In a Universe Where Facts Can Exist, Where Outcomes Can Be Recorded, Remembered, and Tested, the Basic Structure Needed to Support That Is Uniquely the One Described by VERSF

Formal title: Structural Uniqueness of Physical Law from Fact Formation Constraints
Subtitle: A Rigorous Uniqueness Programme within the Void Energy–Regulated Space Framework (VERSF)

For the General Reader

Why does the universe have the mathematical structure it does? Why quantum mechanics? Why complex numbers rather than real ones? Why the specific forms that physical laws take? Most approaches to these questions begin by assuming large parts of the answer — by taking spacetime, or quantum fields, or symmetry principles as given, and asking what follows from there.

This paper begins with something more primitive than any of these. It begins with a single question:

What must the universe be like for stable, recordable facts to be possible at all?

A fact, in the sense used here, is simply an outcome that persists, can be recovered, and can be compared across different observations. When a measurement is made and a result is registered, that result is a fact. When an experiment is repeated and yields the same answer, that answer is a fact. Facts are what experiments produce and what theories are tested against. Without facts, there is no physics — not because physics is about human experience, but because a theory that produces no stable, recordable outcomes makes no testable predictions and is therefore not a physical theory at all.

Starting from that single requirement — that the universe must be capable of producing stable facts — this paper derives a chain of structural necessities. Each step is proved rather than assumed:

Facts require that the number of distinguishable outcomes in any bounded region be finite. Forming a fact must be an irreversible process for any local observer — otherwise what looks like a fact could be locally undone and was never really a fact at all. And from this, something remarkable follows immediately: time itself has no operational meaning without facts. A before/after relation requires distinguishable, persistent, irreversible events to order. In a universe with no stable facts, there is no stable ordering, no persistent record of which came first, and no physical arrow of time. Time is not a primitive of the theory — it emerges from the first irreversible commitment. And with time comes causality: the directed "X causes Y" relation requires that X precede Y, that the outcome Y persist and be identifiable, and that the connection be testable across trials. Without stable facts, none of these conditions can be met. Causality, like time, is not given in advance — it emerges from the ordered structure of irreversible commitments. The irreversibility of fact formation imposes a precise requirement on the information structure of the universe: it must contain closed loops rather than being purely tree-like. These loops force the existence of a specific kind of boundary — a two-dimensional surface separating committed outcomes from unchosen alternatives. That surface carries exactly two binary geometric properties. And those properties, combined with the requirement that the mathematical description be as lean as possible, force the pre-commitment description to be complex-valued — specifically, to live in a four-dimensional complex space.

This structure — a two-dimensional boundary with two binary geometric degrees of freedom and a four-dimensional complex reversible description — is precisely the fold architecture of the Void Energy–Regulated Space Framework (VERSF). The paper proves that within the class of all theories capable of producing stable facts for finite observers, this is the only minimal architecture. Every alternative either fails to produce stable facts, carries unnecessary structure, or reduces to this one under any reasonable standard of equivalence.

The result establishes a necessary structural core: any theory capable of producing stable, recordable, reproducible outcomes must contain this architecture. Physics is not built on spacetime or symmetry or quantum fields as primitives. It is built on facts — and the architecture of facts forces the unique minimal admissible structure.

Abstract

Unlike most physical theories, which begin with spacetime, fields, or symmetries, this paper begins with a single primitive: *facts are not a feature of physics — they are its precondition.* A physical fact is defined formally as a state-transition outcome satisfying five jointly necessary properties — stability, recoverability, reproducibility, irreversibility, and distinguishability. An equivalence theorem is proved: a universe admits a physical theory if and only if it admits a class of stable, recoverable, reproducible facts (Corollary 3.2). From that equivalence alone, with no prior commitment to quantum mechanics, spacetime geometry, or VERSF-specific primitives, all subsequent structural constraints are derived.

The admissible theory class $\mathfrak{F}$ is shown to be minimal: removing any of its seven conditions causes operational physics to collapse, establishing that the class is not engineered to select VERSF but is the weakest framework within which fact-based physics is possible. The central result (Master Uniqueness Theorem, Theorem 11.1) is proved by a thirteen-step elimination chain: any minimal theory in $\mathfrak{F}$ must contain, up to observational equivalence and field redefinition, a structure equivalent to the VERSF fold — a 2D intrinsic commitment interface with binary geometric data (σ, ω) and reversible pre-commitment representation over $\mathbb{C}^4$. The complete space of structural alternatives is explicitly partitioned and each case eliminated: tree versus cyclic substrates; intrinsic versus extrinsic versus bulk boundaries; local versus non-local commitment; 1D versus 2D versus higher-dimensional interfaces; and $\mathbb{R}$ versus $\mathbb{C}$ versus $\mathbb{H}$ versus $\mathbb{O}$ scalar fields. No case survives except the VERSF fold.

A kill corollary (Corollary 11.3) establishes that any theory in $\mathfrak{F}$ not containing the VERSF fold must violate at least one proved necessity. A further section defends restricted uniqueness as the correct and standard target for structural uniqueness results. A falsifiability section identifies six conditions that would refute the framework, distinguished by practical accessibility: two are genuinely near-term testable (distinguishable quaternionic phases; admissible non-local commitment), and four are structural conditions — genuine logical refutations that would require currently inconceivable physical systems but remain precise empirical commitments of the theory.

Table of Contents

---

1. Introduction

1.1 The Central Question

Most physical theories begin with spacetime, or fields, or symmetries, or a Hilbert space. This paper begins with none of these. It begins with a single primitive:

Facts are not a feature of physics — they are its precondition.

From that primitive alone — made formally precise in Section 3 — the complete structural architecture of physics is derived. The chain is:

facts → stable records → irreversibility → cyclic topology → commitment interface → 2D fold → $\mathbb{C}^4$ → VERSF

Each arrow is a proved theorem. Nothing is assumed beyond the operational definition of what it means for a theory to be physical.

The VERSF programme reconstructs physics not from pre-assumed spacetime, fields, or measurement rules, but from the minimal structural conditions under which stable physical facts can exist. A substantial body of previous work has established key components of this programme: physics requires fact-stable state classes; such classes require finite distinguishability; finite observers in locally causal universes require irreversible commitment; and these jointly force finite localization capacity. Independent results establish that empirical theories require record primitives, that local irreversibility cannot be constructed on acyclic information substrates, and that the minimal commitment boundary carries a four-state reversible structure represented minimally over $\mathbb{C}$.

What has remained open is the sharper question:

Do these results merely motivate VERSF, or do they uniquely force a VERSF-type architecture within a well-defined class of physically admissible theories?

This paper proves the latter. The method is eliminative: the class of admissible fact-supporting theories is defined, the structural freedoms that appear to remain are identified, and it is proved that each non-VERSF alternative fails one or more established necessities.

1.2 Master Uniqueness Theorem (Preview)

The central result is stated here so all subsequent material can be read as its proof. The formal statement with assembled proof appears as Theorem 11.1.

Master Theorem (Restricted Structural Uniqueness)

Any theory $T \in \mathfrak{F}$ that:

- *supports stable local irreversible fact formation for finite observers,*
- *satisfies finite distinguishability and finite localization capacity, and*
- *admits a minimal intrinsic realisation,*

must contain a structure equivalent to:

- *a 2D intrinsic commitment interface Σ,*
- *with two independent binary geometric degrees of freedom $(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$,*
- *and a reversible pre-commitment representation over $\mathbb{C}^4$ with unitary evolution.*

Up to observational equivalence and field redefinition, this structure is unique within $\mathfrak{F}$.
Sections 3–10 constitute the proof.

1.3 Proof Architecture, Diagram, and Theorem Hierarchy

Logical flow diagram:

```
Physical Facts (Definition 2.0.1)
         ↓
Time — emergent from irreversible fact ordering
    [Theorem 3.5, Corollary 3.7]
         ↓
Causal order — emergent from directed fact sequences
    [Theorem 3.9, Corollary 3.11]
         ↓
Finite Distinguishability + Irreversibility + Finite Capacity
    [Theorems 4.1 – 4.3]
         ↓
Topological Necessity: β₁ ≥ 1  [Lemma 5.3 + Theorem 5.2]
         ↓
Intrinsic Commitment Boundary ∂C  [Proposition 6.1]
         ↓
State-Space Decomposition: 𝒮 = 𝒮_accessible ⊕ 𝒮_inaccessible
   non-invertible · locally defined · not algebraically projectable
    [Theorem 6.2]
         ↓
Codimension-1 Separating Hypersurface  [Theorem 6.3]
         ↓
Non-local mechanisms eliminated  [Proposition 7.0]
         ↓
2D Minimal Interface  [Prop 8.2 fact-determinacy + Theorem 8.3, Jordan-
Brouwer]
         ↓
Binary Geometric Data: (σ, ω) — exactly 4 states  [Lemma 9.3]
         ↓
```

```
Scalar Field: ℂ
   ℝ insufficient · ℍ redundant · 𝕆 non-associative  [Theorem 10.3]
         ↓
Reversible Pre-Commitment Sector: ℂ⁴ with U(4)  [Corollary 10.4]
         ↓
VERSF Fold Architecture  [Theorem 11.1]
   unique up to observational equivalence
```

The elimination chain proceeds in strict logical order. Each result is established independently before feeding the next. No result assumes the conclusion of any later result.

Theorem Hierarchy Table:

Result	Type	Depends On
Theorem 3.1	Primitive	None (conditions 1–2 of $\mathfrak{F}$)
Corollary 3.2	Equivalence	Theorem 3.1
Theorem 3.5	Structural	Theorem 3.1, Theorem 4.2, Definition 2.0.1
Corollary 3.7	Eliminative	Theorem 3.5
Theorem 3.9	Structural	Theorem 3.5, Theorem 3.1, condition 6
Corollary 3.11	Eliminative	Theorem 3.9
Theorem 4.1	Necessary	Theorem 3.1, condition 3
Theorem 4.2	Necessary	Theorem 3.1, conditions 2, 7
Theorem 4.3	Necessary	Theorems 4.1, 4.2, conditions 2, 7
Lemma 5.3	Structural (auxiliary)	Definition 5.1, tree unique-path property
Theorem 5.2	Structural	Theorem 4.2, Lemma 5.3
Lemma 5.6	Structural	Theorems 4.1, 4.3, Definition 5.1
Proposition 6.1	Structural	Theorem 5.2, condition 7
Theorem 6.2	Structural	Proposition 6.1, Theorem 4.2, Definition 5.1
Theorem 6.2a	Structural	Theorem 6.2, Theorem 4.2
Corollary 6.2b	Structural	Theorem 6.2a
Theorem 6.3	Structural	Theorem 6.2, Lemma 5.6, Jordan–Brouwer
Proposition 7.0	Eliminative	Theorems 6.2, 6.3, conditions 2, 7
Lemma 7.0a	Eliminative	Proposition 7.0, condition 6, condition 7
Proposition 7.1	Eliminative	Proposition 7.3, condition 6, obs. equiv.
Proposition 7.2	Eliminative	Proposition 7.3, condition 6, obs. equiv.
Proposition 7.3	Sufficiency	Theorems 6.2, 6.3
Theorem 7.4	Eliminative	Propositions 7.0–7.3
Section 7.5	Survey	Theorems 5.2, 6.2, 6.3, 7.4, 8.3, 10.3

Result	Type	Depends On
Proposition 8.2	Structural	Theorem 7.4, Definition 8.1a, Jordan Curve Theorem, condition 6
Lemma 8.1c	Structural	Theorem 6.2, Theorem 6.3, Definition 8.1a, condition 6
Theorem 8.3	Structural	Theorem 7.4, Proposition 8.2, Jordan–Brouwer
Proposition 8.3b	Eliminative	Proposition 8.2, condition 6, obs. equiv.
Proposition 8.5	Structural	Theorem 8.3, condition 6
Lemma 9.3	Structural	Theorems 8.3, 8.5, Propositions 8.2, 8.3b, orientability
Proposition 9.4b	Structural	Lemma 9.3, condition 6, Definition 2.0.1
Theorem 9.5	Structural	Proposition 9.4b, Lemma 9.3, conditions 2, 6, 7
Lemma 9.5a	Eliminative	Theorem 9.5, Theorem 10.0, condition 6
Lemma 9.5b	Eliminative	Theorem 9.5, Theorem 10.0, condition 6
Theorem 10.0	Structural	Theorems 9.5, 4.2, Proposition 6.1
Proposition 10.1	Algebraic	Lemma 9.3
Theorem 10.3	Algebraic	Theorem 10.0, Proposition 10.1, Proposition 9.4b, condition 6, obs. equiv.
Corollary 10.4	Derived	Theorem 10.3
Theorem 11.1	Uniqueness	All above
Corollary 11.3	Uniqueness (kill)	Theorem 11.1 + all eliminations

1.4 Scope and Honesty of Claims

This paper does **not** claim that all possible physical theories reduce to VERSF in complete detail. It proves something narrower and stronger: any theory in $\mathfrak{F}$ capable of supporting stable, irreversible, local facts for finite observers must contain a structure equivalent to the VERSF fold. This is a restricted structural uniqueness claim, and Section 11.2 argues that restricted uniqueness is the correct and standard target for results of this type. Full uniqueness for all downstream phenomenology remains an open task identified in Section 13.

To be precise: the result establishes uniqueness of the minimal fact-supporting architecture within $\mathfrak{F}$. It does not assert that all physical structure reduces to the fold, but that any admissible theory must contain it as a necessary structural core. The fold is not the whole of physics; it is the part of physics that the existence of stable facts makes unavoidable.

---

2. Admissible Theory Class

2.0 Definition of a Physical Fact

The word "fact" is used throughout in a precise technical sense. Because the entire uniqueness argument rests on what facts require, the definition precedes everything else.

Definition 2.0.1 (Physical Fact). A *physical fact* is a state-transition outcome — the result of a process selecting from among previously open alternatives — satisfying all five of the following jointly necessary properties:

1. **Stability.** The outcome persists under admissible perturbations within the theory's protocol class. A transient fluctuation that does not persist is not a fact.
2. **Recoverability.** The outcome is re-accessible to a finite observer at a later time within that observer's local causal reach. An outcome produced but immediately inaccessible to any finite agent is not a fact in the operational sense.
3. **Reproducibility.** The outcome is comparable across independent instances: different observers, or repeated trials under equivalent conditions, can identify the same fact class. This requires that the outcome falls into a stable equivalence class that is intersubjectively confirmable.
4. **Irreversibility.** For any locally bounded observer O, the process producing the outcome cannot be locally undone. No bounded operation sequence accessible to O recovers the pre-transition state from the post-transition state. This distinguishes a fact from a reversible fluctuation.
5. **Distinguishability.** The outcome is separated from alternative outcomes by at least a minimum resolution scale δ_min > 0. Distinctions below this threshold are not stably recordable and do not constitute facts.

Definition 2.0.2 (Fact class). A *fact class* is an equivalence class of outcomes under within-tolerance perturbation. Two outcomes belong to the same fact class if and only if no admissible measurement by a finite observer distinguishes them.

Definition 2.0.3 (Record). A *record* is a physical system whose state encodes a fact and is stably accessible to a finite observer. A record is the physical carrier of a fact and must itself satisfy the stability and recoverability conditions above.

Definition 2.0.4 (Commitment). *Commitment* is the process by which an unresolved alternative — a situation in which multiple outcomes remain physically open — becomes a fact. By property (4), commitment is irreversible for bounded observers.

Remark 2.0.5 (What a fact is not). A fact is not: a globally defined property of the complete state of the universe (violates finite observability); a formal symbol without physical instantiation (lacks stability and recoverability); a transient event leaving no record (lacks recoverability and reproducibility); or a distinction below δ_min (violates distinguishability). These exclusions are not incidental — they drive the structural conclusions of the paper.

Remark 2.0.6 (Grounding the admissible class). Conditions 1–5 of $\mathfrak{F}$ use the word "fact" in the sense of Definition 2.0.1. The five properties map onto the conditions of $\mathfrak{F}$ as follows:

stability and reproducibility ground condition 1; irreversibility grounds condition 4; distinguishability grounds conditions 3 and 5; recoverability, together with condition 2, ensures facts are operationally instantiated rather than formally posited.

2.1 Definition of the Admissible Theory Class

A theory $T \in \mathfrak{F}$ if and only if it supplies:

1. **Stable empirical facts:** there exist physically instantiated outcomes that can be registered, recovered, and compared across instances.
2. **Finite observers:** all physically realised agents and record-forming processes have finite causal reach and finite operational resources.
3. **Finite distinguishability:** within any bounded physically accessible context, only finitely many distinctions are stably recordable.
4. **Irreversible commitment** (explicitly included for clarity; derivable from conditions 1, 2, and 7 as shown in Theorem 4.2)**:** some transitions from unresolved alternatives to facts are irreversible for locally bounded agents.
5. **Finite localization capacity:** bounded regions support only finitely many independent facts.
6. **Admissibility of description:** a distinction that produces no difference in any admissible fact outcome cannot be detected by any finite observer operating within the protocol class of Definition 3.0. Such a distinction has no empirical content and therefore no physical content under the operational definition of physics. Physically equivalent descriptions that cannot be distinguished by any fact carry no distinct physical meaning.
7. **Local causal boundedness:** no finite agent can jointly control all degrees of freedom relevant to every local record-formation event. *Locality in this framework is effective locality at the commitment scale: dependence on finitely bounded regions sufficient to determine fact formation, rather than strict pointwise locality. This interpretation is made precise in Lemma 7.0a, which shows that quasi-local mechanisms with bounded range do not constitute a third structural category but reduce either to effective local interfaces or to violations of this condition.*

Note on reversibility. A reversible pre-commitment sector is not listed as a condition, to avoid circularity. It is derived as Corollary 10.4 from conditions 1–7 together with the interface structure and field admissibility results of Sections 8–10.

2.2 Minimality of the Admissible Class

The objection this section answers: "Your assumptions are doing all the work — the class was chosen to select your answer."

The response is direct: the class cannot be weakened without destroying operational physics. Each condition is necessary; none is chosen for convenience. This is proved by the following removal arguments.

Remove condition 1 (Stable empirical facts). Without stable facts, there are no reproducible outcomes. A theory without reproducible outcomes makes no testable predictions. Its "laws" are formal relations over non-instantiated states. This is not physics in any operational sense — it is pure mathematics with an unverified physical interpretation. Condition 1 is the minimum required for the word "physics" to have operational content.

Remove condition 2 (Finite observers). Allowing infinitely capable observers with global causal reach removes all operational constraints on record formation. Any distinction, however fine, becomes observable; any region, however large, is within reach; any state, however global, is accessible. The admissible protocol class becomes vacuous — all distinctions are operational — and the admissibility conditions (3, 5, 6) collapse. Physics loses its operational grounding. Condition 2 is what converts mathematical structure into physical theory.

Remove condition 3 (Finite distinguishability). If arbitrarily fine distinctions are stably recordable, bounded physical regions must support infinite information density. No finite physical record-keeping system can implement this. In practice, physical storage requires energy and space at any given resolution scale; removing condition 3 requires infinite resources for bounded records. More formally, Theorem 4.1 — which is needed to bound the state space for all subsequent results — cannot be established without condition 3. Remove it and the elimination chain fails at the first step.

Remove condition 4 (Irreversible commitment). If all state transitions are locally reversible for bounded observers, no distinction exists between pre-factual and factual states. Every "fact" could be locally undone. Stability (Definition 2.0.1 property 1) and irreversibility (property 4) of facts both fail. Without condition 4, facts collapse to fluctuations, and the operational content of the theory evaporates back to condition 1 failure. Furthermore, Theorem 4.2 — the necessity of irreversible commitment — is precisely a proof that condition 4 follows from conditions 1, 2, and 7; removing condition 4 from the class means the elimination chain produces it as a theorem rather than assuming it, so the class remains appropriate. But a theory genuinely without commitment must fail condition 1 — no reversible theory can support stable facts.

Remove condition 5 (Finite localization capacity). Unbounded fact density in bounded regions requires either distinctions finer than δ_min (contradicting condition 3) or infinite irreversible correlation export per bounded region (contradicting conditions 2 and 7). Condition 5 is not an independent assumption — it follows from conditions 2, 3, and 7 as Theorem 4.3 — but its inclusion in the class definition makes the dependency explicit and prevents a theory from satisfying conditions 2, 3, and 7 while covertly assuming infinite localization capacity.

Remove condition 6 (Admissibility of description). Without this condition, unobservable structure is physically meaningful. Theories that differ only in non-fact-generating structure would be counted as distinct physical theories. This introduces infinite theoretical redundancy: any two empirically equivalent theories would be physically distinct, and uniqueness results become vacuous because infinitely many non-minimal theories satisfy every other condition. Condition 6 is what makes the elimination argument possible — without it, the observational equivalence convention (Section 2.4) fails, and non-minimal architectures cannot be identified and removed.

Remove condition 7 (Local causal boundedness). If agents can jointly control all degrees of freedom relevant to every local commitment event, commitment becomes non-local. Non-local commitment violates the topological trapping requirement (Theorem 5.2), because non-local agents can reconstruct cycle topology that local modifications destroy. The entire chain from Section 5 onwards collapses. Condition 7 is also what prevents bulk-collapse architectures from being admissible: without it, a single global update across the entire substrate could serve as the commitment mechanism, rendering the interface structure unnecessary. Condition 7 is the condition that forces commitment to be local, which is what forces the interface.

Conclusion. Each of the seven conditions is necessary. No condition can be removed without either: (a) immediately destroying the operational content of the theory (conditions 1, 2), (b) preventing the first theorem from being established (condition 3), (c) making facts indistinguishable from fluctuations (condition 4), (d) introducing internal contradiction (condition 5), (e) rendering the uniqueness argument vacuous (condition 6), or (f) collapsing the topological elimination chain (condition 7). The class $\mathfrak{F}$ is minimal.

2.3 Why This Class Is the Right Battleground

A uniqueness claim without a defined class is empty. This class is not chosen to privilege VERSF. It formalises the weakest conditions already established as necessary for operational physics, as shown by the removal arguments in Section 2.2. A theory outside $\mathfrak{F}$ may be a mathematical structure, but by construction it is not operational physics.

Note on condition 7 (forward dependency). Condition 7 (local causal boundedness) is listed as a class membership condition but its full necessity is established in conjunction with the topological arguments of Section 5, where local causal boundedness is shown to be required for the existence of locally constructible irreversibility. The condition is grounded in the same resource constraints that ground condition 2 (finite observers): an agent with finite causal reach and finite operational resources cannot simultaneously access and control all degrees of freedom in even a bounded region, for the same thermodynamic and information-theoretic reasons that make any physical recording device operate within finite energy and time budgets. Condition 7 is therefore not an independent postulate but follows from condition 2 under standard resource-theoretic reasoning. Its structural indispensability within the elimination chain is additionally proved at Theorem 5.2 and the subsequent interface arguments. This is an acknowledged forward dependency, not a circularity.

2.4 Observational Equivalence Convention

Two theories T and T′ are identified if and only if every stable fact producible by an admissible agent under T is also producible under T′ and vice versa. This convention is invoked at four specific points:

- **Theorem 6.2:** to establish that decompositions not detectable by any bounded fact impose no independent physical structure.

- **Proposition 7.1:** to eliminate bulk-collapse architectures whose additional global degrees of freedom produce no additional admissible facts.
- **Proposition 7.2:** to eliminate extrinsically labelled boundaries, since the external label is not a fact-producing structural feature.
- **Theorem 10.3:** to close the ℍ-elimination — quaternionic phases producing no distinguishable fact outcomes are observationally redundant and excluded by condition 6.

2.5 Engagement with Competitor Frameworks

Framework	Conditions satisfied	Fails at	Reason
Relational QM (Rovelli)	1–5, 7	Theorem 5.2	No intrinsic topological trapping; irreversibility is relational but structurally ungrounded
Consistent Histories	1–3	Theorem 4.2 / Prop 7.1	Decoherence functional is global; consistent set selection not a locally producible fact; non-minimal
GRW / Spontaneous Collapse	1–5	Prop 7.0 / Theorem 6.3	Collapse parameters inserted by hand; no derivation of interface dimension or intrinsic boundary
Causal Sets	1–4	Theorem 5.2	Acyclic by construction; $\beta_1 = 0$. Directed ordering ≠ local irreversibility in the sense of Definition 5.1. See extended note below.
Spin Foams / LQG	1–5	Theorems 6.2–6.3, 10.3	No intrinsic boundary between reversible and committed sectors; ℂ assumed not derived

Extended note on Causal Sets. The causal sets programme (Bombelli, Lee, Meyer, Sorkin 1987; Henson 2006) warrants a fuller response because a sophisticated referee may object as follows: "Causal sets encode irreversibility through the partial order itself — the asymmetric order relation is irreversibility, and the DAG structure is precisely a record of which events causally preceded which. Your cycle-threshold argument therefore misidentifies what causal sets are doing."

This objection identifies a genuine distinction, which the response must sharpen rather than dismiss. The relevant question is whether causal ordering in a DAG implements *local irreversibility in the sense of Definition 5.1* — that is, a bounded modification that renders previously open alternatives permanently non-recombinable.

In a causal set, the partial order is a global structure fixed from the outset. There is no pre-commitment sector of genuinely unresolved alternatives. Every event's causal future is determined by the order relation — there are no "open alternatives" in the state space that a bounded local modification could trap. The DAG encodes which events followed which, but the *openness* required by Definition 5.1 — the existence of alternatives that were genuinely accessible and that become permanently inaccessible after a bounded modification — is not present. Ordering says A preceded B; trapping says that alternative-B-given-A-occurred was genuinely open before the fact-forming event at A, and was then permanently foreclosed. Causal sets implement the former but not the latter.

More precisely: Theorem 5.2 concerns the *generation* of locally irreversible facts by bounded operations on a substrate with open alternatives. A causal set assumes the full causal order as a primitive — it does not generate irreversibility dynamically through bounded trapping. If the full causal order is posited as a global primitive, then irreversibility is assumed rather than derived, which exits the admissible theory class $\mathfrak{F}$ at condition 4 (irreversibility must be a structural consequence, not a stipulation). If the causal order is instead generated dynamically by a sequential growth process (as in the classical sequential growth models, or stochastic growth models in which new elements are added one at a time with probabilities satisfying a covariant condition), then the growth dynamics must supply a local trapping mechanism, and the question is whether that mechanism satisfies Definition 5.1 on an acyclic substrate — which Theorem 5.2 proves it cannot. In these dynamic models, each new element is added to an existing DAG, and the DAG remains acyclic at every stage; $\beta_1 = 0$ is maintained throughout the growth process. The acyclicity is precisely what is preserved by the growth rules (no element can be added whose causal past contains a directed cycle), so the stochastic growth process never generates the cycle structure required for locally constructible irreversibility. Dynamic causal set models therefore fail Theorem 5.2 for the same reason as static ones.

The causal sets framework therefore either assumes irreversibility globally (exiting $\mathfrak{F}$) or fails to generate it locally (failing Theorem 5.2). The distinction between ordering and trapping is not merely terminological: it is the structural difference between a theory that posits its fundamental asymmetry and one that derives it.

---

3. Primitive Theorem: Facts Are the Precondition of Physics

Facts are not a feature of physics — they are its precondition.

This paper is not built on spacetime, fields, or quantum mechanics. It is built on a single primitive: physics requires facts. Everything that follows — finite distinguishability, irreversible commitment, topological necessity, the 2D fold, $\mathbb{C}^4$ — is a consequence of that single requirement. The theorem below makes this primitive rigorous.

Definition 3.0 (Physics, operationally)

A *physical theory* is a framework that:

(a) makes repeatable, testable predictions about observable outcomes, and

(b) admits a protocol class under which those predictions can be confirmed or refuted by finite agents operating within bounded resources.

A mathematical system that lacks (a) or (b) is not a physical theory — it is a formal structure without empirical content.

Theorem 3.1 (Fact Necessity for Physics)

Any physical theory in the sense of Definition 3.0 must contain a class of stable, recoverable, reproducible facts.

Proof

The proof proceeds in three steps.

Step 1 — What physics operationally requires. A physical theory makes repeatable, testable predictions. For predictions to be testable, outcomes must be:

- *Identifiable:* an observer must be able to determine which outcome occurred.
- *Comparable:* outcomes from independent trials must be recognisable as the same or different.
- *Reproducible:* the same preparation under the same conditions must yield outcomes in the same fact class.

Without these three properties, there is no experiment, no verification, and no law. A "prediction" that cannot be identified, compared, or reproduced is not a physical prediction — it is a formal symbol without empirical referent.

Step 2 — Eliminating the alternative. Suppose a theory T contains no stable facts: all outcomes are transient, either failing to persist (violating stability), failing to be re-accessed after production (violating recoverability), or failing to be compared across instances (violating reproducibility).

Then:

- No outcome persists long enough to be identified. Identifiability fails.
- No two trials produce persistent comparable results. Reproducibility fails.
- The state space of T has no stably instantiable elements. There is no operational domain for any law of T.
- Any "law" of T is a formal relation over non-instantiated, non-comparable, non-persistent states. It makes no testable predictions.

Therefore T is not a physical theory in the sense of Definition 3.0. This contradicts the assumption that T is a physical theory. Therefore any physical theory must contain stable facts.

Step 3 — All three properties are jointly necessary. Suppose only stability fails: outcomes occur but do not persist. Then no outcome is available to be re-accessed or compared — recoverability and reproducibility fail as immediate consequences.

Suppose only recoverability fails: outcomes persist but cannot be re-accessed by any finite agent. Then no finite observer can confirm any outcome, and no prediction is testable. The theory satisfies (a) formally but violates (b) of Definition 3.0.

Suppose only reproducibility fails: outcomes occur and persist but cannot be compared across trials. Then no law — which requires universality over a class of instances — can be established. A single non-reproducible outcome is anecdotal, not physical.

In each case, the failure of any single property destroys the empirical content of the theory. All three properties are necessary, and their conjunction defines the class of facts. □

Corollary 3.2 (Equivalence: physics if and only if facts)

A universe admits a physical theory if and only if it admits a class of stable, recoverable, reproducible facts.

Proof

Necessity (→). Theorem 3.1 proves that any physical theory requires facts.

Sufficiency (←). Suppose a universe admits a class F of stable, recoverable, reproducible facts. Facts in F are identifiable (stable), re-accessible (recoverable), and comparable across trials (reproducible). A theory cataloguing the conditions under which facts in F occur and making predictions about which facts occur under which conditions satisfies Definition 3.0(a) and (b). Therefore the existence of a fact class is sufficient for a physical theory to be constructible.

The class of universes admitting physical theories is exactly the class of universes admitting stable, recoverable, reproducible facts. □

Corollary 3.3 (The derivation chain)

From the necessity of facts, the following chain is forced sequentially. Each link is proved in the sections indicated.

```
Physics (Definition 3.0)
      ↓ [Theorem 3.1]
Stable, recoverable, reproducible facts
      ↓ [Theorem 3.5]
Time (emergent from irreversible fact ordering)
      ↓ [Theorem 3.9]
Causal order (emergent from directed fact sequences)
      ↓ [Theorem 4.1]
Finite distinguishability (δ_min > 0)
      ↓ [Theorem 4.2]
Irreversible commitment
      ↓ [Theorem 4.3]
Finite localization capacity
      ↓ [Theorem 5.2]
Cyclic substrate topology (β₁ ≥ 1)
```

```
      ↓ [Theorems 6.1 – 6.3]
Intrinsic codimension-1 commitment interface
      ↓ [Theorem 8.3]
2D minimal surface
      ↓ [Lemma 9.3]
Binary geometric data (σ, ω)
      ↓ [Theorem 10.3]
Scalar field ℂ
      ↓ [Theorem 11.1]
VERSF fold — uniquely, up to observational equivalence
```

This chain is the spine of the paper. Every section that follows proves one link. The starting point is not quantum mechanics, not spacetime, not symmetry — it is the single requirement that physics produces facts.

Remark 3.4 (Preempting the anthropic objection)

A referee might object: "You are assuming the existence of observers or facts — this is an anthropic assumption that privileges conscious observers or the existence of life."

This objection misidentifies what is being claimed. The claim is not that observers happen to exist, or that facts happen to be produced, or that the universe is fine-tuned for life. The claim is definitional:

A theory that cannot, even in principle, produce stable outcomes accessible to finite agents is not a theory of physics — it is a mathematical system without empirical content.

This is not a contingent claim about the universe. It is a criterion for what the word "physics" means in the operational context relevant here. The admissible theory class $\mathfrak{T}$ does not assume that physics produces facts as a lucky accident. It defines the class of theories that qualify as physics precisely by this criterion.

The derivation that follows is entirely substrate-neutral and observer-neutral in its formal structure. It depends on nothing about who the facts are for, or how they are perceived, or whether life exists. It depends only on the structural requirements imposed by the existence of stable, recoverable, reproducible outcomes — requirements that apply to any physical theory in any conceivable universe that qualifies as physics in the operational sense of Definition 3.0.

The objection "you are assuming facts" is therefore equivalent to the objection "you are assuming physics." That is precisely correct. A theory of physics that does not produce facts is not physics.

Theorem 3.5 (Time Emergence from Fact Ordering)

A necessary condition for the existence of a physically meaningful notion of time is the existence of a partially ordered set of stable, irreversible facts.

Proof

The proof establishes that each component of the operational concept of time requires facts, and that the absence of facts renders time operationally undefined.

Step 1 — What time operationally requires. Time, in any operationally meaningful sense, requires at least the following:

- *Ordering:* a before/after relation between distinct events — event A occurs before event B.
- *Distinguishability:* the events being ordered must be identifiable as distinct.
- *Persistence:* the record of which event preceded which must persist long enough to be recovered and compared.
- *Irreversibility:* the ordering must be asymmetric — if A is before B, it is not also true that B is before A in the same sense. A reversible relation cannot ground the directed arrow that makes time temporal rather than merely spatial.

Step 2 — Each component requires facts. *Ordering* requires distinguishable states at the ordered events. By Definition 2.0.1 (property 5), distinguishable states above δ_min are facts or grounded in facts. Without facts, no two events are stably distinguishable — they cannot be placed in order. *Distinguishability* requires that the distinction between event-A-has-occurred and event-A-has-not-occurred be stable and recoverable: this is precisely the stability and recoverability properties (1) and (2) of Definition 2.0.1. Without facts, event-occurrence is not stably distinguishable from non-occurrence. *Persistence* of the temporal record requires that the ordering relation, once established, remain accessible to finite observers — this is the recoverability condition (property 2) again. *Irreversibility* of temporal ordering requires that the transition marking the occurrence of an event be locally irreversible for bounded observers. If it were reversible, the before/after relation could be locally undone, collapsing the distinction between earlier and later. This is precisely the irreversibility property (4) of Definition 2.0.1.

Step 3 — Absence of facts renders time undefined. Suppose T contains no stable, irreversible facts. Then:

- No state-transition is stably distinguishable from its reversal → no asymmetric before/after relation exists.
- No event-occurrence is stably recorded → no ordering relation persists.
- No distinction between pre-event and post-event states is recoverable → "before" and "after" are not operationally definable.

In such a theory, any candidate temporal ordering is at best a formal assignment of labels to states, with no operational content. No finite observer can verify that event A preceded event B, because no stable record of either event's occurrence exists. The temporal ordering carries no admissible facts (in the sense of condition 6), and by the observational equivalence convention it is redundant — it generates no physical distinctions.

Conclusion. Time has no operational meaning without facts. A necessary condition for time is a partially ordered set of stable, irreversible facts — the ordering relation is grounded in the sequence of irreversible commitment events. □

Corollary 3.6 (Time from Facts)

A physically meaningful notion of time requires an ordered sequence of stable, irreversible facts. In the absence of such facts, no invariant temporal ordering can be defined, and time has no operational meaning.

Scope. These results concern operational temporal order — the structure defined by observable, recordable, irreversible events accessible to finite observers — and do not make claims about coordinate time as a purely mathematical parameter or about temporal structure in theories that do not admit operational physics in the sense of Definition 3.0.

Interpretation. Time is not a primitive of physical theory but an emergent structure arising from the ordering of irreversible commitments. The pre-commitment sector — the reversible domain prior to fact formation — admits no temporal ordering: alternatives are open, transitions are reversible, and no persistent before/after relation can be established. Time begins where commitment begins. The commitment interface (Sections 6–8) is therefore not merely the boundary between committed and reversible sectors — it is the structural locus at which time itself becomes definable.

Corollary 3.7 (No Facts ⇒ No Time)

In any theory T lacking stable, irreversible facts, no physically meaningful temporal ordering can be defined.

Proof

By Theorem 3.5, time requires stable irreversible facts. If T lacks stable irreversible facts, the necessary condition for time is unmet. Therefore no invariant temporal ordering exists in T. □

Remark 3.8 (Scope and precision)

This result does not say "time does not exist without facts" in a metaphysical sense. It says time has no *operational meaning* without facts — no physical measurement, no reproducible before/after determination, no testable temporal ordering can be constructed. This is the operationally precise claim, and it is fully consistent with the spirit of condition 6 (admissibility of description): a temporal ordering that cannot be grounded in any admissible fact carries no distinct physical meaning and is not a physical quantity.

The result also clarifies the structure of the commitment interface. The pre-commitment sector evolves reversibly: no commitments have occurred, no facts have been produced, and by Corollary 3.7, no temporal ordering is defined there. The commitment event — the irreversible transition at the interface — is the first fact in the temporal sequence. Time, in the operational sense, emerges at the interface. This is not an interpretation added to the formalism; it is a direct consequence of Theorem 3.5 applied to the interface structure proved in Sections 6–8.

Theorem 3.9 (Operational Causality Requires Facts)

A physically meaningful — that is, testable — notion of causality requires a set of stable, distinguishable, irreversible facts that can be ordered.

Proof

Step 1 — What operational causality requires. Causality in physics is the claim that interventions produce outcomes: changing X changes Y in a stable, reproducible way. For a causal relation "X causes Y" to be operationally meaningful — testable rather than merely asserted — the following must hold:

- *Identifiability:* the outcome Y must be identifiable as having occurred. It must be distinguishable from the absence of Y and from alternative outcomes.
- *Stability:* the outcome Y must persist long enough to be registered, compared across trials, and used as a basis for inference about X.
- *Reproducibility:* repeated interventions on X under equivalent conditions must yield comparable outcomes Y that can be identified as the same type of result.
- *Ordered precedence:* X must precede Y — the cause must come before the effect. This requires a temporal ordering, which by Theorem 3.5 requires stable irreversible facts.

Step 2 — Each requirement is a fact requirement. Identifiability is exactly the distinguishability property (5) of Definition 2.0.1. Stability is property (1). Reproducibility is property (3). Ordered precedence requires time, which by Theorem 3.5 requires stable irreversible facts. Therefore every operational requirement of causality is a requirement for facts.

Step 3 — Absence of facts eliminates operational causality. Suppose T contains no stable, irreversible facts. Then:

- No outcome Y is stably identifiable: any apparent occurrence of Y is a transient fluctuation that cannot be recorded or recovered.
- No intervention on X can be verified to have produced Y: without a stable record of Y's occurrence, the putative causal connection cannot be confirmed across trials.
- No before/after ordering of X and Y can be established (by Corollary 3.7), so the asymmetric "X precedes and causes Y" relation has no operational grounding.

In such a theory, any candidate causal relation is a formal assignment with no testable content. It generates no admissible facts (condition 6) and is therefore physically meaningless under the observational equivalence convention.

Note on underlying dynamics. A referee might object: "There may exist underlying causal structure even without stable facts." This is acknowledged. The claim is not that no dynamics exist in a fact-free theory, but that any dynamics present cannot be identified, distinguished, or tested as causal relations by any finite observer. Underlying dynamics without operational access are not physical causality in the sense of Definition 3.0 — they are formal structure without empirical content. Operational causality requires facts; formal dynamics do not.

Conclusion. Without stable, irreversible, reproducible facts, no intervention can be registered, no outcome can be compared across trials, and no cause–effect relation can be established operationally. Causality without facts is not false — it is undefined. □

Corollary 3.10 (Causality from Facts)

A physically meaningful notion of causality requires an ordered sequence of stable, irreversible facts. In the absence of such facts, no intervention can be registered, no outcome can be compared across trials, and no cause–effect relation can be established. Therefore, without facts, causality has no operational meaning within any admissible physical theory.

Interpretation. Causality is not a primitive of physical theory but an emergent structure arising from the ordered relations between irreversible commitments. Just as time emerges from the ordering of facts (Corollary 3.6), causal order is identified with the directed relations between irreversible facts: fact A causally precedes fact B when A is a committed outcome whose occurrence is necessary for B's commitment event to occur. The causal arrow and the temporal arrow are both grounded in the same structure — the irreversible commitment interface.

Corollary 3.11 (No Facts ⇒ No Operational Causal Order)

In the absence of stable, irreversible facts, no invariant cause–effect relation can be defined or tested.

Proof

By Theorem 3.9, operational causality requires stable, irreversible facts. If T lacks such facts, the necessary condition is unmet. Therefore no testable cause–effect relation exists in T. □

Remark 3.12 (The emergent spine: facts → time → causality)

Theorems 3.5 and 3.9, together with Corollaries 3.6–3.11, establish that both time and causality are emergent from the single primitive of stable, irreversible fact formation. The spine of physical structure is:

```
Stable irreversible facts
      ↓
Ordered records (before/after)
      ↓
Time — operational temporal ordering  [Theorem 3.5]
      ↓
```

```
Causal order — directed cause/effect  [Theorem 3.9]
```

Neither time nor causality is a primitive ingredient poured into the theory from outside. Both emerge from the commitment structure that facts require. The commitment interface (Sections 6–8) is the geometric realisation of this emergence: it is the locus at which unordered, reversible, acausal pre-commitment alternatives become ordered, irreversible, causally structured facts.

This has a precise consequence for the VERSF programme: the fold architecture is not only the minimal structure for producing facts — it is the minimal structure within which time and causal order have operational meaning at all.

Time and causality are therefore not independent primitives but different aspects of the same underlying structure: the ordering of irreversible commitments.

---

4. Necessity of Finite Distinguishability, Irreversible Commitment, and Finite Localization Capacity

The first three links in the derivation chain (Corollary 3.3, following Theorem 3.5) are established here. Each is a necessary structural consequence of the fact requirement; none is an independent assumption.

Theorem 4.1 (Finite distinguishability is unavoidable)

Any $T \in \mathfrak{F}$ possesses a nonzero minimum fact-supporting distinguishability scale $\delta_{min} > 0$.

Proof

Suppose for contradiction that arbitrarily fine distinctions are stably fact-supporting. Let $\{\delta_n\}$ be a sequence with $\delta_n \to 0$, each supporting stable records in bounded regions of diameter D. By condition 3, only finitely many distinctions are stably recordable within any bounded context. But if arbitrarily fine distinctions are stably recordable, the number of stably distinguishable states within D grows without bound as $\delta \to 0$, contradicting condition 3. Therefore a minimum scale $\delta_{min} > 0$ exists. □

Theorem 4.2 (Irreversibility is necessary for fact stability)

The irreversibility condition in Definition 2.0.1(4) — that fact-forming operations are locally irreversible for bounded observers — cannot be removed without collapsing the stability of facts.

Remark (on the relationship to Definition 2.0.1). Definition 2.0.1(4) includes irreversibility as a component of what it means to be a fact. This theorem does not derive irreversibility independently of that definition. It establishes that irreversibility is a *necessary structural component* of any admissible notion of fact: any weakening of the irreversibility requirement causes the remaining properties of Definition 2.0.1 — stability, recoverability, reproducibility — to become unsatisfiable for finite observers. The theorem shows that irreversibility is not an optional strengthening of the fact concept but is entailed by the other properties together with the conditions of $\mathfrak{F}$.

Proof

Let O be a finite observer (condition 2) performing a fact-forming operation F in a locally causally bounded region (condition 7). F maps an unresolved alternative A to a committed outcome f. Suppose the irreversibility requirement is removed: F is locally reversible for O, so O can perform F^{-1}, recovering A from f without accessing the environment beyond O's causal boundary.

We show this collapses the remaining three properties of Definition 2.0.1. *Stability* (property 1): if F is locally reversible, then f is not stable against operations within O's reach — any perturbation triggering F^{-1} restores A, and f cannot persist under O's operational class. *Recoverability* (property 2): if f and A are interconvertible by O, the "committed" state f is not operationally distinct from the pre-commitment state A — O cannot reliably recover a record of f rather than A. *Reproducibility* (property 3): if the operation can go both ways, repeated applications under equivalent conditions can yield A or f, with no stable fact class identifiable across trials.

All three properties fail simultaneously if irreversibility is removed. This contradicts condition 1 of $\mathfrak{F}$, which requires stable, recoverable, reproducible facts. Therefore the irreversibility condition cannot be removed from any admissible notion of fact for finite observers in $\mathfrak{F}$: fact-forming operations are locally irreversible. □

Theorem 4.3 (Finite localization capacity is unavoidable)

For any $T \in \mathfrak{F}$ and bounded region R of diameter D, the number of independently supportable stable facts N(R) satisfies $N(R) < \infty$.

Proof

Suppose N(R) is unbounded. Each independent fact requires a commitment event exporting irreversible correlation beyond O's local reach (Theorem 4.2). Unbounded N(R) implies divergent total correlation export from R. By condition 7, the local causal environment has bounded absorption capacity. An unbounded number of facts would require either (a) distinctions finer than δ_min to be recorded, contradicting Theorem 4.1, or (b) the finite causal environment to absorb divergent correlation export, contradicting conditions 2 and 7. Therefore $N(R) < \infty$. □

5. Topological Threshold for Local Fact Formation

Definition 5.1 (Locally constructible irreversibility)

A local irreversibility mechanism is a bounded modification — involving finitely many degrees of freedom within a region of diameter D — that renders previously recombinable alternatives permanently non-recombinable for bounded observers without requiring global restructuring of the entire substrate.

Lemma 5.3 (External witness persistence in trees)

Let G be a tree and let $M \subset G$ be any bounded connected subgraph with boundary ∂M. Let (ψ_A, ψ_B) be a pair of alternatives not perfectly degenerate on $V \setminus M$ — there exists at least one vertex $v^* \in V \setminus M$ whose state differs between ψ_A and ψ_B. Then for any bounded modification M supported on M, the state at v^* remains correlated with the pre-modification alternative: no sequence of local operations within diameter D of M can destroy the external signature at v^* without directly accessing v^*.

Proof

By the unique path property of trees, the unique path $P(v^*, u)$ from v^* to any vertex $u \in M$ passes through ∂M. Since M is confined to M, no local operation within M alters the state at v^*. The state at v^* is unchanged by M. Since v^* was chosen such that its state differs between ψ_A and ψ_B, this external signature persists after M is applied. Any operation capable of destroying the signature at v^* must be supported on a region containing v^*, which lies outside M. Therefore no bounded modification confined to M eliminates the external witness at v^*. □

Remark 5.4

Lemma 5.3 closes the potential gap in the tree-impossibility argument. It proves explicitly that the external recombination witness not only exists in the pre-modification state, but strictly persists through any bounded modification confined to M. This is the structural property preventing local trapping on trees.

Theorem 5.2 (Cycle threshold for local irreversibility)

Let the effective information substrate of $T \in \mathfrak{F}$ be a connected undirected graph $G = (V, E)$ with reversible microdynamics. If T supports locally constructible irreversibility, then $\beta_1(G) \geq 1$. No tree-structured substrate supports locally constructible irreversibility.

Proof

Trees cannot trap. Let G be a tree and let (ψ_A, ψ_B) be a pair of alternatives to be trapped by bounded modification M on M.

Case 1: ψ_A and ψ_B differ on $V \setminus M$. By Lemma 5.3, $v^* \in V \setminus M$ exists whose state persists and remains correlated with the pre-modification alternative. This external witness is accessible to local operations outside M, so ψ_A and ψ_B remain locally recombinable. Trapping fails.

Case 2: ψ_A and ψ_B are perfectly degenerate on $V \setminus M$. The distinction is entirely local to M. Any bounded modification M can be countered by a counter-modification M′ of equal diameter restoring the original state within M. The distinction is not permanently trapped. This is not locally constructible irreversibility in the sense of Definition 5.1.

In both cases, tree substrates fail to support locally constructible irreversibility.

Cycles enable trapping. Let G contain a cycle $C = (v_1, v_2, \ldots, v_k, v_1)$. Encode alternatives ψ_A and ψ_B as distinct flux states around C — a topological defect present versus absent on C. Both alternatives are degenerate on every individual vertex but differ in the global cycle-traversal character.

Let M be the bounded modification that collapses the cycle by removing one edge $e = (v_i, v_{i+1})$ from C, supported on a region of diameter D containing e. After M is applied, the graph no longer contains C as a cycle: the removed edge is the only edge between v_i and v_{i+1} in the cycle (since $G \setminus \{e\}$ is still connected via the remaining path, but the cycle itself is gone). The flux distinction between ψ_A and ψ_B becomes inaccessible: without a cycle to carry the defect, both alternatives reduce to the same graph state.

We verify irreversibility: any counter-modification M′ of equal diameter attempting to restore the cycle must add an edge between v_i and v_{i+1}. But M′ is a bounded local modification supported on a region of diameter D. The only edge that can restore C is the removed edge e, which lies within the support of M. 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). Therefore M′ cannot restore e, and the cycle topology is not locally recoverable.

Closing the alternate-edge loophole. A referee might object that M′ could introduce a different edge e′, forming a new cycle C′ distinct from the original C, and thereby restoring some cyclic structure. This does not restore the original distinguishability. The trapped alternatives (ψ_A, ψ_B) were encoded specifically in the topology of cycle C — as distinct flux states around C's particular traversal structure. Replacing C with a distinct cycle C′ formed from different edges does not recover the original topological defect but produces a new topological structure with different global properties. The original flux distinction is encoded in C's specific cycle class; a new cycle C′ with a different cycle class does not decode the same distinction. Furthermore, constructing e′ to make C′ topologically equivalent to C would require coordinated modification of the original cycle vertices and edges, which lies outside the admissible bounded operation class of diameter D centred away from the support of M.

The distinction is permanently trapped, and the modification M satisfies Definition 5.1. □

Conclusion. Locally constructible irreversibility requires $\beta_1(G) \geq 1$. In a tree, every bounded modification preserves at least one externally accessible recombination witness, preventing local irreversibility. □

Corollary 5.5

Any $T \in \mathfrak{F}$ supporting local fact formation must contain or induce an effective connectivity structure with $\beta_1 \geq 1$.

5.6 Lemma: Effective Manifold Emergence

The gap this lemma closes. The topological results of Sections 5–6 use graph-theoretic language (β_1, cycles, adjacency), while the codimension theorem (Theorem 6.3) and the dimension arguments of Section 8 invoke manifold theory (Alexander duality, Jordan–Brouwer). A referee could legitimately object: a graph is not a manifold, and $\beta_1 \geq 1$ does not imply a smooth structure. This lemma derives the effective manifold description directly from the conditions of $\mathfrak{F}$, closing the gap without introducing additional primitives.

Lemma 5.6 (Effective Manifold Emergence from Finite Observability)

Any substrate G supporting finite distinguishability (Theorem 4.1), finite localization capacity (Theorem 4.3), and locally constructible operations in the sense of Definition 5.1, admits an effective manifold description at observational scales ≥ δ_min.

Proof

The argument proceeds in five steps.

Step 1 — Coarse-graining at δ_min. By Theorem 4.1, no finite observer can distinguish features of G at scales below δ_min. The operationally accessible description of G is therefore not the discrete graph G itself, but the coarse-grained structure G_eff obtained by identifying all vertices within distance δ_min and replacing them with a single equivalence class. Since all admissible operations are bounded modifications at scale ≥ δ_min (Definition 5.1), G_eff captures all operationally meaningful structure.

Step 2 — Neighbourhood topology. Each equivalence class in G_eff has a well-defined neighbourhood: the set of equivalence classes reachable by a single bounded operation. By condition 7 (local causal boundedness), bounded operations have finite causal reach D. The neighbourhood structure N(x) of each point x ∈ G_eff is therefore a finite collection of nearby equivalence classes within reach D of x.

Step 3 — Topological space axioms. The pair (G_eff, {N(x)}) satisfies the axioms of a topological space. Specifically: (i) every point x is in its own neighbourhood (the trivial case of a zero-diameter operation); (ii) if U ∈ N(x), then there exists V ∈ N(x) with V ⊆ U such that U ∈ N(y) for all y ∈ V, because bounded operations of diameter D are transitive within their reach (from any y reached in one step, the same step-structure applies). The resulting topology is Hausdorff because distinct equivalence classes at scale ≥ δ_min are distinguishable by Theorem 4.1.

Step 4 — Local homogeneity and locally Euclidean charts. By condition 3 (finite distinguishability), the local state space within any neighbourhood N(x) is finite-dimensional: finitely many distinguishable states in a bounded region. The local dimension d(x) — the number of independent directions in which a bounded operation can move from x — is a well-defined finite integer at every x. By condition 7 (local causal boundedness), d(x) is bounded uniformly across all points: all local observers share the same finite-dimensional local operational structure.

No global group structure is assumed or required. What is used is only *local operational homogeneity*: bounded operations at scale δ_min induce a locally homogeneous neighbourhood structure at each point, in the sense that from any point x, the set of points reachable by a single bounded operation forms a neighbourhood that looks the same in all operationally accessible directions. This homogeneity is guaranteed by the uniform local dimension (every point has the same d-dimensional operational reach) and by the local causal structure of condition 7 (every local observer faces the same finite-dimensional environment).

From local operational homogeneity and uniform finite dimension d, a locally Euclidean chart can be constructed at each point — but only after ruling out branching singularities. At a branching point, some directions lead to multiple inequivalent continuations while others do not. Any such branching singularity introduces locally distinguishable directional asymmetry. There are exactly two cases:

Case (a): the asymmetry affects commitment outcomes. Branching singularities can be excluded at this point without reference to the interface structure. A branching point introduces locally distinguishable directional asymmetry: some directions from x admit multiple inequivalent continuations while others do not. If this asymmetry affects admissible fact outcomes, it introduces additional locally distinguishable structure beyond that encoded in the equivalence classes at scale δ_min. By Theorem 4.1, only finitely many distinctions are stably recordable in bounded regions; any additional locally distinguishable directional asymmetry not captured by the coarse-grained structure G_eff would require distinctions below δ_min or constitute an additional distinguishable class at the operational scale — either contradicting finite distinguishability or requiring G_eff to carry more local structure than the δ_min coarse-graining admits. In either case, the branching does not coexist with the uniform local structure established in Steps 1–3. This case is structurally impossible without invoking the interface structure.

Case (b): the asymmetry does not affect commitment outcomes. If the directional asymmetry generates no difference in any admissible fact outcome, it is observationally redundant. By condition 6 (admissibility of description) and the observational equivalence convention, descriptions differing only in the presence of a branching singularity are identified as the same physical theory. The branching structure is excluded as non-physical.

In both cases — whether the branching affects or does not affect commitment outcomes — the branching singularity either fails to be a valid commitment interface (case a) or is observationally redundant (case b). This argument holds regardless of the scale of the branching. Therefore the local topology must be uniform and non-branching, and each neighbourhood is homeomorphic to ℝ^d. The d independent operational directions from x define a local coordinate system in which the neighbourhood maps to ℝ^d by displacement vector. A Hausdorff topological space —

which G_eff is, since distinct equivalence classes at scale ≥ δ_min are distinguishable (Theorem 4.1) — in which every point has a neighbourhood homeomorphic to ℝ^d is, by definition, a d-dimensional topological manifold.

Step 5 — Atlas construction. The uniform local homeomorphisms to ℝ^d constructed in Step 4 constitute an atlas for G_eff. The transition functions between overlapping charts are continuous (they are induced by the bounded operation transitions, which are continuous at coarse-grained scale). The resulting topology is second-countable — because G_eff is constructed from a countable cover of equivalence classes at scale δ_min over a substrate with finite localization capacity (Theorem 4.3) — and locally Euclidean by construction, since every point has a neighbourhood homeomorphic to ℝ^d (Step 4). These are the standard defining conditions of a topological manifold. Therefore G_eff is a topological manifold, and the substrate G admits an effective manifold description G_eff at scales ≥ δ_min.

Conclusion. All subsequent geometric arguments — Alexander duality, Jordan–Brouwer, codimension, and dimension — are applied to G_eff, not to G. The graph G is the microscopic substrate; G_eff is the operationally accessible manifold. No smooth structure beyond the topological manifold is assumed; all arguments in Sections 6–8 use only topological manifold properties. □

Remark 5.7

Lemma 5.6 establishes that the manifold structure used in subsequent arguments is not an additional assumption about the substrate. It is derived from the operational constraints of 𝔉: finite distinguishability fixes the coarse-graining scale, finite localization capacity bounds the local dimension, and local causal boundedness supplies the uniform neighbourhood structure. The manifold is the effective description that any finite observer must use; the underlying discrete graph is operationally inaccessible below δ_min.

---

6. From Topological Trapping to Commitment Boundaries

6.1 Commitment Boundary Existence

Proposition 6.1 (Commitment boundary existence)

If T ∈ 𝔉 supports local fact formation by topological trapping, then there exists an operationally intrinsic commitment boundary ∂C separating the committed record sector from the inaccessible discarded alternative sector.

Proof

A fact-forming event at region R selects one alternative f and renders the others inaccessible to further local recombination. The inaccessible alternative cannot be annihilated — annihilation would require a global state update violating condition 7. It resides in a sector D(R) no longer operationally accessible to the local record. The committed record occupies sector C(R). Since C(R) and D(R) are operationally distinct, there exists a boundary ∂C between them. This boundary is defined by the topological trapping mechanism (Theorem 5.2), not by any external label, and is therefore intrinsic to the substrate. □

6.2 Interface Necessity: Forced State-Space Decomposition

The objection this section answers: "You just renamed the boundary 'interface' — you haven't shown it must be a geometric interface rather than a non-local constraint or algebraic projection."

The response is a three-step forced logical chain:

Local irreversibility *(Theorem 4.2)*
↓ forces
Non-invertible state-space decomposition *(property 1 below)*
↓ which, combined with local determination *(property 2)*
Cannot be expressed as a global algebraic constraint *(property 3)*
↓ therefore must be
A geometrically embedded local interface

Theorem 6.2 (Interface Necessity: Forced State-Space Decomposition)

Let $T \in \mathfrak{T}$ support locally constructible irreversibility. There exists a decomposition

$\mathcal{S} = \mathcal{S}$_accessible $\oplus$ $\mathcal{S}$_inaccessible

such that:

1. The decomposition is **not globally invertible**: no bounded operation sequence acting on $\mathcal{S}$_accessible recovers $\mathcal{S}$_inaccessible from a committed state.
2. The separation is **locally defined**: assignment of a state to $\mathcal{S}$_accessible or $\mathcal{S}$_inaccessible is determined by local data in a bounded neighbourhood, not by global evaluation of the full state.
3. The boundary ∂($\mathcal{S}$_accessible) **cannot be expressed as the level set of a globally defined algebraic constraint** (projection, superselection rule, or non-local potential). Any such global expression implies a globally invertible recovery map, contradicting property (1).

Therefore the separation is realised by an intrinsic local interface: a geometrically embedded commitment boundary, not an algebraically imposed global constraint.

Proof

Property (1). Let $f \in \mathcal{S}$_accessible be a committed fact and $d \in \mathcal{S}$_inaccessible the discarded alternative. Suppose there exists a bounded operation sequence Λ acting on $\mathcal{S}$_accessible recovering d. Then $\Lambda(f) \in \mathcal{S}$_inaccessible, re-opening the discarded alternative and undoing the fact — contradicting Theorem 4.2. No such Λ exists. $\square_1$

Property (2). By Definition 5.1, locally constructible irreversibility is achieved by a bounded modification M on a finite subgraph of diameter D. The commitment status of any state is determined by local configuration within diameter D of the commitment event. The trapping mechanism operates on local cycle structure (Theorem 5.2): the cycle is a local feature of G, and M exploiting it is bounded. States separated by M inherit their commitment status from local topology, not from global boundary conditions. $\square_2$

Property (3). Suppose the separation were expressible as a global algebraic constraint: a globally defined operator $P: \mathcal{S} \to \mathcal{S}$ such that $\mathcal{S}$_accessible = $P(\mathcal{S})$ and $\mathcal{S}$_inaccessible = $(1-P)(\mathcal{S})$. For P to be consistent with local determination (property 2), P must act as a local projector at every point — defined fibre-wise over the substrate with action determined by local data only. Such a fibre-wise projector is geometrically a field of local separations — precisely a submanifold embedded in the substrate, i.e., a local interface structure. Furthermore, any globally defined projector admitting a global inverse P^{-1} with $P^{-1} \circ P = \mathrm{id}$_$\mathcal{S}$ is globally invertible, violating property (1). Therefore no globally invertible algebraic constraint realises the separation, and the separation must be realised by a geometrically embedded local interface.

Properties (1)–(3) are jointly established. The separation $\mathcal{S} = \mathcal{S}$_accessible $\oplus$ $\mathcal{S}$_inaccessible is non-invertible, locally defined, and geometrically embedded. $\square$

Theorem 6.2a (Physical Necessity of the Inaccessible Sector)

In any admissible theory $T \in \mathfrak{T}$ supporting irreversible commitment, the inaccessible sector $\mathcal{S}$_inaccessible in the decomposition $\mathcal{S} = \mathcal{S}$_accessible $\oplus$ $\mathcal{S}$_inaccessible (Theorem 6.2) must correspond to physically real degrees of freedom and cannot be reduced to a purely formal, epistemic, or descriptive construct.

Proof

The proof is by contradiction using only previously established structural results.

Step 1 — Assume the inaccessible sector is not physically real. Suppose $\mathcal{S}$_inaccessible is not physically real: it represents a formal remainder of description, an epistemic uncertainty over underlying states, or a coarse-graining artifact of $\mathcal{S}$_accessible. Under this assumption the total physical state is fully contained within $\mathcal{S}$_accessible and the decomposition is descriptive rather than ontological.

Step 2 — Consequence: reversibility in principle. If $\mathcal{S}$_inaccessible is purely formal, no physical degrees of freedom exist outside $\mathcal{S}$_accessible to which information can be transferred during commitment. Therefore all information about pre-commitment alternatives remains, in principle, recoverable from the accessible sector; any apparent loss of alternatives is due only to coarse-graining or description choice. Hence the commitment process is reversible in principle, even if not operationally convenient.

Step 3 — Contradiction with irreversibility. This contradicts Theorem 4.2, which establishes that fact-forming transitions are locally irreversible for bounded observers. Irreversibility requires that information distinguishing pre-commitment alternatives is not recoverable by any admissible bounded operation. But if no physically real sector exists outside $\mathcal{S}_{\text{accessible}}$, no such information can leave the accessible sector, and the process cannot be fundamentally irreversible.

Step 4 — Non-invertibility requires externalisation. By Theorem 6.2, property (1), the decomposition is non-invertible: no bounded operation on $\mathcal{S}_{\text{accessible}}$ recovers $\mathcal{S}_{\text{inaccessible}}$. Non-invertibility of this form requires that the discarded alternatives correspond to degrees of freedom not contained within the accessible sector — that they are externalised to physically real structure outside local control. A purely formal or epistemic remainder cannot produce true non-invertibility: it can always be inverted by refinement of description or acquisition of the hidden epistemic variable, restoring the pre-commitment state in principle. This contradicts property (1). This argument is structural and applies prior to any specification of the pre-commitment algebra; it is not a restatement of the hidden-variable exclusion in Theorem 10.0, which concerns the nature of pre-commitment states after the algebraic structure of $\mathcal{S}_{\text{pre}}$ has been established in Sections 9–10. The present step concerns only the structural decomposition of $\mathcal{S}$ itself, at the level of Theorem 6.2.

Step 5 — Conclusion. The assumption that $\mathcal{S}_{\text{inaccessible}}$ is not physically real leads to reversibility in principle (Step 2), contradicting required irreversibility (Step 3), and fails to support non-invertibility (Step 4). Therefore $\mathcal{S}_{\text{inaccessible}}$ must correspond to physically real degrees of freedom. □

Corollary 6.2b (Irreversibility Requires External Degrees of Freedom)

Irreversible commitment in any admissible theory $T \in \mathfrak{F}$ requires that information distinguishing pre-commitment alternatives is transferred to degrees of freedom that are physically real but inaccessible to bounded observers after commitment.

Remark 6.2c (Identification with the VERSF Void)

Theorem 6.2a establishes the necessity of a physically real, irreversibly inaccessible sector as a structural consequence of fact formation — not as a postulate. Within the VERSF framework, this sector is identified with the Void. The theorem does not assume the existence of the Void; it derives the necessity of a sector with precisely the defining properties attributed to the Void: physically real, inaccessible to bounded observers, and the repository of information transferred in irreversible commitment. The designation "Void" is therefore applied to a structure already forced by the admissibility conditions of $\mathfrak{F}$.

6.3 Codimension Theorem: The Interface Is a Separating Hypersurface

Theorem 6.3 (Codimension Theorem)

Let the effective substrate of $T \in \mathfrak{F}$ be an n-dimensional manifold M ($n \geq 2$). The commitment interface ∂C established by Theorems 6.1–6.2 is a codimension-1 submanifold of M — a separating hypersurface — and cannot be realised as a codimension-k structure for $k \geq 2$.

Proof

A codimension-k submanifold $\Sigma \subset M$ has dimension $n-k$. For Σ to function as a commitment interface it must separate M into two connected components — one accessible, one inaccessible — so $M \setminus \Sigma$ is disconnected.

By the Jordan–Brouwer Separation Theorem, a compact connected codimension-1 submanifold $\Sigma \subset S^n$ separates S^n into exactly two connected components. Alexander duality provides the corresponding homological condition for the general case: a compact connected submanifold $\Sigma \subset S^n$ of dimension $n-k$ has nontrivial $H_{n-1}(\Sigma; \mathbb{Z})$ if and only if $k = 1$. For dimension $d = n-k$, separation requires $n-k = n-1$, i.e., $k = 1$.

The substrate manifold M is not necessarily compact or simply connected. The application is as follows: by Lemma 5.6, M is a topological manifold, and the argument applies locally — within any bounded region of M of diameter D (which is the operationally relevant scale), the local topology is equivalent to an open region of $\mathbb{R}^n$, which is locally indistinguishable from an open region of S^n. Jordan–Brouwer applies to compact codimension-1 submanifolds within this local region: any such submanifold separates the local region into two components. This local separation is all that is required for the commitment interface to separate $\mathcal{S}_{\text{accessible}}$ from $\mathcal{S}_{\text{inaccessible}}$ within the bounded causal region of any finite observer. The global topology of M is not invoked.

For $k \geq 2$: the submanifold carries trivial H_{n-1} and does not separate M. Removing a codimension-2 or higher submanifold from a connected n-manifold ($n \geq 3$) leaves it path-connected — the two candidate sectors are not separated, and the interface fails to realise the decomposition of Theorem 6.2. The commitment interface is necessarily codimension-1. □

Corollary 6.4

The commitment interface ∂C is a hypersurface embedded intrinsically in the substrate of $T \in \mathfrak{F}$. It is not a bulk region, not a codimension-2 or higher structure, and not a globally defined algebraic constraint.

7. Elimination of Non-Interface Architectures

7.0 No Non-Local Commitment Realisation

Proposition 7.0 (No non-local commitment realisation)

Any commitment mechanism requiring information about, or operations on, degrees of freedom outside a bounded observer's causal reach either:

(a) violates condition 7 (local causal boundedness), making it physically inaccessible to any finite observer in $\mathfrak{F}$; or

(b) is observationally equivalent to an effective local interface, in which case it reduces to the interface architecture under the observational equivalence convention.

There is no third option.

Proof

Let M_nl be a non-local commitment mechanism requiring access to degrees of freedom in R_nl exceeding O's causal boundary. There are exactly two sub-cases.

Sub-case (a). No finite observer O can execute M_nl, since condition 7 prohibits joint control over all degrees of freedom in R_nl. The mechanism is not physically realisable within $\mathfrak{F}$.

Sub-case (b). If M_nl produces the same admissible fact outcomes as a local interface, the observational equivalence convention identifies M_nl with the local interface. Since the local interface posits no inaccessible non-local degrees of freedom, it is strictly minimal. M_nl is eliminated by condition 6.

No mechanism lies outside sub-cases (a) and (b).

No mechanism lies outside sub-cases (a) and (b). □

Lemma 7.0a (Quasi-Local Commitment Does Not Evade the Interface Result)

Any quasi-local commitment mechanism — one depending on degrees of freedom extending beyond the bounded neighbourhood of the commitment event but not to the full global substrate — is either reducible to an effective local interface description under observational equivalence, or physically distinct only by introducing additional fact-relevant degrees of freedom beyond the local commitment region, in which case it violates local causal boundedness and admissibility.

Proof

Let M_ql be a quasi-local commitment mechanism whose operation depends on degrees of freedom in a region R_ql extending beyond the local neighbourhood R_loc of the commitment event, but remaining finite.

Case 1: The extra quasi-local degrees of freedom do not alter admissible fact outcomes. Then replacing M_ql by a mechanism depending only on R_loc leaves all admissible fact distributions unchanged. By condition 6 and the observational equivalence convention, the extra quasi-local dependence carries no independent physical meaning. The mechanism reduces to an effective local interface description. $\square_1$

Case 2: The extra quasi-local degrees of freedom do alter admissible fact outcomes. Then those degrees of freedom are part of the physically relevant commitment mechanism. A finite observer determining or controlling the commitment event must access not only R_loc but the enlarged quasi-local region R_ql. If this enlarged dependence can vary without bound across commitment events, the mechanism ceases to be locally bounded in the sense required by condition 7. If instead R_ql has a fixed finite range, then it simply enlarges the local neighbourhood entering the effective interface description and does not define a genuinely new class of mechanism — it is still effectively local, now with a larger effective local region. $\square_2$

In both cases, quasi-locality gives no third option: it is either observationally redundant (Case 1) or reduces to effective locality with a possibly larger neighbourhood (Case 2). Quasi-local commitment does not evade Proposition 7.0 or Theorem 7.4. $\square$

Remark 7.0b

The force of this argument is that "local" in the present framework means "bounded and fact-complete at the commitment scale," not "mathematically pointlike." Finite-range dependence does not define an alternative to the interface architecture unless it introduces irreducible additional fact-relevant structure beyond what the local interface captures; when it does, it is already constrained by condition 7. The $\mathbb{H}$ elimination (Theorem 10.3) and the interface uniqueness result (Theorem 7.4) are both protected: neither relies on strict pointlike locality but only on the bounded-and-fact-complete notion formalised here.

Proposition 7.3 (Interface architecture is sufficient)

An intrinsic 2D commitment interface is sufficient to realise all conditions of $T \in \mathfrak{F}$ relevant to fact formation.

Proof

A 2D surface: (i) supports nontrivial $\beta_1 \geq 1$; (ii) intrinsically separates regions without external labelling; (iii) supports a finite-dimensional reversible state representation; and (iv) localises commitment events without global substrate updates. All fact-formation conditions of $\mathfrak{F}$ are satisfiable by the interface architecture. $\square$

Proposition 7.1 (Bulk irreversible collapse is non-minimal)

A theory realising fact formation only as a bulk global update is non-minimal within $\mathfrak{F}$.

Proof

By Proposition 7.3, the interface architecture realises all fact-formation conditions with local boundary participation only. A bulk-collapse theory requires all substrate degrees of freedom per commitment event. Under the observational equivalence convention, the additional global degrees of freedom produce no additional admissible facts. By condition 6 they are redundant. Bulk-collapse architectures are non-minimal. □

Proposition 7.2 (Extrinsically labelled boundaries fail admissibility)

Any architecture in which committed and reversible sectors are distinguished only by auxiliary external labelling violates condition 6.

Proof

Suppose ∂C is defined only by external labelling L. An alternative labelling L′ assigns the same substrate states to different sectors. Since L and L′ generate the same admissible facts, the observational equivalence convention identifies them. But if ∂C is label-dependent, commitment depends on description rather than structure, violating conditions 1 and 6. □

Theorem 7.4 (Intrinsic interface uniquely survives elimination)

Within the class of boundary architectures for fact formation in $\mathfrak{F}$, the unique minimal admissible realisation is an intrinsic commitment interface.

Proof

By Proposition 7.3, the intrinsic interface is sufficient. By Proposition 7.0, all non-local mechanisms are prohibited or reduce to the local interface. By Proposition 7.1, bulk collapse is non-minimal. By Proposition 7.2, extrinsic labelling violates admissibility. By Theorem 6.2, the decomposition must be locally defined and geometrically embedded. By Theorem 6.3, the embedded interface is codimension-1. These results exhaust the space of boundary architectures. The intrinsic interface is the unique survivor. □

7.5 Complete Structural Elimination

The objection this section answers: "You have not eliminated all alternatives — exotic algebras, non-local theories, and unusual topologies might survive."

The response is an explicit exhaustive partition of the structural decision space and an elimination argument for each case. Every combination of substrate topology, boundary type, interface dimension, and scalar field is covered. No case survives except the VERSF fold.

Topology of the substrate:

Case	Substrate	Result	Eliminated by
T1	Tree ($\beta_1 = 0$)	Cannot support locally constructible irreversibility	Lemma 5.3, Theorem 5.2
T2	Cyclic ($\beta_1 \geq 1$)	Passes	Proceeds to boundary classification

No other topological class exists for connected graphs: a connected graph is either acyclic (a tree) or cyclic. In the effective manifold description (Lemma 5.6), the graph-theoretic distinction "tree vs cyclic" corresponds precisely to the topological invariant $\beta_1 = 0$ *versus* $\beta_1 \geq 1$*, which is the appropriate formulation at the manifold level: the first Betti number measures the number of independent cycles and is well-defined for both graphs and manifolds.*

Type of commitment boundary:

Case	Boundary type	Result	Eliminated by
B1	Bulk collapse (no intrinsic boundary)	Non-minimal — no admissible facts inaccessible to local modification	Proposition 7.1
B2	Extrinsic labelling	Violates admissibility of description	Proposition 7.2
B3	Non-local constraint or projection	Either violates condition 7 or reduces to local interface	Proposition 7.0
B4	Intrinsic local interface	Passes	Proceeds to dimension classification

This partition is exhaustive: a boundary is either intrinsic (B4), extrinsically imposed (B2), non-local (B3), or replaced by a global bulk update (B1). No fourth category exists.

Dimension of the commitment interface:

Case	Interface dimension	Result	Eliminated by
D1	0D (point)	Cannot separate any region	Codimension argument in Theorem 6.3
D2	1D (curve or arc)	Cannot encode intrinsic commitment polarity	Proposition 8.1
D3	2D (surface)	Minimal — passes	Proceeds to binary data classification
D4	3D or higher	Non-minimal — admits continuous families of geometrically inequivalent separating structures producing no additional admissible facts; boundary description becomes underdetermined	Proposition 8.2 (condition ii), Proposition 8.3b, condition 6

This partition is exhaustive over dimension: interface dimension is a non-negative integer, and the argument covers all cases.

Scalar field for the reversible representation:

Case	Scalar field	Result	Eliminated by
A1	ℝ	O(n) contains no continuous U(1) phase subgroup → no continuous relative phase → interference condition fails	Theorem 10.3
A2	ℂ	Passes: U(4) contains U(1); all three admissibility conditions satisfied	Theorem 10.3
A3	ℍ (quaternions)	Quaternionic phases generate transformations producing no distinguishable fact outcomes for finite observers → unobservable → redundant under condition 6	Theorem 10.3
A4	𝕆 (octonions)	Non-associative → cannot support a consistent composition of reversible transformations	Theorem 10.3
A5	Any other normed division algebra	The normed division algebras are exactly {ℝ, ℂ, ℍ, 𝕆} (Hurwitz's theorem). Cases A1–A4 are exhaustive.	Hurwitz + Theorem 10.3

Any additional algebraic structure beyond ℂ that does not change admissible fact outcomes is eliminated by representational invariance — condition 6 together with the observational equivalence convention. The partition over normed division algebras is complete by Hurwitz's theorem.

Overall elimination:

The complete set of structural alternatives is the Cartesian product of the four classifications above. Every combination either: (a) fails at the topology stage (Case T1); (b) fails at the boundary type stage (Cases B1–B3); (c) fails at the dimension stage (Cases D1, D2, D4); or (d) fails at the algebra stage (Cases A1, A3, A4). The only surviving combination is: cyclic substrate (T2) × intrinsic local interface (B4) × 2D (D3) × ℂ (A2). This is the VERSF fold architecture.

The classification over substrate topology, boundary realisation, interface dimensionality, and scalar algebra is complete, and these four axes are mutually independent — a choice on any one axis does not constrain the available choices on any other, so their Cartesian product exhausts all admissible structural possibilities. No structural alternative lies outside this product.

While the four axes are logically distinct, their independence is structural rather than merely definitional: each axis corresponds to a necessary condition whose violation cannot be compensated by choices on any other axis. Topology constrains the existence of local trapping (Theorem 5.2): no choice of boundary type, dimension, or algebra can generate locally constructible irreversibility on an acyclic substrate. Boundary type constrains the localisation of commitment (Theorem 6.2): no topological cycle, dimensional choice, or scalar field can make an extrinsically labelled boundary physically intrinsic. Dimension constrains the determinacy of polarity encoding (Proposition 8.2): no topological, boundary, or algebraic structure can make a 1D interface encode intrinsic commitment polarity, nor make a $d > 2$ interface avoid underdetermination of the commitment locus. Algebra constrains interference (Theorem 10.0): no topological, boundary, or dimensional structure can make $\mathbb{R}$ or $\mathbb{H}$ support the continuous phase freedom that interference requires. These cross-compensations are structurally blocked, not merely logically excluded. The Cartesian product of the four axes therefore exhausts the admissible space, and the product structure is non-redundant. The classification is complete and exhaustive.

7.6 Completeness of the Structural Classification

A potential objection is that the classification of structural alternatives over the four axes is asserted rather than proven exhaustive. The completeness follows from the observation that any physical theory capable of producing stable facts must specify exactly four functional elements, and the four classification axes correspond precisely to these four functions.

The four necessary functions. Any admissible theory $T \in \mathfrak{F}$ must specify:

1. *Connectivity* — how information propagates through the substrate, determining whether local irreversibility is constructible (addressed by substrate topology).
2. *Commitment mechanism* — how the transition from open alternatives to committed facts is localised and physically implemented (addressed by boundary realisation type).
3. *Encoding capacity* — what geometric structure carries the binary polarity and orientation data that distinguish committed from reversible sectors (addressed by interface dimension).
4. *Reversible composition structure* — what algebraic structure governs the pre-commitment evolution that is reversible prior to commitment (addressed by scalar algebra).

These four functions are individually necessary — a theory lacking a specification of any one of them cannot produce stable, locally irreversible facts for finite observers. They are collectively sufficient to characterise the minimal fact-producing architecture, since together they determine how facts are topologically enabled, locally realised, geometrically encoded, and algebraically represented. No further functional element is required.

The four classification axes (topology, boundary type, dimension, algebra) are therefore not an arbitrary partition but the unique partition into the four necessary functional components of any admissible fact-producing theory. The classification contains no gaps because the four functions are exhaustive.

The completeness follows from the functional roles these axes represent.

Topology encodes the possibility of local irreversibility via cycle structure (Theorem 5.2). Any substrate supporting fact formation must instantiate a connectivity structure, which is either acyclic or cyclic — no third category exists for connected graphs. The binary partition is complete.

Boundary realisation encodes the localisation of commitment (Theorem 6.2). A separation between accessible and inaccessible sectors must be realised either intrinsically in the substrate, extrinsically by labelling, non-locally, or via global update. These four cases — corresponding to Cases B1–B4 in the classification — exhaust all possibilities for implementing a partition of state space. No further structural category for "how a separation is realised" exists.

Interface dimension encodes the minimal geometric capacity required to represent commitment polarity intrinsically (Proposition 8.2). Interface dimension is a non-negative integer, and all cases are explicitly analysed: dimension 0 (a point) cannot separate, dimension 1 fails intrinsic polarity, dimension 2 is minimal, and dimension 3 or higher fails fact-determinacy. The integer-valued classification is exhaustive.

Scalar algebra encodes the reversible composition structure (Section 10). By Hurwitz's theorem, the normed division algebras are exhausted by $\{\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}\}$. Any further algebraic structure either reduces to one of these under admissibility conditions or is non-associative and cannot support consistent reversible composition. The Hurwitz classification is complete.

These axes are functionally independent: no choice on one axis can compensate for failure on another (Section 7.5). Their Cartesian product therefore exhausts all admissible structural possibilities within $\mathfrak{F}$. The classification contains no gaps.

Definition 7.7 (Fact-Determinacy)

A physical structure S is *fact-determinate* if it is uniquely reconstructible, up to observational equivalence, from the admissible fact distributions it generates: there is no other structure S′, inequivalent to S under the observational equivalence convention, that produces the same admissible fact outcomes for all finite observers.

A theory or architectural component that is not fact-determinate contains structure that either (a) is not fixed by any admissible fact — making it physically underdetermined — or (b) is redundant relative to a leaner description — making it non-minimal under condition 6. Both cases are inadmissible within $\mathfrak{F}$.

Fact-determinacy is used explicitly in the dimensional argument of Section 8 (Proposition 8.2, condition (ii)), where it is proved that $d > 2$ interfaces fail fact-determinacy: the same committed/reversible fact pattern is compatible with continuously many inequivalent separating structures, making the interface underdetermined by the fact distributions it supports.

---

8. Minimal Dimension of the Commitment Interface

Proposition 8.1 (One-dimensional interfaces are insufficient)

A 1D interface cannot serve as the minimal intrinsic commitment boundary.

Proof

An intrinsic commitment boundary must encode commitment polarity without external reference.

Case 1: 1D open arc. An open arc bounds no region intrinsically. Any separation is a property of the ambient embedding, not of the arc. Commitment polarity is extrinsic, violating Proposition 7.2.

Case 2: 1D closed loop (S^1). S^1 separates an ambient 2D space into two components only by virtue of the ambient embedding. On the 1D manifold itself, S^1 is a connected manifold without boundary, with no intrinsic interior or exterior. Commitment polarity is extrinsic, violating Proposition 7.2. □

Definition 8.1a (Fact-Determinacy — Formal)

A physical structure X in an admissible theory $T \in \mathfrak{T}$ is *fact-determinate* if X is uniquely reconstructible, up to observational equivalence, from the full set of admissible fact distributions generated under all admissible protocols accessible to finite observers.

Equivalently: X is fact-determinate if whenever two candidate structures X and X′ generate identical admissible fact distributions under all admissible protocols, they are physically identical under the observational equivalence convention.

Fact-determinacy is not an epistemic condition but a structural one: it asserts that the mapping from physical structure to admissible fact distributions is injective up to observational equivalence. A structure that is not fact-determinate either lies in the kernel of this mapping (generating no distinguishing facts, hence redundant under condition 6) or in the non-injective region (multiple structures mapping to the same fact distributions, making the structure underdetermined by its own physical consequences).

Remark 8.1b

Fact-determinacy is not an additional postulate. It is the structural form of condition 6 applied to load-bearing physical architecture. A structure that is not fact-determinate either carries surplus unobservable content (violating minimality under condition 6) or fails to be uniquely specified by the physical facts it is supposed to ground (violating local definability, Theorem 6.2). The preliminary statement in Definition 7.7 is the informal version; Definition 8.1a is the formal version used in Section 8.

Lemma 8.1c (Commitment Interface Must Be Fact-Determinate)

The commitment interface established in Theorems 6.2–6.3 must be fact-determinate in the sense of Definition 8.1a.

Proof

By Theorem 6.2, the commitment interface is not an auxiliary mathematical convenience but the locally defined physical boundary separating accessible from inaccessible sectors. By Theorem 6.3, it is the codimension-1 hypersurface realising that separation in the effective manifold description.

Suppose the interface were not fact-determinate. Then there exist two geometrically inequivalent candidate interfaces Σ and Σ' that generate identical admissible fact distributions under all admissible protocols yet are not observationally equivalent.

Two possibilities arise:

(i) *The distinction between Σ and Σ' is physically meaningful.* Then the same complete admissible fact data fail to determine which physical commitment boundary is present. The physical locus of commitment is underdetermined by the facts it is supposed to generate. But the interface is locally defined by Theorem 6.2: the commitment status of any state is determined by local data in a bounded neighbourhood. A locally defined physical boundary must be reconstructible from the local fact structure it produces. Otherwise locality of definition fails — the same local fact data are compatible with multiple inequivalent physical boundaries, destroying the local definability property proved in Theorem 6.2 property (2).

(ii) *The distinction between Σ and Σ' is not physically meaningful.* Then the difference between them carries no admissible physical content and is excluded by condition 6. In that case Σ and Σ' are physically identical under observational equivalence after all — contradicting the assumption that they are inequivalent.

In either case, non-fact-determinacy is inadmissible: it either violates local definability (case i) or collapses to observational equivalence (case ii). The commitment interface must be fact-determinate. □

Proposition 8.2 (Dimensional Minimality from Fact Determinacy)

The minimal dimension of the commitment interface is the lowest dimension d such that the interface is fact-determinate in the sense of Definition 7.7 — uniquely reconstructible from admissible fact distributions — while supporting intrinsic separation of committed from reversible sectors. Formally, this requires:

(i) intrinsic separation of committed from reversible sectors is possible (the interface admits a closed separating curve), and

(ii) the separation is **uniquely determined by local structure** — that is, there is no continuous family of geometrically inequivalent separations that produce the same admissible fact outcomes.

Condition (i) fails for $d = 1$ and holds for $d \geq 2$. Condition (ii) holds for $d = 2$ and fails for $d > 2$. Therefore $d = 2$ is the unique dimension satisfying both conditions.

Proof

Condition (i) for d = 1. On a 1D manifold, a separating curve is the full manifold itself (as shown in Proposition 8.1). No proper sub-region separation is intrinsically possible. Condition (i) fails.

Condition (i) for d = 2. By the Jordan Curve Theorem, every simple closed curve on S^2 separates S^2 into exactly two components. Condition (i) holds.

Condition (ii) for d = 2. On S^2 (minimal genus, established by Proposition 8.5 below), the separation is uniquely determined by the Schoenflies theorem: any simple closed curve on S^2 bounds a topological disc, and any two such curves are ambient isotopic via an orientation-preserving homeomorphism of S^2. The moduli space of separating curves on S^2 therefore consists of a single isotopy class — there is no continuous family of geometrically inequivalent separating curves on S^2 producing distinct embedded structures. The separation is determined, up to the polarity assignment $\sigma \in \{0,1\}$, entirely by the two components it creates. No additional moduli freedom exists on S^2 that could distinguish one separation from another while producing the same committed and reversible domains. Condition (ii) holds.

Condition (ii) fails for d > 2. For an interface manifold Σ of dimension $d > 2$, the space of embedded separating codimension-1 submanifolds within Σ forms a continuous moduli space. Concretely: a codimension-1 submanifold $\Sigma' \subset \Sigma$ separating Σ into two components can be continuously deformed — not merely by isotopy but by geometry-changing deformations — into a family of geometrically inequivalent separating submanifolds $\{\Sigma'_t\}$ that all produce exactly the same partition of Σ into committed and reversible domains. The deformation changes the intrinsic geometry of Σ'_t (its curvature, embedding class within Σ, and normal bundle structure) without changing which points of Σ belong to which sector.

Each element of this continuous family is a geometrically distinct separating structure. By condition 6 (admissibility of description), geometrically distinct structures that produce no additional admissible fact outcomes are observationally equivalent and therefore redundant. But since the family is continuous — parametrised by a real modulus — it constitutes infinitely many inequivalent descriptions of the same commitment separation, none of which is singled out by any admissible fact. There is no fact-grounded criterion for selecting a unique element of the moduli family. The interface description is underdetermined: the separation boundary is ambiguous up to a continuous family of equivalent deformations.

Why condition (ii) fails for $d > 2$.

In dimensions $d > 2$, separating codimension-1 submanifolds generically admit a positive-dimensional family of inequivalent embeddings compatible with the same local sector assignment. The important issue is not merely that this family exists, but that the same committed/reversible fact pattern fails to determine a unique separator.

The ambiguity is not resolved by ambient isotopy. While isotopic embeddings represent the same topological separator — the same global topological partition of Σ — the relevant distinction here is not topological equivalence but geometric realisation within the effective manifold. The commitment interface is locally defined (Theorem 6.2, property 2), so its realisation must be fixed by local geometric data, not merely by global topological class. In dimensions $d > 2$, there exist continuously many geometrically distinct local embeddings compatible with the same topological partition, differing in curvature, normal bundle structure, and local embedding geometry. These differences are not removable by isotopy without altering local geometry.

Concrete example for $d = 3$. Consider a 3-dimensional interface manifold Σ^3 and a separating 2-sphere $\Sigma^2 \subset \Sigma^3$. There exists a one-parameter family of embeddings $\{\Sigma^2_t\}_{t \in [0,1]}$ related by conformal rescaling of the induced metric on Σ^2_t — for instance, embeddings that uniformly inflate or deflate the sphere's intrinsic curvature while keeping the same partition of Σ^3 into two components. Each Σ^2_t separates the same committed and reversible domains (the same two connected components of Σ^3), so the committed/reversible fact pattern is identical across the family. Yet the embeddings are geometrically distinct: Σ^2_t and Σ^2_s ($t \neq s$) have different Gaussian curvature at every point and different normal bundle geometry, so no isotopy carries one to the other without altering local geometric data. No admissible fact (which is determined by which alternatives are committed, not by the curvature of the interface boundary) distinguishes Σ^2_t from Σ^2_s. The separator is therefore geometrically underdetermined by the fact partition, violating fact-determinacy (Definition 8.1a). The same argument applies for any $d > 2$: there always exists a positive-dimensional family of geometrically inequivalent embeddings compatible with the same topological partition.

Therefore the ambiguity persists at the level of local geometric definition and is not resolved by topological equivalence or isotopy class.

This is fatal. The interface is the physical locus of commitment (Theorems 6.2–6.3), not an auxiliary description. By Theorem 6.2 property (2), the commitment boundary is locally defined: reconstructible from local data within a bounded neighbourhood. A continuous ambiguity in its geometric realisation — where the same local fact data are compatible with continuously many inequivalent separators — directly contradicts local definability. Either the specific geometry is physically meaningless (surplus non-derived structure, violating condition 6) or it is physically meaningful (same facts underdetermine the boundary, violating Theorem 6.2). Both options are inadmissible.

The problem in $d > 2$ is therefore not descriptive redundancy but failure of fact-determinacy: the local commitment data no longer fix a unique interface up to observational equivalence. Condition (ii) fails for $d > 2$.

Clarification — this is not gauge redundancy. The ambiguity identified for $d > 2$ is not analogous to gauge redundancy. In gauge theories, different configurations correspond to the same physical state, and the gauge equivalence class is itself well-defined and unique — one may freely pass between representatives without physical consequence. Here, by contrast, the interface is not a redundant description layered over physical structure but is the physical locus of commitment itself (Theorems 6.2–6.3). If multiple geometrically inequivalent separators correspond to the same fact pattern, then either: (i) the separator geometry is unphysical, in which case the interface carries surplus non-derived structure and violates minimal admissibility (condition 6); or (ii) the separator geometry is physical, in which case the commitment boundary is not uniquely determined by local fact data, violating locality of definition (Theorem 6.2, property 2). In a gauge theory, the gauge orbit is the physical object and any representative suffices. Here there is no analogous gauge orbit — the interface must be a specific locally reconstructible physical locus, not an equivalence class of geometrically inequivalent structures. The ambiguity cannot be treated as gauge freedom and is genuinely inadmissible. □

Theorem 8.3 (Minimal Separating Manifold)

The minimal dimension of an intrinsic commitment interface — encoding commitment polarity as an internal geometric property, uniquely determined by local structure — is 2.

Proof

Why 1D fails. On a compact connected 1D manifold (S^1), $H_0(S^1) = \mathbb{Z}$ implies S^1 is connected with no proper non-empty clopen subsets. No closed curve on S^1 bounds a proper sub-region intrinsically. Condition (i) of Proposition 8.2 fails.

Why 2D is minimal. By the Jordan Curve Theorem, every simple closed non-null-homotopic curve on a 2D surface bounds two intrinsic components (Proposition 8.2, condition (i)). On S^2 specifically, by the Schoenflies theorem, the moduli space of separating curves consists of a single isotopy class: the separation is uniquely determined up to polarity (Proposition 8.2, condition (ii)). Both conditions of Proposition 8.2 are satisfied, and 2D is the minimum for which this holds.

Codimension compatibility. By Theorem 6.3, the commitment interface must be codimension-1 in the substrate. The interface as a carrier of intrinsic commitment polarity is therefore a 2D surface: codimension-1 in a 3D substrate, or — in the minimal substrate case — a 1D curve within a 2D interface manifold, where the 2D manifold itself carries the binary polarity data.

Why $d > 2$ fails (the correct argument). In dimensions $d > 2$, the space of separating hypersurfaces within the interface manifold contains continuous families of geometrically inequivalent embeddings that produce no additional distinguishable fact outcomes. Higher-dimensional interfaces fail fact-determinacy: the same local commitment data are compatible with continuously many inequivalent separators, so the interface ceases to be uniquely recoverable as the physical locus of commitment. The commitment boundary is underdetermined by the fact structure. By condition 6 and Proposition 8.2 (condition (ii)), this underdetermination is inadmissible. Higher-dimensional interfaces are therefore eliminated — not merely non-minimal, but fact-underdetermining. □

Corollary 8.4

The minimal intrinsic commitment interface in $\mathfrak{F}$ is a 2-dimensional surface: the VERSF fold.

Corollary 8.3b (No Higher-Dimensional Minimal Interfaces)

Any interface of dimension $d > 2$ fails fact-determinacy (Definition 8.1a) and is inadmissible within $\mathfrak{F}$. This follows directly from Proposition 8.2 (condition (ii) failure for $d > 2$): the moduli family of geometrically inequivalent separators maps to the same fact distribution, violating injectivity, so $d > 2$ interfaces are either non-minimal (stripping moduli yields a lower-dimensional interface) or inadmissible (the moduli cannot be removed without destroying the interface). □

Proposition 8.5 (Minimal Genus)

The minimal commitment interface is genus 0 (topologically S^2, the 2-sphere).

Proof

A closed orientable 2D surface Σ of genus g has first homology $H_1(\Sigma; \mathbb{Z}) = \mathbb{Z}$^{2g}. Each generator of $H_1(\Sigma; \mathbb{Z})$ corresponds to an independent non-contractible cycle class on Σ — a distinct homotopy class of closed loops on Σ that cannot be continuously deformed to a point.

Each independent non-contractible cycle class constitutes a topologically distinguishable structural feature of the interface: two loops in different homotopy classes cannot be deformed into each other, so they represent distinct and independently identifiable properties of Σ. By condition 6 (admissibility of description) and the minimality requirement, any distinguishable structural feature of the interface that does not generate additional admissible fact outcomes is redundant and must be excluded.

The commitment interface requires exactly one cycle class to encode commitment polarity: the class of the closed curve bounding the committed sub-domain of Σ. This is the structure established in Section 9 (the binary data (σ, ω) are determined by this curve and its orientation). No additional cycle class is required for this function.

For $g \geq 1$, Σ carries 2g additional independent non-contractible cycle classes beyond those required for commitment polarity. Each additional class is a distinguishable feature of the interface generating no additional admissible fact outcomes. By condition 6, these are redundant. Any interface of genus $g \geq 1$ is therefore non-minimal.

The unique minimal genus is $g = 0$. The minimal commitment interface is topologically S^2. Any additional homotopy class constitutes an independent distinguishable structure not required for fact formation, violating minimality under condition 6. □

Corollary 8.6

The minimal commitment interface is S^2, carrying exactly the two binary geometric structures (σ, ω) established in Section 9 and no additional topological data.

---

9. Geometric Binary Data of the Interface

Definition 9.1 (Commitment polarity)

Intrinsic region separation on the 2D interface induces a binary polarity $\sigma \in \{0, 1\}$, corresponding to the assignment of committed versus reversible status to the two domains separated by the commitment curve.

Definition 9.2 (Reversible orientation)

A 2D orientable interface admits a local orientation structure. Orientation reversal is a binary symmetry $\omega \in \{-1, +1\}$, reflecting the two inequivalent choices of local traversal orientation.

Lemma 9.3 (Exactly two independent binary geometric structures)

The minimal 2D commitment interface carries exactly the two binary structures (σ, ω), which are independent. No additional independent binary geometric degree of freedom exists on the minimal interface.

Proof

Independence. Reversing ω does not exchange committed and reversible domains: orientation is a property of boundary traversal, not domain assignment. Reversing σ does not alter orientation. The two structures act at distinct geometric levels and are independent.

Completeness. All candidate additional binary structures are addressed:

- *Global topology (orientable vs. non-orientable):* The binary ω requires orientability. Non-orientable surfaces cannot support a consistent global ω and exit the minimal orientable class.
- *Embedding class (knotting, linking):* Extrinsic properties of the ambient embedding, excluded by Proposition 7.2.
- *Spin structure:* On S^2 (topologically minimal closed orientable surface), $H^1(S^2; \mathbb{Z}_2) = 0$, giving a unique spin structure. No independent binary datum.
- *Framing:* Integer-valued ($\mathbb{Z}$), not binary.
- *Chirality:* On an orientable 2D surface, intrinsic handedness is captured by orientation, equivalent to ω. Not independent.
- *Boundary orientation consistency:* Encoded by the combination of σ and ω. Not independent.

The complete set of independent binary geometric data on the minimal 2D interface is exactly {σ, ω}. □

Corollary 9.4

The minimal geometric interface state set is

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

containing exactly four states.

Proposition 9.4b (No Privileged Binary Structure)

Neither σ nor ω may be privileged over the other in the reversible representation of the interface. The two binary degrees of freedom must enter symmetrically.

Proof

Both σ and ω are intrinsic geometric degrees of freedom of the same minimal 2D interface (Lemma 9.3). They are both defined by the interface's own structure: σ by the intrinsic region separation (Definition 9.1) and ω by the intrinsic orientation structure (Definition 9.2). Neither is defined by reference to an external embedding or an external labelling scheme — both are grounded in the same interface geometry.

Suppose the reversible representation privileges σ over ω — that is, the two binary degrees of freedom enter the representation asymmetrically, with σ playing a structurally distinct role. Such an asymmetry would constitute additional structure in the representation beyond what is determined by the interface geometry: the representation would distinguish between σ and ω in a way that the interface itself does not. By condition 6 (admissibility of description), this additional structure carries distinct physical meaning only if it generates distinct admissible fact outcomes.

To see that no such distinction exists, note that σ and ω are geometrically different in character: σ is a region assignment (which domain is committed, which is reversible), and ω is a directional property (the traversal orientation of the bounding curve). They are not the same kind of geometric object. However, from the perspective of fact formation, both binary choices bear on the commitment event through exactly the same mechanism: by determining the partition of the interface into committed and reversible domains. The committed fact is determined by which domain is assigned as accessible (σ) and by how the boundary of that domain is oriented (ω). Neither binary choice produces a fact outcome of a different *type* — both contribute to the same single act of specifying a committed region on the interface. There is no class of admissible facts whose production requires σ to be defined but ω to be undefined, or vice versa. Both choices are necessary to specify the commitment event, and both choices are of the same structural kind: a binary selection on the interface that participates in grounding the committed sector. The partition into committed and reversible domains is the only fact-relevant structure, and both σ and ω enter it symmetrically in the sense that neither can be omitted or privileged without changing the set of definable commitment events.

Therefore, no fact outcome distinguishes a representation that privileges σ over ω from one that privileges ω over σ. There is no class of admissible facts that can be produced by σ-type commitment but not ω-type commitment or vice versa.

Therefore, any privileging of σ over ω in the reversible representation generates structure that produces no additional admissible fact outcomes. By condition 6 and the observational equivalence convention, such privileging is redundant and excluded. The two binary degrees of freedom must enter the reversible representation symmetrically: the representation must be invariant under the exchange $\sigma \leftrightarrow \omega$.

This is the isotropy condition (Definition 10.2.2). It is now derived as a structural consequence of the interface geometry and condition 6, not assumed as an admissibility axiom. □

9.5 Necessity of Linear Composition

9.5.0 Pre-Commitment vs Committed Distinguishability

Before establishing the structure of the pre-commitment state space, a potential tension must be resolved. Theorem 4.1 establishes that only finitely many distinctions are stably recordable in bounded regions — a finite δ_{min} exists below which no distinction constitutes a stable fact. The pre-commitment state space $\mathcal{S}_{\text{pre}}$, derived below, is a continuous vector space. These two claims may appear to conflict.

They do not, because the finite distinguishability condition and the continuous pre-commitment structure apply to entirely different sectors of the theory:

Finite distinguishability (Theorem 4.1) applies to *committed facts* — outcomes of commitment events that are stably recorded, recovered, and compared by finite observers. The minimum scale δ_{min} bounds the resolution of the fact record: only finitely many committed outcomes can be distinguished in bounded regions.

The pre-commitment sector represents *unresolved alternatives that have not yet been committed to as facts*. Its structure encodes the possibility space of outcomes prior to commitment, not the realised outcomes themselves. Pre-commitment states are not directly observable as distinct stable facts: they are the open possibilities from which the commitment event selects. Their continuous structure does not contradict finite distinguishability, because they are precisely the states that have not yet resolved into facts.

Therefore, finite distinguishability constrains the *output* of commitment — the discrete record produced — not the *structure* of the pre-commitment state space. The latter may admit a continuous representation without violating any condition of $\mathfrak{F}$. This is the standard relationship between a continuous quantum state space and a discrete measurement record, here derived from operational principles rather than assumed.

The gap this section closes. Section 10 derives which scalar field F underlies the reversible pre-commitment representation. But before asking which field, the paper must establish *why the representation must be linear at all*. Without this, a referee could object that the pre-commitment structure could be nonlinear — a nonlinear manifold of states rather than a vector space. This section proves that linearity is forced.

Theorem 9.5 (Necessity of Affine, Hence Linear, Composition)

Any admissible pre-commitment composition law for $T \in \mathfrak{F}$ that combines unresolved alternatives, preserves reversible recombination, is observer-order independent, and yields identical committed fact classes for operationally equivalent preparation procedures, must be affine in its arguments. After fixing the neutral unresolved state as the origin, the composition law becomes linear, and the pre-commitment sector is a vector space over a field F.

The theorem does not assume Hilbert-space linearity at the outset; it derives linearity as the unique representation compatible with preparation equivalence and decomposition-independence for unresolved alternatives.

Proof

Let $\mathcal{S}$_pre denote the pre-commitment state space. Let $C(\psi, \varphi; \lambda)$ denote the admissible composition of unresolved alternatives $\psi, \varphi \in \mathcal{S}$_pre with continuous weighting parameter λ. The proof imposes only operational requirements already established in the paper.

(R1) Closure. $C(\psi, \varphi; \lambda) \in \mathcal{S}$_pre for all admissible ψ, φ, λ.

(R2) Recombination invariance. If two preparation procedures produce the same unresolved alternative, replacing one by the other inside any larger composition must not change the resulting fact class. Operationally equivalent preparations are interchangeable in all larger compositions.

(R3) Order independence. Multi-step composition of unresolved alternatives must not depend on the order in which equivalent intermediate compositions are formed.

(R4) Continuity. Small changes in preparation weights produce small changes in the pre-commitment state.

Deriving affineness. The same physical unresolved alternative can be prepared through different decomposition routes. A state prepared as a combination of ψ_1 and ψ_2 then combined with ψ_3 must be operationally equivalent to a state prepared by first combining ψ_2 and ψ_3 then combining with ψ_1, whenever the two procedures correspond to the same unresolved alternative. If the composition rule were nonlinear in a route-dependent way, the final pre-commitment state would depend on the decomposition history rather than only on the physical alternative prepared. That violates R2 (operationally equivalent preparations cease to be interchangeable inside larger compositions) and R3 (grouping dependence survives into the final state description).

Therefore admissible composition must preserve interpolation between preparation procedures: the state assigned to a weighted mixture of preparations must be the corresponding weighted mixture of the assigned states. For any admissible preparations P_1, P_2 with state assignments ψ_1, ψ_2 and any $\lambda \in [0,1]$, the preparation-to-state map Γ must satisfy

$$\Gamma(\lambda P_1 + (1-\lambda)P_2) = \lambda\, \Gamma(P_1) + (1-\lambda)\, \Gamma(P_2).$$

If this failed, operationally equivalent preparation mixtures would map to inequivalent states, violating admissibility of description (condition 6). The admissible state representation is therefore affine.

Clarification — affineness does not assume classical mixtures. The affineness condition does not assume that all pre-commitment combinations are classical probability mixtures. Rather, it follows from the requirement that any two preparation procedures that are operationally indistinguishable at the level of committed facts must remain interchangeable under further composition. This forces the representation to preserve convex structure at the level of preparation equivalence classes, regardless of whether the underlying ontology is classical or non-classical. The argument is about equivalence-class structure, not about probabilistic interpretation.

From affine to linear. The pre-commitment sector contains a distinguished neutral unresolved state $0 \in \mathcal{S}$_pre corresponding to the null preparation contrast — the state in which no alternative is preferentially weighted. The existence of such a neutral state follows from the requirement that preparation procedures can be compared only up to relative weighting of alternatives; the zero vector corresponds to the absence of differential preparation bias, and must exist as the unique fixed point of the trivial preparation (weight zero everywhere). Once this origin is fixed, affine structure reduces to linear structure: differences of states are well-defined, scalar weighting extends by continuity from $[0,1]$ to $\mathbb{R}$ via R4, and composition becomes linear combination. Extension to negative coefficients follows from reversibility: if a state difference can be created by admissible composition, reversibility (R2) implies it can be undone, generating the additive inverse and completing the vector space structure over $\mathbb{R}$.

Finite dimensionality. By Corollary 9.4, the interface carries exactly four independent basis states — the four elements of $\{0,1\} \times \{-1,+1\}$ corresponding to the (σ, ω) pairs. The pre-commitment sector is spanned by superpositions of these four states: $\mathcal{S}$_pre has dimension at most four over the underlying field, with the four interface states as its natural basis. The identification of the dimension as exactly four is established separately in Proposition 10.1.

Eliminating nonlinear alternatives. Suppose composition is nonlinear: $C(\psi, \varphi; \lambda)$ is not affine in ψ, φ. Then route-dependence arises immediately (different decomposition paths for the same alternative yield different states), violating R2. Furthermore, path-dependence in state construction means different observers combining the same alternatives in different orders arrive at different descriptions of the same physical situation — violating representational invariance (condition 6 and Definition 10.2.3). Nonlinear composition is excluded.

Scalar field. The continuous parametric family (R4) defines a map $\lambda \mapsto \lambda \cdot \psi$ satisfying the distributivity axioms $(\lambda + \mu) \cdot \psi = (\lambda \cdot \psi) \oplus (\mu \cdot \psi)$ and $\lambda \cdot (\psi \oplus \varphi) = (\lambda \cdot \psi) \oplus (\lambda \cdot \varphi)$, which follow from compatibility with the group operation. Associativity $(\lambda\mu) \cdot \psi = \lambda \cdot (\mu \cdot \psi)$ and the identity axiom $1 \cdot \psi = \psi$ complete the vector space axioms over $\mathbb{R}$. The determination of the admissible field extension $F \supseteq \mathbb{R}$ is deferred to Section 10.

$\mathcal{S}$_pre is a vector space over a field F. The pre-commitment combination structure is linear. □

Lemma 9.5a (Nonlinear composition violates preparation equivalence)

Any non-affine composition law on the pre-commitment state space induces preparation-route dependence that produces distinct committed fact distributions, violating admissibility of description (condition 6).

Proof

Let $\mathcal{S}$_pre be the pre-commitment state space and suppose the composition law $C(\psi, \varphi; \lambda)$ is nonlinear in its arguments.

Consider three unresolved alternatives ψ_1, ψ_2, ψ_3 and two preparation procedures:

- P_A: first form $C(\psi_1, \psi_2; \lambda)$, then combine with ψ_3.
- P_B: first form $C(\psi_2, \psi_3; \mu)$, then combine with ψ_1.

These procedures can be arranged to produce the same preparation weights over the underlying alternatives and are therefore operationally equivalent at the level of preparation statistics. If the composition law is nonlinear, then in general:

$C(C(\psi_1, \psi_2; \lambda), \psi_3; \nu) \neq C(\psi_1, C(\psi_2, \psi_3; \mu); \nu')$,

so the resulting pre-commitment states differ as elements of $\mathcal{S}$_pre.

Now consider the subsequent commitment process. Since commitment produces fact distributions determined by the pre-commitment state (Theorem 10.0, proved subsequently in Section 10), distinct pre-commitment states generically yield distinct distributions over fact classes. Therefore P_A and P_B, which are operationally equivalent preparation procedures, produce different committed fact distributions.

This violates condition 6 (admissibility of description): operationally indistinguishable preparations must not produce distinguishable fact outcomes. Hence nonlinear composition is inadmissible and the composition law must be affine. □

Lemma 9.5b (Context-Independence of Admissible Composition)

Any admissible composition law on the pre-commitment state space must be context-independent: if two preparation procedures are operationally equivalent, then replacing one by the other inside any larger admissible composition cannot change the resulting committed fact distribution.

Proof

Let P and P′ be two preparation procedures producing the same unresolved alternative in the sense of the operational equivalence convention: no admissible finite observer can distinguish them by any direct protocol prior to commitment. Suppose, for contradiction, that composition is context-dependent. Then there exists an admissible extension context K[·] such that

$K[P] \neq K[P']$

at the level of resulting pre-commitment states, and consequently the subsequent commitment process yields different fact distributions for K[P] and K[P′].

But this means that P and P′, although operationally equivalent in isolation, become distinguishable once embedded in a larger admissible preparation procedure. Their equivalence was therefore not physical equivalence but merely accidental equivalence relative to a restricted protocol class. This contradicts condition 6, which requires that physically equivalent descriptions remain indistinguishable under all admissible fact-producing protocols.

Equivalently: if operational equivalence were not stable under admissible extension, then the fact class associated with a preparation would depend on external compositional context rather than on the preparation itself. The identity of a fact would no longer be intrinsic to the state produced, but contingent on how that state was later embedded. That destroys the reproducibility and comparability conditions of Definition 2.0.1, since repeated preparation of the "same" alternative could yield different fact classes depending on extension context.

Therefore admissible composition must preserve operational equivalence under all admissible extensions. Composition is context-independent. □

Corollary 9.6

The pre-commitment state space 𝒮_pre is a finite-dimensional vector space over a field F, and the reversible pre-commitment evolution group acts on it by linear automorphisms. The question of which field F is admissible is addressed in Section 10.

---

10. Minimal Reversible Representation of the Interface

10.0 Theorem: Necessity of Interference

The gap this theorem closes. Definition 10.2 lists interference as an admissibility condition on the scalar field F. Without deriving interference from prior structure, this condition appears as an assumption. Theorem 10.0 proves that interference is not assumed but forced: any non-classical pre-commitment sector with genuinely open alternatives must admit it.

Theorem 10.0 (Necessity of Interference)

Any admissible pre-commitment sector for T ∈ 𝔉 supporting genuinely unresolved alternatives must admit interference. Any composition rule for pre-factual alternatives that excludes interference is observationally equivalent to a classical stochastic mixture, contradicting the genuine openness of the pre-factual alternatives.

Proof

Step 1 — Pre-factual alternatives are not classical probability mixtures. A classical probability mixture represents a situation in which the outcome is already determined but unknown: the system is in state ψ_A with probability p or in state ψ_B with probability 1−p, but the actual state is fixed. Measurement merely reveals the pre-existing outcome.

Pre-factual alternatives in T ∈ 𝔉 are genuinely unresolved: the commitment event (Proposition 6.1, Theorem 6.2) produces the outcome. The outcome is not pre-existing — it is created by commitment. A classical mixture would mean commitment merely reveals a pre-existing fact — contradicting Theorem 4.2, which proves commitment is locally irreversible. Local irreversibility implies the pre-commitment state is not an epistemic mixture over fixed outcomes, but a genuinely open state from which an irreversible selection is made.

One might object that a deterministic hidden-variable theory could produce locally irreversible commitment while maintaining pre-determined outcomes: the hidden variable fixes the result, and commitment merely discloses it. This possibility is excluded by the commitment-boundary structure, independently of Definition 2.0.4.

A hidden-variable completion with pre-determined outcomes is incompatible with the non-invertible decomposition proved in Theorem 6.2. If the outcome is already fixed prior to commitment, then the discarded alternatives are not genuinely open alternatives that become inaccessible through commitment — they are merely epistemic possibilities relative to an observer's incomplete knowledge of the hidden variable. But Theorem 6.2 requires a physical decomposition into accessible and inaccessible sectors generated by the commitment process itself. A merely epistemic decomposition is invertible in principle by acquisition of the hidden variable, and therefore does not realise the non-invertible commitment structure required by $\mathfrak{F}$. The hidden-variable option is therefore excluded on structural, not interpretational, grounds: it fails to instantiate the non-invertible state-space decomposition that Theorem 6.2 proves is necessary.

Step 2 — Linear non-classical combinations must admit cancellation. By Theorem 9.5, the pre-commitment combination structure is linear. In a linear vector space over field F, the combination $\alpha\psi\_A + \beta\psi\_B$ has squared norm:

$|\alpha\psi\_A + \beta\psi\_B|^2 = |\alpha|^2|\psi\_A|^2 + |\beta|^2|\psi\_B|^2 + 2\mathrm{Re}(\alpha\beta^*\langle\psi\_A, \psi\_B\rangle)$.

The cross-term $2\mathrm{Re}(\alpha\beta^*\langle\psi\_A, \psi\_B\rangle)$ is the interference term.

Step 3 — Excluding interference implies classical mixture. Suppose the composition rule excludes interference: the cross-term vanishes identically for all pairs (ψ_A, ψ_B). Then $|\alpha\psi\_A + \beta\psi\_B|^2 = |\alpha|^2|\psi\_A|^2 + |\beta|^2|\psi\_B|^2$ — exactly the law of total probability for a mixture with weights $|\alpha|^2$ and $|\beta|^2$. The rule is observationally equivalent to a classical mixture. But Step 1 proves the pre-factual sector is not a classical mixture. Contradiction. Therefore interference is necessary.

Any composition rule without interference reduces to classical stochastic mixture, violating nontriviality of the pre-factual sector. □

Proposition 10.1 (Four-dimensional state space is necessary)

Any admissible linear reversible representation of two independent binary structures requires a state space of dimension exactly 4 over F, with tensor product structure $F^2 \otimes F^2$.

Proof

Each independent binary structure (σ and ω, from Lemma 9.3) requires a two-dimensional reversible module over F — the minimum for a reversible representation of a binary choice supporting superposition across both values.

Why tensor product, not direct sum. The joint state space must carry a representation of the group $\{0,1\} \times \{-1,+1\}$ acting simultaneously on every state. The two binary structures are independent (Lemma 9.3), meaning both degrees of freedom must act *nontrivially on every state*: for any state $|\psi\rangle$, the σ-action and the ω-action each produce a genuine transformation. A direct sum $F^2 \oplus F^2$ represents the two binary structures on disjoint subspaces: a state in the first summand has zero projection onto the second, so σ acts nontrivially but ω acts trivially on it, and vice versa. This makes simultaneous specification of σ and ω on any given state impossible —

the two degrees of freedom cannot both be active on the same state. This directly contradicts the independence requirement of Lemma 9.3, which demands that both binary degrees of freedom can be simultaneously specified and simultaneously varied for any state.

The unique representation in which both σ and ω act nontrivially on every state is the tensor product $F^2 \otimes F^2$, which carries the full product representation: the σ-action is $\mathrm{id} \otimes \hat{\sigma}$ and the ω-action is $\hat{\omega} \otimes \mathrm{id}$, and both act nontrivially on all of $F^2 \otimes F^2 \cong F^4$. This gives dimension $2 \times 2 = 4$. No lower-dimensional space carries the full product representation without conflating the two actions. Any higher-dimensional representation containing F^4 as an invariant subspace carrying the σ and ω actions introduces additional degrees of freedom not corresponding to any admissible fact distinction — by condition 6, such extensions are redundant and excluded. Therefore the minimal admissible representation has dimension exactly 4. □

Definition 10.2 (Admissibility conditions on F)

The scalar field F must satisfy:

1. **Interference:** superposition states must yield interference patterns with cross-terms varying continuously with a phase parameter θ. The squared norm of $\alpha|\psi_A\rangle + \beta|\psi_B\rangle$ must contain a term varying continuously with $\arg(\alpha\beta^*)$.
2. **Isotropy:** the representation must be invariant under relabelling of $(\sigma \leftrightarrow \omega)$, with no privileging of one binary sector over the other.
3. **Representational invariance:** physically equivalent reversible evolutions producing the same admissible facts correspond to the same transformation class on F^4. Any additional algebraic structure that does not change admissible fact outcomes is eliminated by representational invariance, via condition 6 and the observational equivalence convention.

Theorem 10.3 ($\mathbb{C}$ is the unique minimally admissible scalar field)

The unique minimally admissible scalar field for the reversible pre-commitment representation of $T \in \mathfrak{F}$ is $F = \mathbb{C}$.

Proof

Elimination of $F = \mathbb{R}$.

Over $\mathbb{R}$, the reversible transformation group is $O(n)$. $O(n)$ contains no continuous $U(1)$ phase subgroup: its only phase-like elements are reflections with eigenvalues ± 1. The absence of a continuous $U(1)$ subgroup in $O(n)$ prevents the existence of a continuous phase degree of freedom, which is required for interference. Consequently, there is no continuous relative phase between superposed components. Interference patterns take the form $2\alpha\beta\langle\psi_A, \psi_B\rangle$ with $\alpha, \beta \in \mathbb{R}$ — real-valued modulation with no continuous angular parameter θ. The interference condition (Definition 10.2.1) fails: the cross-term cannot vary continuously as $e^{i\theta}$ for $\theta \in [0, 2\pi)$.

The isotropy condition also fails. The exchange of two independent binary sectors of equal dimension as a continuous symmetry requires a transformation that: (a) maps the σ-sector into the ω-sector and vice versa; (b) preserves norm; and (c) does so via a smooth one-parameter family of transformations interpolating between identity and full exchange.

We first show that no real one-parameter subgroup of $O(4)$ achieves this. 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. The eigenvalues of A are of the form $\pm i\varphi$ for real φ, so the eigenvalues of $\exp(tA)$ are $e^{\pm it\varphi}$. 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, because $\exp(tA)$ restricted to either summand is an $SO(2)$ rotation within that summand. At $t = 0$ we have the identity (σ-sector maps to σ-sector); the smooth path $\exp(tA)$ then keeps the σ-sector mapping to the σ-sector for all t. There is no value of t for which $\exp(tA)$ maps the σ-sector to the ω-sector, because the orbit of the σ-subspace under a real one-parameter orthogonal group lies within the Grassmannian connected component containing the σ-subspace, not crossing to the ω-subspace. More precisely: the Grassmannian $Gr(2,4)$ of 2-planes in $\mathbb{R}^4$ has disconnected components under the $O(4)$ action preserving orientation, and the σ-subspace and ω-subspace lie in different components; no continuous path in $O(4)$ maps between these components, preventing continuous interpolation from the σ-sector to the ω-sector.

A full exchange (swap) of two equal-dimensional subspaces corresponds to a quarter-turn rotation in the combined 4D space. This requires the phase $e^{i\pi/2} = i$ as a rotation between the two planes — complex eigenvalues i and $-i$ acting on the combined space. In $O(4)$, such eigenvalues are unavailable as independent phase rotations on separate 2D subspaces; they require complex linearity. By contrast, $U(4)$ realises this symmetry directly via the complex phase rotation $\mathrm{diag}(1, i, 1, i)$ (in appropriate block form), which maps the σ-sector to the ω-sector with the required i-phase. The isotropy condition therefore cannot be satisfied over $\mathbb{R}$. $\mathbb{R}$ **is eliminated.** $\square_{\mathbb{R}}$

Edge case — $O(2) \cong U(1)$: one might object that $O(2)$ contains a continuous $U(1)$ component and therefore supports a phase degree of freedom. This observation applies only to a single 2D sector in isolation. The required representation acts on $\mathbb{R}^4$ (two independent binary sectors, each 2D, by Proposition 10.1), so the relevant group is $O(4)$, not $O(2)$. $O(4)$ does not contain a subgroup isomorphic to $U(1) \times U(1)$ acting as independent continuous phases on two orthogonal 2D subspaces simultaneously: any $U(1) \subset O(4)$ is a rotation in a single 2D subplane and leaves its orthogonal complement fixed. It cannot simultaneously phase-rotate both the σ-sector and the ω-sector independently. The $O(2)$ observation therefore does not rescue the $\mathbb{R}$ case. $\mathbb{R}$ **is eliminated.** $\square_{\mathbb{R}}$

Elimination of $F = \mathbb{H}$.

The elimination of $\mathbb{H}$ can be made entirely internal to the interface structure. By Proposition 10.1, the complete fact-producing degrees of freedom at any commitment event are exhausted by the four-state interface basis (σ, ω), and all admissible pre-commitment states are superpositions within the minimal representation space $\mathbb{C}^4$. Any extension to $\mathbb{H}$ must therefore act in one of two ways:

Extension beyond the interface basis. This introduces additional degrees of freedom not spanned by the four-state interface basis. By Proposition 10.1 and condition 6, such degrees of freedom are non-minimal unless they generate new admissible fact distinctions. But by construction, the interface basis already exhausts all fact-producing distinctions established in Section 9: σ and ω are the complete set of independent binary geometric degrees of freedom (Lemma 9.3), and the tensor product $F^2 \otimes F^2$ is the unique minimal representation carrying both (Proposition 10.1). No new fact-producing degrees of freedom exist for a quaternionic extension to generate. Therefore such extensions are inadmissible by condition 6.

Action within the $\mathbb{C}^4$ representation. In this case, quaternionic phases act as transformations internal to the existing state space. If these transformations do not change admissible fact outcomes, they are observationally redundant and excluded by condition 6 and the observational equivalence convention. If they do change fact outcomes, they must be representable within the $\mathbb{C}^4$ structure — but then they are already part of the complex representation, contradicting the assumption that they are genuinely quaternionic phases beyond $\mathbb{C}$.

In both cases, $\mathbb{H}$ does not introduce new admissible structure: it either extends beyond the interface basis (inadmissible by condition 6) or acts redundantly within it (eliminated by condition 6). The exclusion of $\mathbb{H}$ follows directly from the completeness of the interface basis established in Proposition 10.1, independent of any external measurement assumptions. This argument is entirely internal to the VERSF framework.

All admissible fact-producing observables are local by construction of $\mathfrak{F}$ (condition 7 and Proposition 7.0), which reinforces the conclusion: any observable capable of detecting genuinely quaternionic phases beyond $\mathbb{C}^4$ would require coupling to degrees of freedom outside the local interface sector, which is already exhausted by the complex representation. Such observables lie outside the admissible class of $\mathfrak{F}$. The Adler–Peres–Stueckelberg results (Adler 1995; Peres 1979; Stueckelberg 1960) confirm this for locally accessible observables but are not the primary load-bearing argument here.

$\mathbb{H}$ is not eliminated because it is mathematically inconsistent, but because within the admissible fact-based theory class it is strictly non-minimal: it introduces no new admissible structure relative to $\mathbb{C}$. **$\mathbb{H}$ is eliminated as non-minimal.** $\square_{\mathbb{H}}$

$\mathbb{C}$ satisfies all conditions.

Over $\mathbb{C}$, the state space is $\mathbb{C}^4$ and the reversible transformation group is U(4), which contains U(1) as a continuous phase subgroup.

- *Interference:* $U(1) \subset U(4)$ provides continuous phase rotations $e^{i\theta}$. The interference term in $|\alpha|\psi_A\rangle + e^{i\theta}\beta|\psi_B\rangle|^2$ varies as $2\mathrm{Re}(\alpha\beta^* e^{i\theta}\langle\psi_A|\psi_B\rangle)$, continuous in θ over $[0, 2\pi)$. ✓
- *Isotropy:* U(4) contains elements permuting the σ and ω sectors symmetrically, realising the $\sigma \leftrightarrow \omega$ exchange without privileging either. ✓

- *Representational invariance:* Unitarily equivalent evolutions on $\mathbb{C}^4$ produce identical fact outcomes. All admissible equivalences are captured within U(4). Any structure beyond U(4) that produces no new fact outcomes is excluded by condition 6. ✓

Minimality. By Hurwitz's theorem, the normed division algebras are exactly $\{\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}\}$. $\mathbb{R}$ is insufficient; $\mathbb{H}$ is non-minimal; $\mathbb{O}$ is non-associative and cannot support a consistent composition of reversible transformations. $\mathbb{C}$ is the unique minimally admissible scalar field. □

Corollary 10.4 (Reversible pre-commitment sector is derived, not assumed)

Any $T \in \mathfrak{F}$ possessing a 2D commitment interface must support a reversible pre-commitment sector with state space $\mathbb{C}^4$ and unitary evolution group U(4) (or its relevant subgroup). This follows from the interface structure and admissibility conditions. It is not assumed as a class membership condition, resolving the circularity present in earlier formulations.

10.5 Definition: The VERSF Fold

The elimination chain of Sections 5–10 has now assembled all the structural components forced by the existence of stable facts in an admissible physical theory. Before stating the uniqueness theorem, these components are collected into a single named definition.

Definition 10.5 (The VERSF Fold)

The *VERSF fold* is a commitment interface satisfying the following jointly necessary and sufficient conditions, each established as a proved necessity in the sections indicated:

1. **Cyclic substrate** (Theorem 5.2): the effective substrate carries nontrivial cycle structure, $\beta_1 \geq 1$.
2. **Intrinsic boundary** (Proposition 6.1, Theorem 6.2): the committed and reversible sectors are separated by a boundary that is operationally intrinsic — not extrinsically labelled and not globally algebraically imposed.
3. **Codimension-1 hypersurface** (Theorem 6.3): the boundary is a separating hypersurface of codimension 1 in the effective substrate manifold.
4. **Local commitment mechanism** (Proposition 7.0, Theorem 7.4): fact formation is locally constructible; no non-local mechanism replaces or supplements the interface.
5. **2D minimal surface** (Theorem 8.3, Proposition 8.5): the interface is topologically S^2 — a compact, connected, orientable 2D surface of genus 0.
6. **Binary geometric data** (Lemma 9.3, Corollary 9.4): the interface carries exactly two independent binary geometric degrees of freedom, the commitment polarity $\sigma \in \{0,1\}$ and the reversible orientation $\omega \in \{-1,+1\}$, giving a four-state geometric structure $(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$.

7. **Complex reversible representation** (Theorems 9.5, 10.0, 10.3): the pre-commitment sector is a linear vector space $\mathbb{C}^4$ — four-dimensional over the field $\mathbb{C}$ — with reversible evolution by unitary transformations $U \in U(4)$.

The VERSF fold is the unique minimal structure satisfying all seven conditions simultaneously within the admissible theory class 𝔉. Up to observational equivalence and field redefinition, no alternative structure satisfies all seven conditions, as established by Theorem 11.1 and Corollary 11.3.

---

11. Master Uniqueness Theorem

Theorem 11.1 (Master Restricted Structural Uniqueness)

Within the admissible theory class 𝔉, any minimal theory supporting stable irreversible local facts for finite observers is necessarily equivalent, up to observational equivalence and field redefinition, to a theory whose minimal fact-supporting structure is the VERSF fold: a 2D intrinsic commitment interface with two independent binary geometric degrees of freedom $(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$, carrying a reversible pre-commitment state space $\mathbb{C}^4$ with unitary evolution.

Proof (Assembly)

Step	Result	Reference
1	Facts are necessary and must be fact-stable	Theorem 3.1, Corollary 3.2
2	Finite distinguishability, irreversibility, finite capacity are necessary	Theorems 4.1–4.3
3	Substrate requires $\beta_1 \geq 1$; external witnesses persist in trees	Lemma 5.3, Theorem 5.2
4	Intrinsic commitment boundary exists	Proposition 6.1
5	𝒮 = 𝒮_accessible ⊕ 𝒮_inaccessible: non-invertible, locally defined, not algebraically projectable	Theorem 6.2
6	Decomposition boundary is codimension-1	Theorem 6.3
7	No non-local commitment mechanism survives	Proposition 7.0
8	Intrinsic interface is the unique minimal realisation	Theorem 7.4
9	Complete structural elimination: all alternative substrate / boundary / dimension / algebra combinations fail	Section 7.5
10	Minimal interface dimension is 2 (Jordan–Brouwer)	Theorem 8.3
11	Interface carries exactly (σ, ω): two independent binary structures, four states	Lemma 9.3
12	Unique minimally admissible scalar field is $\mathbb{C}$	Theorem 10.3

Step	Result	Reference
13	Reversible pre-commitment sector is $\mathbb{C}^4$ with unitary evolution	Proposition 10.1, Corollary 10.4

Any $T \in \mathfrak{F}$ must satisfy steps 1–13. The structure assembled in steps 4–13 — an intrinsic 2D interface with (σ, ω) binary data and $\mathbb{C}^4$ reversible representation — is the VERSF fold architecture. Up to observational equivalence and field redefinition, no other structure satisfies all thirteen requirements simultaneously. □

Remark 11.2 (Scope of what is proved and what is not)

Proved:

- No minimal theory in $\mathfrak{F}$ avoids the 2D fold interface.
- The algebraic representation is uniquely $\mathbb{C}^4$ up to equivalence.
- All identified competitor frameworks fail at a specific named step.
- The admissible class is proved minimal (Section 2.2): no condition can be removed without destroying operational physics.
- The complete space of structural alternatives is explicitly partitioned and each case eliminated (Section 7.5).
- Non-local projection alternatives are eliminated by Proposition 7.0.
- $\mathbb{R}$ elimination: absence of $U(1)$ in $O(n)$. $\mathbb{H}$ elimination: quaternionic phases are observationally redundant under condition 6.

Not proved:

- That the full gauge sector, particle spectrum, and coupling constants are uniquely determined by the fold structure alone.
- That no theory outside $\mathfrak{F}$ could realise stable facts by a radically different mechanism.
- That VERSF's field equations and derived numerical constants follow immediately without additional analysis.

11.2 Why Restricted Uniqueness Is the Correct Target

A common objection to restricted uniqueness results is: "You didn't prove full uniqueness — so what?" This objection reflects a misunderstanding of how structural uniqueness works in physics.

Every significant uniqueness result in physics is restricted to a domain:

- **Wigner's theorem** classifies particles as unitary representations of the Poincaré group — within the class of relativistic quantum theories. It does not classify all possible mathematical structures.

- **Axiomatic QFT** (Wightman axioms) uniquely constrains the structure of local relativistic quantum field theory — within that axiomatic class. It does not claim uniqueness over all formal theories.
- **Effective Field Theory** uniqueness results constrain the operator content of an EFT — given symmetries, power-counting, and UV completion assumptions.
- **General Relativity** uniqueness (Lovelock's theorem) constrains gravitational theories to GR — given four dimensions, diffeomorphism invariance, second-order field equations, and the equivalence principle.

In every case, the domain is not a defect — it is the content of the theorem. A uniqueness claim without a domain says nothing. A uniqueness claim within a well-motivated domain says something precise and useful.

The domain $\mathfrak{F}$ is not arbitrary. Section 2.2 proves that it is minimal: it is the smallest class within which operational physics makes sense. A theory outside $\mathfrak{F}$ either produces no stable facts (in which case it is not physics) or produces stable facts by a mechanism that, by the arguments of Sections 3–6, instantiates the same structural requirements. The restriction is therefore not a concession but a precision: the uniqueness theorem is as strong as a uniqueness theorem in physics can legitimately be. Any claim of "full uniqueness" without a defined admissible class is mathematically ill-posed — it is a statement about all conceivable formal systems, not about physics.

The remaining open question is not whether an alternative minimal architecture exists within $\mathfrak{F}$ — Corollary 11.3 below proves that none does — but whether the fold architecture, once identified, uniquely determines the full downstream physics of the Standard Model and gravity. That is the next task in the programme.

Corollary 11.3 (No Alternative Minimal Architecture)

Any theory $T \in \mathfrak{F}$ that does not contain a VERSF-equivalent fold structure must violate at least one proved necessity.

Proof

Suppose $T \in \mathfrak{F}$ satisfies all conditions of $\mathfrak{F}$ but does not contain the VERSF fold. Then T must lack at least one of the following structural features, each of which is necessary:

Missing feature	Proved necessary by
Cyclic substrate ($\beta_1 \geq 1$)	Theorem 5.2
Intrinsic commitment boundary	Proposition 6.1
Non-invertible state-space decomposition	Theorem 6.2 (property 1)
Locally defined separation	Theorem 6.2 (property 2)
Codimension-1 interface	Theorem 6.3
Local (not non-local) commitment	Proposition 7.0

Missing feature	Proved necessary by
2D interface (not 1D or 3D+)	Theorem 8.3
Binary data (σ, ω) — exactly four states	Lemma 9.3
Scalar field ℂ (not ℝ, ℍ, or 𝕆)	Theorem 10.3
Reversible pre-commitment sector $\mathbb{C}^4$	Proposition 10.1, Corollary 10.4

Each missing feature corresponds to a violation of a proved theorem. Since T ∈ 𝔉 satisfies all class conditions, any such violation is a contradiction. Therefore T cannot lack any of these features — T must contain the VERSF fold. No alternative minimal architecture exists within 𝔉. □

---

12. What Would Falsify This Framework

A theory that cannot be falsified is not a physical theory. The following conditions would refute the framework. They are not all equally near-term: some are directly testable with existing or near-future experimental methods; others are structural claims that would require remarkable and currently inconceivable physical systems to test, but are nonetheless genuine falsifiability conditions in the sense that their realisation would constitute a logical refutation of a proved theorem. Both categories are listed, distinguished by their practical accessibility.

Genuinely near-term testable:

F3 — Distinguishable quaternionic phases. Detection of physically distinguishable quaternionic phases in local measurements accessible to finite observers in bounded regions — phases produced by right-multiplication by j or k in ℍ, generating outcomes distinct from any $\mathbb{C}^4$ evolution — would refute the ℍ-elimination argument in Theorem 10.3. This has been extensively sought in experimental tests of quaternionic quantum mechanics (Kaiser, George, and Warner 1981; Peres 1979) with null results to date. This is the most directly testable condition: precision quantum interferometry in principle allows sensitive searches for quaternionic deviations from standard $\mathbb{C}^4$ predictions.

F4 — Admissible non-local commitment. Identification of a fact-forming mechanism that is genuinely non-local — requiring degrees of freedom outside any finite observer's causal reach — yet is not eliminable by observational equivalence (i.e., it produces fact outcomes distinguishable from those of any local interface) would refute Proposition 7.0. Quantum non-locality experiments (Bell inequality violations and their successors) are relevant context, though current results are consistent with local commitment; a definitive refutation would require a non-local mechanism producing fact outcomes not reproducible by any local description.

Structural falsifiability conditions (genuine logical refutations requiring currently inconceivable physical systems, but not thereby vacuous):

F1 — Locally reversible stable facts. Observation of stable, reproducible fact classes that are locally reversible for finite bounded observers would directly refute Theorem 4.2. This is the foundation of the entire chain: if facts can be locally undone, the entire topological and geometric argument collapses. Such an observation would require constructing a physical record that persists and is reproducible, yet can be completely reversed by a bounded local operation without accessing the environment.

F2 — Fact formation on tree-like substrates. Construction of a physical system that forms stable, locally irreversible records on an effectively tree-structured information substrate — one with $\beta_1 = 0$ — would refute Theorem 5.2. This would require demonstrating that the unique-path property of trees does not prevent local trapping, i.e., that Lemma 5.3 fails for some physical substrate. While this is in principle testable for specific substrate architectures, no known physical system exhibits tree-structured information substrates at scales relevant to fact formation.

F5 — Intrinsic 1D commitment boundary. Construction of a 1D interface that encodes commitment polarity without reference to any ambient space — i.e., a 1D manifold that intrinsically separates two commitment sectors — would refute Proposition 8.1 and Theorem 8.3. This would require overturning the application of the Jordan Curve Theorem to the interface dimension argument, which would constitute a result in pure mathematics as well as physics.

F6 — Third independent binary geometric structure. Identification of a third independent binary geometric structure on the minimal 2D interface — one not reducible to σ or ω by the completeness argument of Lemma 9.3 — would unsettle Corollary 9.4 and force a re-examination of the four-state basis. This would have to involve a binary datum not captured by: orientation, domain assignment, spin structure (on S^2), framing, or chirality.

Each of these six conditions is a genuine falsifiability condition in the sense that its realisation would constitute a logical refutation of a specific proved theorem. The distinction between near-term and structural conditions reflects practical experimental accessibility, not the strength of the logical commitment: all six would refute the framework if realised.

---

13. Open Tasks

1. **Extension to gauge sector.** Whether the specific gauge groups of the Standard Model follow uniquely from the $\mathbb{C}^4$ fold structure — or require additional constraint analysis — is the primary remaining open question in the uniqueness programme.
2. **Ambient substrate dimension.** The present results establish that the commitment interface is 2-dimensional and codimension-1 in the effective substrate manifold (Lemma 5.6, Theorem 6.3), which implies that the ambient substrate dimension n satisfies $n \geq 3$. However, the results do not uniquely determine n: the elimination chain is compatible with $n = 3$, 4, or higher. The derivation of the ambient substrate dimension from the admissibility conditions of $\mathfrak{F}$ remains open and is not claimed here. This has a direct empirical consequence: a 3D, 4D, or higher-dimensional substrate would differ at the

level of observable physics (field content, particle spectrum, causal structure), so any specific instantiation of the VERSF programme carries an empirical commitment to a particular substrate dimension not yet fixed by the present results. The determination of n is therefore both an open theoretical task and an empirical commitment of the framework.

3. **Cycle topology specification.** Theorem 5.2 requires $\beta_1 \geq 1$ but does not specify the cycle structure beyond this minimum. Whether the specific topological features of the VERSF substrate are forced by additional minimality arguments, or whether a family of cycle topologies are compatible with $\mathfrak{F}$, requires further analysis.
4. **Quantitative coupling constants.** The derivation of numerical predictions (fine-structure constant, proton-to-electron mass ratio, cosmological constant) from the uniquely derived interface structure requires further steps beyond structural equivalence.
5. **Non-minimal theories within $\mathfrak{F}$.** Theories in $\mathfrak{F}$ satisfying all conditions but carrying redundant additional structure are not eliminated by Theorem 11.1. Whether they converge empirically to VERSF predictions, or are further eliminable by a tighter minimality criterion, warrants investigation.
6. **Relaxation of admissibility conditions.** Whether weakening any of conditions 1–7 opens space for alternative architectures — and whether the uniqueness result is stable under small perturbations of those conditions — remains open.

---

14. Conclusion

This paper began with a single primitive: facts are not a feature of physics — they are its precondition. That primitive is made formally precise in Theorem 3.1 and elevated to a biconditional in Corollary 3.2: a universe admits a physical theory if and only if it admits a class of stable, recoverable, reproducible facts. From that equivalence alone — with no prior commitment to quantum mechanics, spacetime geometry, or VERSF-specific primitives — a strict thirteen-step derivation chain establishes the complete structure of the VERSF fold.

The chain runs: physics → facts → time (emergent) → causal order (emergent) → finite distinguishability → irreversible commitment → finite localization capacity → cyclic substrate ($\beta_1 \geq 1$) → intrinsic codimension-1 commitment interface → 2D minimal surface → binary geometric data (σ, ω) → scalar field $\mathbb{C}$ → VERSF fold. Each arrow is a proved theorem. No step is assumed. The starting point is not quantum mechanics, not spacetime, not symmetry — it is the single requirement that physics produces stable, recoverable, reproducible facts.

The admissible theory class $\mathfrak{F}$ is proved minimal (Section 2.2): removing any condition collapses operational physics. The complete space of structural alternatives is partitioned and exhaustively eliminated (Section 7.5), covering all combinations of substrate topology, boundary type, interface dimension, and scalar field. The only surviving combination is the VERSF fold. Corollary 11.3 states the kill theorem: any theory in $\mathfrak{F}$ not containing the VERSF fold must violate at least one proved necessity. Section 11.2 establishes that restricted uniqueness within $\mathfrak{F}$ is the correct and standard target for structural uniqueness results. Section 12 identifies six falsifiability conditions defining the empirical commitments of the framework.

The result establishes the unique minimal fact-supporting architecture within $\mathfrak{F}$: within the only class within which operational physics is possible — the class of theories that produce stable, recoverable, reproducible facts — any admissible theory must contain the VERSF fold as a necessary structural core. More precisely, the present result establishes a necessary-condition uniqueness: any theory producing stable facts must contain the fold architecture. It does not yet establish sufficient-condition uniqueness for the full physical theory above this structure — that is, it does not yet prove that the fold architecture, together with the admissibility conditions, is sufficient to uniquely determine all downstream phenomenology including gauge structure, particle content, and coupling constants. The fold is not the whole of physics; it is the part of physics that the existence of stable facts makes unavoidable. The specific open tasks constituting the remaining programme are identified in Section 13 as items 1 (extension to gauge sector), 2 (ambient substrate dimension), and 3 (cycle topology specification) — these are the three primary gaps between the structural uniqueness proved here and the full empirical content of the Standard Model.

References

Topology and algebra

Hatcher, A. (2002). *Algebraic Topology*. Cambridge University Press. [Jordan–Brouwer Separation Theorem, Alexander Duality, Schoenflies theorem]

Munkres, J.R. (2000). *Topology* (2nd ed.). Prentice Hall. [Topological manifolds, Jordan Curve Theorem]

Hewitt, E. & Ross, K.A. (1963). *Abstract Harmonic Analysis, Vol. I*. Springer. [Pontryagin–van Kampen structure theorem for locally compact abelian groups]

Hurwitz, A. (1923). "Über die Komposition der quadratischen Formen." *Mathematische Annalen* 88, 1–25. [Classification of normed division algebras: $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$, $\mathbb{O}$]

Quaternionic quantum mechanics

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

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

Stueckelberg, E.C.G. (1960). "Quantum theory in real Hilbert space." *Helvetica Physica Acta* 33, 727–752.

Causal sets

Bombelli, L., Lee, J., Meyer, D. & Sorkin, R.D. (1987). "Space-time as a causal set." *Physical Review Letters* 59, 521–524.

Henson, J. (2006). "The causal set approach to quantum gravity." In Oriti, D. (ed.), *Approaches to Quantum Gravity*. Cambridge University Press.

VERSF programme

Taylor, K. (2024–2025). VERSF Research Papers (140+ papers). AIDA Institute. versf-eos.com