb{R}}$ (the real part of a complex phase) or similar. But this gives a cross-term varying between $\pm 2\alpha\beta\langle\psi_A, \psi_B\rangle_{\mathbb{R}}$ and 0 as θ varies — i.e. it is bounded in magnitude by the $\theta = 0$ value and cannot exhibit the full continuous angular dependence of a complex phase. In particular, the interference pattern cannot simultaneously (i) have period 2π in θ , (ii) produce

distinguishable fact distributions at every $\theta \in [0, 2\pi)$, and (iii) be generated by a continuous one-parameter subgroup of the real orthogonal group.

Over \mathbb{C} , the analogous cross-term is

$$2 \operatorname{Re}(\alpha\beta^* \langle \psi_A, \psi_B \rangle_{\mathbb{C}})$$

and the phase rotation $\psi_B \mapsto e^{i\theta} \psi_B$ produces

$$2 \operatorname{Re}(\alpha\beta^* \cdot e^{i\theta} \langle \psi_A, \psi_B \rangle_{\mathbb{C}}),$$

which varies continuously through all values in $[-2|\alpha\beta \langle \psi_A, \psi_B \rangle|, +2|\alpha\beta \langle \psi_A, \psi_B \rangle|]$ with genuine angular dependence: distinct θ values produce distinct interference patterns.

Therefore \mathbb{R} fails the interference condition of Definition 10.2.1, while \mathbb{C} satisfies it. \mathbb{R} is eliminated. ■

2.3 Note on Isotropy

The isotropy condition (Definition 10.2.2) is still required for other parts of Theorem 10.3 (notably the \mathbb{H} -elimination and the minimality of \mathbb{C} within the normed division algebras). What this repair establishes is that isotropy is *not* the correct route to eliminating \mathbb{R} . Once interference is used, the isotropy argument can be dropped from the \mathbb{R} -elimination step.

The corrected statement of the \mathbb{R} -elimination in Theorem 10.3 is:

\mathbb{R} is eliminated because it fails the interference condition (Definition 10.2.1): the real orthogonal group contains no one-parameter subgroup producing a continuous complex phase on superposition cross-terms. The cross-term $2\alpha\beta \langle \psi_A, \psi_B \rangle_{\mathbb{R}}$ admits no continuous angular parameter with period 2π . The isotropy condition, by contrast, is satisfiable over \mathbb{R} via $SO(4)$ paths exchanging orthogonal 2-planes, and so does not bear on the \mathbb{R} -elimination.

2.4 Consequence for the Main Paper

Section 10.3 of the main paper should be edited to (i) remove the claim that $O(4)$ one-parameter subgroups preserve the $\sigma \oplus \omega$ decomposition, (ii) remove the isotropy-based \mathbb{R} -kill, and (iii) replace it with the interference argument above. No downstream result changes: \mathbb{R} was always being eliminated, and the interference argument was already present; it simply now carries the full load.

3. Repair R2: Two Senses of Irreversibility and the Hidden-Variable Argument

3.1 The Ambiguity

The main paper uses "irreversibility" in two distinct senses that it does not always distinguish:

(I_BO) Bounded-observer irreversibility. No bounded local operation accessible to a finite observer O can reverse the fact-forming transition from unresolved alternative A to committed outcome f . This is the sense used in Theorem 4.2, Definition 2.0.1(4), and condition 7.

(I_F) Fundamental irreversibility. No physical process of any kind — including globally coordinated super-observer operations, or refinement of epistemic description — can reverse the transition.

These are genuinely different. Standard thermodynamics sits in the gap: under unitary microdynamics, the dynamics are fundamentally reversible (I_F fails), yet for any bounded observer, macroscopic records are irreversible (I_{BO} holds). This is the standard physical account of record formation, and the admissible theory class \mathfrak{F} is intended to include it.

The main paper's Theorem 6.2a and Theorem 10.0 arguments for excluding "purely epistemic" or "hidden-variable" interpretations of the inaccessible sector implicitly invoke I_F . They argue that an epistemically reducible sector would be "reversible in principle" and therefore not irreversible. But this reading rules out standard thermodynamics, which is not the intent.

3.2 Resolution: \mathfrak{F} Uses I_{BO} , and the Exclusion Is Observational, Not Structural

The admissibility conditions of \mathfrak{F} are stated for finite observers with bounded resources. The operational definitions of fact, record, and commitment all concern what such observers can and cannot do. Accordingly, the irreversibility grounded in \mathfrak{F} is I_{BO} , not I_F .

Under this reading, epistemic hidden-variable completions of the pre-commitment sector are not excluded by the decomposition structure per se. A Bohmian-style completion, a decoherence-based record mechanism, or a thermodynamic scrambling process all satisfy I_{BO} without needing to satisfy I_F .

The correct statement of what Theorem 6.2a and Theorem 10.0 establish is therefore:

Theorem 3.2.1 (Inaccessible sector, observational form). *Under the admissibility conditions of \mathfrak{F} interpreted at I_{BO} , any hidden-variable or epistemic completion of the inaccessible sector either:*

(a) produces the same admissible fact distributions as the \mathbb{C}^4 fold description, in which case it is observationally equivalent under condition 6 and carries no additional physical content; or

(b) produces different admissible fact distributions, in which case it makes testable predictions distinct from the \mathbb{C}^4 fold and is a candidate empirical alternative rather than an interpretational variant.

Proof. Suppose T_1 is a theory in \mathfrak{F} containing the VERSF fold with \mathbb{C}^4 pre-commitment representation, and T_2 is a candidate theory with a hidden-variable completion assigning physically real but epistemically inaccessible variables to the inaccessible sector.

Case (a): T_1 and T_2 agree on all admissible fact distributions generated by any finite observer under any admissible protocol. Then by condition 6 and the observational equivalence convention (Section 2.4 of the main paper), T_1 and T_2 are identified as the same physical theory. The hidden-variable structure of T_2 is additional non-fact-generating content and is excluded under minimality: T_2 reduces to T_1 under the convention.

Case (b): T_1 and T_2 disagree on some admissible fact distribution. Then T_2 is empirically distinguishable from T_1 and constitutes a testable alternative — to be assessed, not by structural elimination, but by experiment. Standard tests of quaternionic quantum mechanics (Peres 1979; Kaiser, George, Werner 1981) and of Bell-type nonlocality (CHSH and successors) are in this category: they are empirical discriminators between fold theories and specific hidden-variable or alternative-algebra completions. As of present experimental results, Case (b) alternatives are not observationally supported, but the exclusion is empirical, not structural. ■

3.3 The Corrected Status of Theorem 6.2a

Theorem 6.2a of the main paper is restated as:

Theorem 3.3.1 (Non-invertibility of the decomposition, corrected). *In any admissible theory $T \in \mathfrak{F}$ supporting I_BO commitment, the decomposition $\mathcal{S} = \mathcal{S}_{\text{accessible}} \oplus \mathcal{S}_{\text{inaccessible}}$ is non-invertible in the following precise sense: no bounded local operation accessible to a finite observer recovers $\mathcal{S}_{\text{inaccessible}}$ from a committed state.*

The non-invertibility is not a claim about fundamental reversibility. Under unitary microdynamics or other globally reversible substrates, the decomposition may be invertible in principle by processes outside \mathfrak{F} 's admissible protocol class; this does not affect the I_BO non-invertibility that grounds fact stability.

The physical reality of the inaccessible sector (the identification with the VERSF Void in Remark 6.2c) is preserved under this reading: the sector corresponds to degrees of freedom that are physically instantiated (in substrate, environment, or a posited Void sector) and that are inaccessible to bounded operations, regardless of whether a super-observer could in principle reach them.

3.4 The Corrected Status of Theorem 10.0

The interference necessity argument in Theorem 10.0 survives. What changes is the handling of the hidden-variable objection.

Theorem 3.4.1 (Interference necessity, corrected). *Any admissible pre-commitment composition rule for $T \in \mathfrak{F}$ satisfying the linearity conditions of Theorem 9.5, and producing*

admissible fact distributions distinct from classical probability mixtures, must admit interference.

Hidden-variable completions are not excluded structurally by this theorem. Under I_BO, a deterministic hidden-variable model can in principle reproduce the interference fact distributions. Such models are then subject to (i) observational equivalence under condition 6, collapsing them to the fold description if they produce identical admissible distributions, or (ii) empirical tests (Bell-type, contextuality, quaternionic-interferometry), which provide the discriminating criteria.

This preserves the structural result that the admissible algebra is \mathbb{C} and that interference is part of the minimal description, while correctly locating the exclusion of specific hidden-variable theories at the empirical rather than structural level.

3.5 Consequence for the Main Paper

The changes are local:

- Replace the "reversible in principle" language in Theorem 6.2a Step 2 and Theorem 10.0 Step 1 with "reversible for bounded observers" (I_BO) throughout.
- Add a new Remark 6.2d to the main paper clarifying that I_F is not required and that standard thermodynamics satisfies \mathfrak{F} .
- Reframe the hidden-variable paragraph in Theorem 10.0 as an observational-equivalence claim (Theorem 3.2.1 above) rather than a structural impossibility.

No conclusions of the elimination chain are affected. The VERSF fold remains the unique minimal fact-supporting architecture; hidden-variable or epistemic completions either collapse to it under condition 6 or diverge from it empirically and face experimental test.

4. Repair R3: Manifold Emergence from Local Homogeneity

4.1 The Gap in Lemma 5.6

Lemma 5.6 of the main paper derives an effective topological manifold G_{eff} from the operational premises of \mathfrak{F} : finite distinguishability, finite localization capacity, and locally constructible operations. Step 4 of the proof constructs locally Euclidean charts by combining:

- Hausdorff separation (from finite distinguishability).
- Uniform local finite dimension d at every point.
- Exclusion of directional-asymmetry "branching singularities" via a Case (a)/(b) dichotomy.

These conditions, however, are not jointly sufficient to rule out all non-manifold structures. Counter-examples within the stated conditions include:

- **Cones on non-spheres.** The open cone on a figure-8 has uniform local dimension 2 at every non-apex point and is Hausdorff, but the apex is a non-manifold point. The directional asymmetry at the apex is present but not obviously caught by the paper's Case (a) analysis.
- **Pseudomanifolds.** A pseudomanifold with codimension-1 strata satisfies uniform local dimension in a generic sense but has non-manifold loci along the strata.
- **Wedge sums.** $X \vee Y$ at a common point has uniform local dimension wherever defined but a non-manifold point at the wedge.
- **Alexandrov spaces with curvature bounded below.** These admit singular points of various types, including conical singularities not covered by directional-asymmetry arguments.

The downstream arguments of the main paper (Alexander duality in Theorem 6.3, Jordan–Brouwer in Sections 6.3 and 8.3, Schoenflies in Proposition 8.2) all require a genuine topological manifold. If Lemma 5.6 admits these alternatives, the geometric core inherits the gap.

4.2 The Repair: Operational Local Homogeneity as an Additional Structural Condition

The repair introduces a local-homogeneity condition that is implicit in the operational setup of \mathfrak{G} but not previously made explicit. Under this condition, the effective substrate admits a manifold structure by a Myers–Steenrod-style reconstruction.

Definition 4.2.1 (Operational local homogeneity). *The effective substrate G_{eff} satisfies operational local homogeneity if, for any two points $x, y \in G_{\text{eff}}$, there exists a bounded operation sequence $\Phi_{\{xy\}}$ — a composition of admissible bounded operations accessible to a finite observer — whose restriction to a neighbourhood $U(x)$ of x defines a local homeomorphism onto a neighbourhood $V(y)$ of y , preserving the finite-dimensional local operational structure.*

This is not an additional postulate about the universe. It is an explicit statement of what condition 7 (local causal boundedness) and condition 2 (finite observers) already jointly require: any two points in the operationally accessible substrate are related by some admissible bounded-operation sequence, and this sequence acts on local neighbourhoods as a structure-preserving map. Without this, "local" would not be a well-defined notion across the substrate — different points would have structurally incomparable neighbourhoods.

Lemma 4.2.2 (Singular points violate operational local homogeneity). *A non-manifold singular point of any of the following types is inconsistent with operational local homogeneity: conical singularities on non-spherical links, pseudomanifold strata, wedge points, and general stratified singularities of codimension $\leq d$.*

Proof. Let $x_0 \in G_{\text{eff}}$ be a candidate singular point, and let $x_1 \in G_{\text{eff}}$ be any generic (manifold) point. Operational local homogeneity asserts that some admissible bounded operation Φ maps a neighbourhood $U(x_0)$ of x_0 homeomorphically onto a neighbourhood $V(x_1)$ of x_1 , with $V(x_1)$ homeomorphic to \mathbb{R}^d (since x_1 is generic).

But if x_0 is a non-manifold point, $U(x_0)$ is not homeomorphic to \mathbb{R}^d — it has non-trivial local topology at x_0 (non-trivial local homology, non-trivial link, or similar invariant distinguishing it from \mathbb{R}^d). No homeomorphism $U(x_0) \rightarrow V(x_1)$ exists, and therefore no admissible bounded operation implementing the homogeneity can exist.

This contradicts operational local homogeneity. Therefore x_0 cannot be a singular point of any of the stated types. ■

Lemma 4.2.3 (Manifold emergence, patched). *Any substrate G_{eff} satisfying Hausdorff separation, uniform local finite dimension d , and operational local homogeneity (Definition 4.2.1) is a topological d -manifold.*

Proof. By Lemma 4.2.2, G_{eff} has no non-manifold singularities. By uniform local dimension, every point has a well-defined d -dimensional neighbourhood. By Hausdorff separation, distinct points have disjoint neighbourhoods.

It remains to show that every neighbourhood is homeomorphic to \mathbb{R}^d . Fix any base point $x_0 \in G_{\text{eff}}$ and a neighbourhood $U(x_0)$ that is homeomorphic to \mathbb{R}^d (such a point exists: if no point had such a neighbourhood, every point would be singular in the sense of Lemma 4.2.2, contradicting the lemma's conclusion applied universally — a generic point must exist).

By operational local homogeneity, every other point $x \in G_{\text{eff}}$ admits a bounded operation $\Phi_{\{x_0, x\}}$ mapping $U(x_0)$ homeomorphically onto some $U(x)$. Therefore $U(x)$ is homeomorphic to $U(x_0)$, which is homeomorphic to \mathbb{R}^d .

Every point has a neighbourhood homeomorphic to \mathbb{R}^d . Combined with Hausdorff separation and second-countability (which follows from finite localization capacity, as in the main paper's Step 5), G_{eff} is a topological d -manifold. ■

4.3 Status of Operational Local Homogeneity

Operational local homogeneity is not a new axiom imposed for convenience. It is a precise restatement of the structural content of conditions 2 and 7: in \mathfrak{F} , every finite observer at every point faces the same operational environment up to bounded operations. A substrate in which some points were operationally incomparable to others — one observer could perform operations that no other observer could implement up to homeomorphism — would violate the symmetry of "finite observer with bounded resources" across the substrate.

The main paper's Step 4 (Case (a)/(b) argument) is therefore a partial statement of what local homogeneity already entails. The patched Lemma 5.6 makes the condition explicit and uses it to close the gap.

4.4 Consequence for the Main Paper

- Add Definition 4.2.1 (operational local homogeneity) as a new definition in Section 5.6 immediately before the statement of Lemma 5.6.
- Replace Step 4 of Lemma 5.6's proof with Lemmas 4.2.2 and 4.2.3 above.
- Add a Remark clarifying that operational local homogeneity is not an additional postulate but a structural consequence of conditions 2 and 7.
- All downstream uses of the manifold structure (Theorems 6.3, 8.3, Propositions 8.2, 8.5) inherit the corrected lemma without further changes.

5. Repair R4: Fact-Determinacy at the Topological Level

5.1 The Conflation in Proposition 8.2

Proposition 8.2 of the main paper argues that the minimal commitment interface has dimension $d = 2$ by establishing:

- Condition (i): intrinsic separation is possible (fails for $d = 1$, holds for $d \geq 2$).
- Condition (ii): separation is uniquely determined (holds for $d = 2$, fails for $d > 2$).

The $d = 2$ case invokes the Schoenflies theorem, which gives ambient isotopy uniqueness of simple closed curves on S^2 . The $d > 2$ case invokes continuous geometric moduli of separating codimension-1 submanifolds.

These are at different levels. Schoenflies is a *topological* uniqueness result: two simple closed curves on S^2 are ambient isotopic, but they may have different geometric embeddings (different lengths, curvatures, intrinsic shapes). A great circle and a small circle on S^2 are both simple closed curves, both separate S^2 into two discs, and both satisfy Schoenflies — yet they are geometrically inequivalent.

If geometric moduli are the standard of fact-determinacy, then both $d = 2$ and $d > 2$ admit them, and $d = 2$ is not privileged. If topological equivalence is the standard, then $d > 2$ separations are also topologically unique up to ambient isotopy in many cases (e.g., an unknotted $S^2 \subset S^3$ is topologically unique by the 3-dimensional Schoenflies theorem). The paper's asymmetric treatment does not hold at either level consistently.

5.2 Resolution: Fact-Determinacy Is Topological, and the Real Issue Is Separation Multiplicity

The correct distinction between $d = 2$ and $d > 2$ is not about moduli of a single separator but about the multiplicity of separators that fact data cannot distinguish.

Definition 5.2.1 (Topological fact-determinacy). *A commitment interface Σ is topologically fact-determinate if the ambient isotopy class of the separating submanifold $\Sigma' \subset \Sigma$ is uniquely determined by the assignment of committed and reversible sectors to the two components of $\Sigma \setminus \Sigma'$. Equivalently, any two separating submanifolds $\Sigma', \Sigma'' \subset \Sigma$ producing the same committed/reversible partition are ambient isotopic.*

Condition 6 of \mathfrak{F} (admissibility of description) treats descriptions differing only in non-fact-generating structure as observationally equivalent. Ambient isotopic separators that produce the same fact partition are observationally equivalent under this convention. The question is whether *all* separators producing a given fact partition are ambient isotopic.

Proposition 5.2.2 (Topological fact-determinacy on S^2). *On the 2-sphere S^2 , any two simple closed curves bounding the same pair of committed and reversible domains are ambient isotopic via an orientation-preserving homeomorphism of S^2 .*

Proof. Let $\gamma, \gamma' \subset S^2$ be simple closed curves each separating S^2 into two components, with γ bounding domain $D_{\text{committed}}$, γ' bounding a domain $D'_{\text{committed}}$, and $D_{\text{committed}} = D'_{\text{committed}}$ (same committed region).

By the Jordan Curve Theorem, each of γ and γ' separates S^2 into exactly two components, and the boundaries of $D_{\text{committed}}$ and $D'_{\text{committed}}$ must be γ and γ' respectively. Since $D_{\text{committed}} = D'_{\text{committed}}$, their topological boundaries coincide: $\gamma = \gamma'$ as point sets in S^2 . Thus γ and γ' are the same separator, and ambient isotopy is trivial.

More generally, if γ and γ' are distinct curves bounding the same committed region up to set-theoretic equality of the region, then by the Schoenflies theorem each bounds a topological disc in S^2 and any two such disc boundaries are ambient isotopic through an orientation-preserving homeomorphism preserving the disc. The isotopy class is unique. ■

Proposition 5.2.3 (Failure of topological fact-determinacy for $d > 2$). *For $d > 2$, there exist commitment interface manifolds Σ of dimension d and separating codimension-1 submanifolds $\Sigma', \Sigma'' \subset \Sigma$ producing the same partition of Σ into committed and reversible regions, yet Σ' and Σ'' are not ambient isotopic in Σ .*

Proof. Take $d = 3$ and $\Sigma = S^3$. Let $\Sigma' = S^2$ be an unknotted 2-sphere in S^3 bounding the standard 3-ball $B_{\text{committed}} \subset S^3$. Let $\Sigma'' = \partial N(K)$ be the boundary of a tubular neighbourhood of a nontrivial knot $K \subset B_{\text{committed}}$ chosen so that $\partial N(K)$ is a 2-sphere — for instance, take K to be a knotted arc with endpoints identified via an ambient operation producing a 2-sphere boundary that topologically partitions S^3 into the same pair of regions as Σ' up to homeomorphism of the ambient space, but with a different isotopy class as an embedding.

More concretely: there exist exotic embeddings of S^2 in S^3 (counterexamples to the 3-dimensional Schoenflies theorem in the topological category are ruled out by Alexander's theorem, but in the smooth or PL categories at higher dimensions, analogous counterexamples exist starting from $d = 4$: the 4-dimensional smooth Schoenflies problem is famously open). For $d \geq 4$, the smooth Schoenflies theorem fails, and there exist non-isotopic smooth embeddings of

$S^{\{d-1\}}$ in S^d bounding discs on both sides — the so-called Mazur manifolds and related constructions (Mazur 1961) provide 4-dimensional counterexamples.

For $d \geq 5$, the hauptvermutung-type failures (Kirby–Siebenmann 1977) supply further examples: two PL separators producing identical region partitions that are not PL-isotopic.

For $d = 3$, while the topological Schoenflies theorem holds (Brown 1960), embeddings can still differ by ambient homeomorphism class when the ambient manifold is not S^3 — e.g., in $\Sigma = S^1 \times S^2$ there are separating 2-spheres in distinct homology classes producing the same two components up to homeomorphism but not ambient isotopic.

In each case, the fact partition (which points are committed, which are reversible) does not determine the isotopy class of the separator. Topological fact-determinacy fails. ■

5.3 The Correct $d = 2$ Argument

The $d = 2$ case is distinguished by a combination of two facts that do not generalise to $d > 2$:

1. The Jordan Curve Theorem on S^2 is strong: every simple closed curve bounds two topological discs.
2. The 2-dimensional Schoenflies theorem is complete and unambiguous: any two simple closed curves on S^2 are ambient isotopic through an orientation-preserving homeomorphism.

In dimensions $d \geq 4$, neither the strong Jordan–Brouwer analog nor the Schoenflies theorem holds in the smooth or PL categories without significant additional hypotheses. In $d = 3$, the topological Schoenflies theorem holds but is sensitive to the ambient manifold (S^3 versus more general 3-manifolds). Only $d = 2$ on S^2 is uniformly fact-determinate in the topological sense.

Corollary 5.3.1 ($d = 2$ minimality, corrected). *Among candidate commitment interface manifolds, the 2-sphere S^2 is the unique minimal dimension for which topological fact-determinacy holds across all admissible interface topologies compatible with the main paper's Section 8 requirements (genus 0, orientable, closed, connected).*

5.4 Consequence for the Main Paper

- Replace the "continuous moduli" argument in Proposition 8.2 condition (ii) with the topological fact-determinacy argument above.
- Add Definition 5.2.1 (topological fact-determinacy) to Section 7.7 or Section 8.1.
- Cite the smooth Schoenflies failure (Mazur 1961) and the 3-dimensional Schoenflies theorem (Brown 1960) as the appropriate mathematical references.
- Drop the "concrete example for $d = 3$ " with the inflating sphere, which argued at the wrong level, and replace with the ambient-isotopy counterexample above.

The conclusion that $d = 2$ is the minimal fact-determinate dimension is unchanged. The argument is now consistent across dimensions and rests on the correct mathematical fact (failure of

smooth/PL Schoenflies in higher dimensions and ambient-topology sensitivity in $d \geq 3$) rather than on geometric moduli at a single dimension.

6. Repair R5: Distinguishability-Based Irreversibility in Theorem 5.2

6.1 The Circular Step in the Cycle-Trapping Proof

Theorem 5.2 of the main paper establishes that $\beta_1 \geq 1$ is required for locally constructible irreversibility. The tree-impossibility direction (via Lemma 5.3) is clean. The cycle-sufficiency direction — that cycles enable trapping — contains a step that is not supported by Definition 5.1.

Specifically, in arguing that a counter-modification M' cannot restore the trapped cycle, the main paper states:

However, M' is required to act outside the support of M by Definition 5.1 (it must be a counter-modification accessible to bounded observers who cannot access the committed region).

This restriction is not in Definition 5.1, which only specifies that M is a bounded modification of diameter D . The justification given — that M' must be accessible to observers who cannot access "the committed region" — refers to the committed region as an already-defined structure, but the committed region is precisely what Theorem 5.2 is in the process of establishing.

The argument is close to circular: to prove that cycle-breaking produces local irreversibility, it assumes that post- M , the region containing the broken cycle is committed and therefore inaccessible to M' .

6.2 The Repair: Irreversibility as Post-Commitment Indistinguishability

The correct formulation of local irreversibility, for the purposes of Theorem 5.2, is not "no bounded operation can undo M " but "no bounded operation can produce a state distinguishable from the post- M state in a way that recovers the pre- M alternatives."

Definition 6.2.1 (Local irreversibility by indistinguishability). *Let M be a bounded modification of the substrate G , transforming G into G_M . The modification M is locally irreversible if no bounded operation sequence A of diameter $\leq D$, applied to G_M , produces a state G_A such that the pre- M alternatives ψ_A and ψ_B are distinguishable in G_A by any bounded local observation.*

This is a post-hoc criterion: it does not require M to be literally un-doable, only that the distinction between the pre- M alternatives cannot be recovered by any bounded operation.

Theorem 6.2.2 (Cycle-trapping by indistinguishability). *Let G contain a cycle C with alternatives ψ_A and ψ_B encoded as distinct flux states around C . Let M be the bounded modification removing one edge e of C . Then no bounded operation Λ of diameter $\leq D$ applied to G_M produces a state in which ψ_A and ψ_B are distinguishable by any bounded local observation.*

Proof. The alternatives ψ_A and ψ_B differ in flux around C — a topological invariant of the cycle C as a loop in the substrate. After M removes edge e , the loop C is no longer present in G_M : the graph G_M contains the vertices and remaining edges of C but no cycle traversing the original edge e .

The flux distinction between ψ_A and ψ_B was carried by the integral of a gauge connection, or an analogous cycle-dependent quantity, along C . With C broken, no cycle of the same homology class exists in G_M at diameter $\leq D$ from the support of M (by construction — the only such cycle was C itself, and it has been broken).

Consider any bounded operation Λ of diameter $\leq D$ acting on G_M . Two sub-cases:

Case (a): Λ adds no new edge replacing e . Then G_Λ still contains no cycle in the original homology class of C , and the flux distinction between ψ_A and ψ_B remains absent from G_Λ . No bounded local observation at diameter $\leq D$ can distinguish ψ_A from ψ_B in G_Λ .

Case (b): Λ adds a new edge e' creating a new cycle C' . Then C' is a cycle with different homology or geometry from C . The flux distinction originally encoded in C was specific to the homotopy class of C as a loop in the original substrate. A new cycle C' in a different homotopy class carries flux of a different invariant: the integral of the connection along C' is a function of the gauge data and the specific path C' , not the original path C .

For Λ to produce a G_Λ in which ψ_A and ψ_B become distinguishable again, the new cycle C' would have to carry the original flux distinction — that is, C' would have to be homotopic (in the relevant sense) to C in the original substrate. But C was the cycle in G containing edge e ; the only homotopy class containing e as a non-trivial loop representative was represented by C itself. A new cycle C' formed by adding a different edge e' at bounded distance from the support of M has a different loop representative and represents a different homotopy class.

Therefore in Case (b), G_Λ contains a cycle but carries a flux distinction of a different invariant, not the original ψ_A - ψ_B distinction. Bounded local observation of this new flux does not recover the original alternative distinction.

In both cases, no bounded operation Λ of diameter $\leq D$ produces a G_Λ in which ψ_A and ψ_B are distinguishable by any bounded local observation. The modification M is locally irreversible by the criterion of Definition 6.2.1. ■

6.3 Why This Formulation Is the Right One

The indistinguishability formulation of local irreversibility captures the operational content of fact formation without any circular reference to "committed regions." It says: once the alternatives are no longer distinguishable by any bounded operation, the distinction between them is operationally lost. This is exactly what fact formation requires (Definition 2.0.1 properties 1, 2, and 4).

The formulation also aligns with physical irreversibility in thermodynamics: a broken egg does not become un-broken by time-reversal in the relevant sense not because the microstates cannot be reversed, but because the macroscopic distinguishability between "broken" and "unbroken" is no longer accessible to bounded observation after the broken state is established.

6.4 Consequence for the Main Paper

- Add Definition 6.2.1 to Section 5.1 of the main paper, immediately after Definition 5.1.
- Replace the final paragraphs of the Theorem 5.2 proof (the M' restriction argument) with the proof of Theorem 6.2.2 above.
- Add a Remark noting that Definition 6.2.1 is the operational formulation and Definition 5.1 is the structural formulation; the two are equivalent under the \mathfrak{F} conditions.

The main paper's conclusion — that $\beta_1 \geq 1$ is required for locally constructible irreversibility — is preserved. The argument is now non-circular and rests on distinguishability rather than on a restriction on M' that was not supported by definitions.

7. Integration with the Main Paper

The five repairs are local and can be inserted at the following points in the main paper:

Repair	Main paper section	Insertion type
R1 (R-elimination)	Theorem 10.3, R-subsection	Replace isotropy argument with interference argument
R2 (Two senses of irreversibility)	Theorem 6.2a, Theorem 10.0	Reframe as observational equivalence; add Remark 6.2d
R3 (Manifold emergence)	Lemma 5.6, Step 4	Add Definition (local homogeneity); replace Step 4 proof
R4 (Fact-determinacy)	Proposition 8.2, condition (ii)	Replace geometric-moduli argument with topological-isotopy argument
R5 (Distinguishability-based irreversibility)	Theorem 5.2, cycle-sufficiency	Replace M' -restriction argument with indistinguishability proof

No result of the main paper is overturned. The dependency table in Section 1.3 is preserved. The Master Uniqueness Theorem (Theorem 11.1) and the kill corollary (Corollary 11.3) are unchanged.

The falsifiability section (Section 12) of the main paper is *strengthened* by these repairs: R2 clarifies that the framework is compatible with standard thermodynamics and decoherence, so falsifiability condition F1 (locally reversible stable facts) is the correct structural falsifier rather than being confused with fundamental reversibility. R3 clarifies that the framework's commitment to a manifold substrate is explicit and testable — non-manifold singular substrates would constitute empirical refutation.

8. Conclusion

The VERSF uniqueness programme's Master Theorem is preserved under close referee scrutiny, but five sub-arguments require repair. This companion paper supplies the repairs:

- **R1** replaces a flawed isotropy-based \mathbb{R} -elimination with a clean interference-based kill.
- **R2** disambiguates two senses of irreversibility, preserving compatibility with standard thermodynamics and correctly locating the hidden-variable question at the empirical rather than structural level.
- **R3** closes the manifold-emergence gap via operational local homogeneity and a Myers–Steenrod-style reconstruction.
- **R4** resolves the $d = 2$ versus $d > 2$ fact-determinacy argument at the topological level, using Schoenflies and its higher-dimensional failures as the correct mathematical foundation.
- **R5** replaces a circular M' -restriction in the cycle-trapping proof with a non-circular distinguishability formulation of local irreversibility.

Each repair is self-contained and can be inserted directly into the main paper. The 13-step elimination chain, the complete structural partition of Section 7.5, and the Master Uniqueness Theorem all survive intact. In several cases (notably R2 and R3), the repaired version is stronger than the original, because it makes explicit structural commitments (compatibility with thermodynamics, operational local homogeneity) that were previously implicit.

The remaining open tasks identified in Section 13 of the main paper — extension to the gauge sector, ambient substrate dimension, cycle topology specification, quantitative coupling constants, non-minimal theories within \mathfrak{F} , and relaxation of admissibility conditions — are unaffected by these repairs. The structural core of the VERSF uniqueness result is secure; the work ahead is to extend that core to the full downstream phenomenology of the Standard Model and gravity.

References

Adler, S.L. (1995). *Quaternionic Quantum Mechanics and Quantum Fields*. Oxford University Press.

Brown, M. (1960). "A proof of the generalized Schoenflies theorem." *Bulletin of the American Mathematical Society* 66, 74–76.

Chiribella, G., D'Ariano, G.M., and Perinotti, P. (2011). "Informational derivation of quantum theory." *Physical Review A* 84, 012311.

Hardy, L. (2001). "Quantum theory from five reasonable axioms." arXiv:quant-ph/0101012.

Hatcher, A. (2002). *Algebraic Topology*. Cambridge University Press.

Kaiser, H., George, E.A., and Werner, S.A. (1984). "Neutron interferometric search for quaternions in quantum mechanics." *Physical Review A* 29, 2276–2279.

Kirby, R.C. and Siebenmann, L.C. (1977). *Foundational Essays on Topological Manifolds, Smoothings, and Triangulations*. Princeton University Press.

Masanes, L. and Müller, M.P. (2011). "A derivation of quantum theory from physical requirements." *New Journal of Physics* 13, 063001.

Mazur, B. (1961). "A note on some contractible 4-manifolds." *Annals of Mathematics* 73, 221–228.

Myers, S.B. and Steenrod, N.E. (1939). "The group of isometries of a Riemannian manifold." *Annals of Mathematics* 40, 400–416.

Peres, A. (1979). "Proposed test for complex versus quaternion quantum theory." *Physical Review Letters* 42, 683–686.

Taylor, K. (2024–2026). *VERSF Research Papers*. AIDA Institute. versf-eos.com.

Taylor, K. (2026). *Structural Uniqueness of Physical Law from Fact Formation Constraints*. [Main paper, this companion.]