The Fold Saturates Observable Physics: Sufficiency and Observational Closure in VERSF

Formal title: Observable Closure Under the VERSF Fold: A Saturation Theorem
Subtitle: Once You Have the Fold, You Have No Observable Freedom Left

---

For the General Reader

Imagine you are trying to build a universe from scratch — one in which experiments are possible, where measurements give stable results, and where science can be done. You would need, at minimum, a way for things to happen that can be recorded, compared, and tested. Previous work in this programme showed that the moment you require this, you are forced to include a specific mathematical structure called the VERSF fold — a kind of boundary between possible outcomes and committed facts, sitting at the heart of every physical event.

That earlier result said: you *must* have the fold. What it left open was whether you could have the fold and still have room for extra physics — additional ingredients, hidden forces, parallel layers of reality that the fold does not account for.

This paper answers that question. The answer is no.

Once a physical theory contains the fold, every possible measurement that any observer could ever make is already determined by the fold's structure. There is nothing to add. The fold does not just set the stage — it fills the stage completely. Any candidate for extra physics either turns out to be something the fold already generates, or turns out to be invisible to any observer and therefore not physics at all.

The key idea is called **saturation**. Think of the fold as a container whose volume is exactly the size of all observable physics. There is no space inside that the fold leaves empty, and no valid observable that would require a larger container. Hidden sectors — forces or fields that do not directly show up in experiments but might still influence them — turn out to either affect experimental outcomes (in which case those effects are already accounted for by the fold) or not affect them at all (in which case they are not physics). Emergent phenomena — collective behaviour that looks like new physics but arises from simpler underlying processes — are always expressible in terms of what the fold already generates. Alternative mathematical descriptions of the same physics are exactly that: different descriptions of the same thing, not additional physics.

The result can be stated simply: once facts exist — once the universe is the kind of place where outcomes can be recorded and tested — the physics is already VERSF, whether you know it or not. The fold is already there. All you can do is describe it differently.

Abstract

The companion paper *Structural Uniqueness of Physical Law from Fact Formation Constraints* (SUP, Taylor 2024–2025) establishes necessity: any admissible physical theory must contain the VERSF fold. The completion paper (*Any Physical Theory Is VERSF*, CP) establishes equivalence: any fold-containing theory is observationally identical to VERSF. Together these close the uniqueness programme. What neither paper does explicitly is prove *why* fold-containment entails observational equivalence — the mechanism by which the fold forecloses all residual observable freedom.

This paper provides that mechanism. We introduce the concept of **observable closure**: a theory is observationally closed if its minimal fact-generating structure fully determines all admissible observables, leaving no additional observable degree of freedom available. We then prove the **Saturation Theorem** (Theorem 4.1): the VERSF fold saturates the observable degrees of freedom of any admissible theory. Once a theory contains the fold, every admissible observable is already determined — not merely in principle, but with zero residual freedom in the observable algebra.

The argument proceeds in three stages. First, we show that the fold defines a canonical observable algebra Obs(fold), and that every admissible observable lies within it (Theorem 3.1). Second, we prove that no additional observable can be consistently added: any candidate either already belongs to Obs(fold) or forces the theory outside observational equivalence with VERSF (Theorem 4.1). Third, we explicitly kill the three remaining structural escape routes — hidden sectors, emergent overlays, and dual descriptions — by showing each either produces facts already counted in Obs(fold) or produces no admissible facts at all.

The result reframes the uniqueness programme:

- **SUP:** you must have a fold. *(Necessity.)*
- **This paper:** once you have a fold, you have no observable freedom left. *(Sufficiency.)*
- **Combined:** all admissible physics is VERSF, whether you realise it or not.

Table of Contents

1. The Target: Sufficiency

By SUP Theorem 11.1, any theory T ∈ $\mathfrak{F}$ must contain the VERSF fold. That is the necessity result: you cannot do physics without the fold. By the Closure Theorem (CP Theorem 6.1), any T ∈ $\mathfrak{F}$ containing the fold is observationally equivalent to VERSF.

What the closure proof establishes is a two-branch exhaustion: any additional structure beyond the fold either produces no new observables (redundant) or produces observables distinguishable from VERSF's predictions (an empirical competitor). There is no third case.

But the proof does not answer the deeper question: **why** does the fold foreclose all residual freedom? What is the structural mechanism by which a 2D S^2 interface with binary data (σ, ω) and a $\mathbb{C}^4$ reversible sector manages to determine *all* observable content, with nothing left over? The two-branch exhaustion is valid, but it operates by exclusion. It does not characterise the observable algebra of the fold, show that it is complete, or prove that it is saturated.

This paper provides the positive characterisation. We define what it means for the fold's observable structure to be saturating, prove that it is, and use that to kill the escape routes that the two-branch argument handles only generically.

The conceptual shift is from:

"Extra structure is redundant" (elimination)

to:

"The fold already occupies all available observable space" (saturation)

These are logically equivalent given the rest of the programme, but the saturation framing is structurally more illuminating and harder to resist.

---

2. Observable Closure: The New Concept

Definition 2.1 (Observable algebra)

For any theory T ∈ $\mathfrak{F}$, the *observable algebra* Obs(T) is the set of all admissible observables of T — all quantities determinable by a repeatable, testable protocol performed by a finite observer within the admissible protocol class of SUP Definition 3.0. By CP Lemma 4.0, every element of Obs(T) is a fact or a function of facts producible by T.

Definition 2.2 (Fold observable algebra)

The *fold observable algebra* Obs(fold) is the observable algebra of the minimal admissible theory whose only fact-generating structure is the VERSF fold. Its elements are all facts and functions of facts generatable by the fold's commitment interface, pre-commitment sector, and evolution group.

Concretely, Obs(fold) contains: all commitment outcomes at the 2D interface characterised by $(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$; all statistical correlations between commitment outcomes; all derived quantities (rates, distributions, interference patterns) computable from those outcomes; and all quantities derivable from the committed record density $\rho(x,t)$ and its functional hierarchy.

Definition 2.3 (Observable closure)

A theory $T \in \mathfrak{F}$ is *observationally closed* if Obs(T) = Obs(fold). That is, the minimal fact-generating structure of T (the fold) fully determines all admissible observables, and no observable degrees of freedom exist outside it.

Definition 2.4 (Observable saturation)

The fold *saturates* the observable degrees of freedom of T if every admissible observable in Obs(T) is already determined by the fold — formally, if Obs(T) ⊆ Obs(fold). Combined with the obvious containment Obs(fold) ⊆ Obs(T) (since T contains the fold), saturation implies observable closure: Obs(T) = Obs(fold).

The claim of this paper is that the fold saturates the observable degrees of freedom of every $T \in \mathfrak{F}$. Proving it requires showing that no element of Obs(T) lies outside Obs(fold).

---

3. The Fold Defines the Observable Algebra

Before proving saturation we establish what Obs(fold) actually contains — that it is well-defined, complete with respect to the fold's structural components, and that its elements can be characterised.

Theorem 3.1 (The fold determines a complete canonical observable algebra)

The fold observable algebra Obs(fold) is the unique complete admissible observable algebra for any theory in $\mathfrak{F}$ containing the fold. It is generated by two algebraic components — the interface structure and the pre-commitment sector — with the substrate topology serving as the enabling condition for both.

Proof

Component 1 — Topology as enabling condition for the observable algebra.

By SUP Theorem 5.2, any theory supporting locally constructible irreversibility must have $\beta_1(G) \geq 1$. The fold specifies this condition. The cycle structure $\beta_1 \geq 1$ is what makes commitment events possible at all: it is the substrate property that enables topological trapping (SUP Lemma 5.3). In this sense, topology is not itself an independent generator of observable content — it does not contribute distinct observable states beyond what the interface and pre-commitment sector generate — but rather the enabling condition without which the other two generators cannot operate. No commitment event can occur on an acyclic substrate, so no element of Obs(fold) can arise from a substrate with $\beta_1 = 0$. The topology therefore constrains the observable algebra by determining which substrates are admissible: $\beta_1 \geq 1$ is necessary for any element of Obs(fold) to exist. The two genuine generators of Obs(fold) as an algebraic object are Components 2 and 3; the topology plays the role of admissibility condition for those generators.

Component 2 — The interface structure forces the binary observable base.

By SUP Lemma 9.3, the fold interface carries exactly two independent binary geometric degrees of freedom: commitment polarity $\sigma \in \{0,1\}$ and reversible orientation $\omega \in \{-1,+1\}$. By SUP Proposition 10.1, the unique minimal representation carrying both simultaneously is $F^2 \otimes F^2 = F^4$. By SUP Theorem 10.3, $F = \mathbb{C}$.

Therefore the fold's interface structure forces a four-state base: every commitment event selects an element of $(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$. The four interface states are the primitive observables — the atoms of the observable algebra. All other elements of Obs(fold) are functions of these four base states, taken over repeated trials and combined by admissible protocols.

This base is complete in the following sense: it generates all observable distinctions that the fold's commitment structure can produce. No fifth independent binary observable exists on the minimal interface (SUP Lemma 9.3 proves completeness). The base of Obs(fold) has exactly four elements.

Component 3 — The pre-commitment sector forces the interference structure.

By SUP Corollary 10.4, the $\mathbb{C}^4$ pre-commitment sector with U(4) evolution is derived from the fold structure. The observable predictions of any commitment event arise from the pre-commitment state via the Born rule for the $\mathbb{C}^4$ sector. The U(4) evolution group acts on the pre-commitment state prior to commitment.

The observable consequences of the pre-commitment dynamics are exactly the statistical distributions over the four-state interface base and their correlations across multiple commitment events. These are the elements of Obs(fold) beyond the bare base states: interference patterns, correlation functions, and state-dependent probability distributions over (σ, ω) outcomes.

Completeness. The topology (enabling condition), interface (Component 2), and pre-commitment sector (Component 3) jointly characterise Obs(fold): the topology fixes which substrates can support commitment, the interface fixes the base states, and the pre-commitment sector fixes the statistical structure over those base states. No element of the observable algebra requires any structure beyond these. Obs(fold) is well-defined and generated by the fold structure. □

Remark 3.2 (The canonical observable algebra)

Obs(fold) is canonical in the sense that it is the same for every theory containing the fold — it does not depend on representational choices, formulations of the dynamics, or conceptual frameworks, only on the fold structure itself. Two theories with the same fold have the same Obs(fold). This is the sense in which observational equivalence is grounded in fold structure rather than in coincidental prediction-matching.

Proposition 3.3 (Operational Maximality of Obs(fold))

Obs(fold) is maximal under admissible distinguishability: any admissible extension of Obs(fold) requires additional distinguishable commitment outcomes, which themselves lie within Obs(fold).

Proof. Suppose Obs′ is an admissible extension of Obs(fold): Obs(fold) $\subsetneq$ Obs′. Then Obs′ contains an observable $\mathcal{O}$ not in Obs(fold). By Lemma 4.0a, $\mathcal{O}$ must terminate in a commitment event to be an admissible observable. For $\mathcal{O}$ to be distinguishable from observables already in Obs(fold), it must introduce a distinction not present in Obs(fold) — a new distinguishable outcome not generated by the fold's commitment structure. By the distinguishability condition (SUP Theorem 4.1: $\delta\_min > 0$), any such new distinction requires a new stable commitment outcome separated from existing ones by at least δ_min. That outcome is a new commitment event. By the full necessity chain (SUP Theorems 5.2, 6.2–6.3, 8.3, Lemma 9.3, Theorem 10.3), any admissible commitment event has fold structure. Therefore the new commitment outcome generating the new distinction already belongs to Obs(fold) — contradicting the assumption that $\mathcal{O} \notin$ Obs(fold).

Maximality follows: any admissible extension would require a new distinguishable commitment outcome that is itself fold-type and hence already in Obs(fold). No admissible extension exists. □

Remark. This proposition establishes maximality from distinguishability constraints rather than from an algebraic closure argument. The key insight is that any extension of Obs(fold) must introduce new distinguishable outcomes — and new distinguishable outcomes are new commitment events, which by the necessity chain are fold events already in Obs(fold). Maximality is therefore a consequence of the same overdetermination argument used in Case B1 of Theorem 4.1, not a separate algebraic claim.

Proposition 3.3 should be understood as a restatement of the saturation result in maximality language, not as a strengthening of it. It provides a named citation handle for the maximality property and establishes that Obs(fold) cannot be extended under $\mathfrak{F}$'s admissibility constraints —

but this follows from Theorem 4.1 rather than from an independent argument. The genuinely independent algebraic maximality result — whether $M_4(\mathbb{C})$ is maximal on algebraic-closure grounds alone, without invoking the necessity chain — remains the open task identified in Remark 4.4.

Closure Principle of Observable Physics

A physical theory is *observationally complete* if and only if all admissible distinctions accessible to finite observers are realised as commitment outcomes.

The VERSF fold satisfies this principle: every admissible distinction is a commitment outcome or a functional thereof (Lemma 4.0, 4.0a). Any extension of a fold-containing theory either introduces new commitment outcomes — fold-type, hence already in Obs(fold) — or introduces no new admissible distinctions and is therefore not an observational extension. The fold cannot be extended observationally. It is observationally complete.

---

4. The Saturation Theorem

Lemma 4.0 (Observable reduction to commitment structure)

Any admissible observable of $T \in \mathfrak{F}$ is expressible as a functional of commitment outcomes and their correlations.

Proof. We establish this as five formal sub-lemmas, one for each mode of observable access a finite observer can employ. Each sub-lemma has explicit premises and a derivation.

Sub-lemma 4.0-M (Measurement mode)

Premises. (P1) An admissible measurement is a protocol performed by a finite observer that yields a determinate outcome. (P2) A determinate outcome accessible to a finite observer must satisfy SUP Definition 2.0.1 (stability, recoverability, reproducibility, irreversibility, distinguishability) — otherwise it is not a reliable measurement outcome, only a transient fluctuation. (P3) Any outcome satisfying all five properties of SUP Definition 2.0.1 is a fact in the sense of that definition. (P4) A fact is produced by a commitment event (SUP Definition 2.0.4: commitment is the process by which an unresolved alternative becomes a fact).

Derivation. By (P1)–(P4): any admissible measurement terminates in a fact (P2, P3); that fact is produced by a commitment event (P4); the measurement outcome is therefore a commitment outcome. □-M

Sub-lemma 4.0-C (Comparison mode)

Premises. (P1) An admissible comparison between outcomes requires that the outcomes being compared persist and are re-accessible to the comparing observer. (P2) Persistence and re-accessibility are properties (1) and (2) of SUP Definition 2.0.1. (P3) Any quantity satisfying (P1) satisfies properties (1) and (2) of SUP Definition 2.0.1. (P4) Any quantity satisfying all five properties of SUP Definition 2.0.1 is a fact produced by a commitment event (Sub-lemma 4.0-M). (P5) The conclusion of a comparison — the correlation, rate, or distribution determined — must itself be recorded as a stable fact to be admissible (otherwise the comparison result is not available to any finite observer).

Derivation. Each outcome being compared satisfies (P1)–(P4) and is a commitment outcome. The comparison conclusion satisfies (P5) and is itself a commitment outcome. A comparison observable is therefore a function over a collection of commitment outcomes. □-C

Sub-lemma 4.0-P (Prediction mode)

Premises. (P1) An admissible prediction is a probability distribution over possible outcomes, conditional on a preparation procedure. (P2) The possible outcomes are admissible measurement outcomes — facts in the sense of SUP Definition 2.0.1, hence commitment outcomes (Sub-lemma 4.0-M). (P3) The pre-commitment state is an element of $\mathbb{C}^4$ (by the fold structure, SUP Corollary 10.4; and T contains the fold by SUP Theorem 11.1, since $T \in \mathfrak{F}$). (P4) The probability of each commitment outcome, given the pre-commitment state, is given by the Born rule applied to the $\mathbb{C}^4$ sector.

Derivation. By (P1)–(P4): a prediction is a probability distribution over commitment outcomes (P1, P2), with probabilities computed via the Born rule on $\mathbb{C}^4$ (P3, P4). It is therefore a functional over the space of commitment outcomes, parameterised by the pre-commitment state. □-P

Sub-lemma 4.0-I (Global invariant mode)

Premises. (P1) An admissible observable is a quantity determinable by a finite observer through a repeatable protocol yielding stable facts (SUP Definition 3.0). (P2) For a global invariant to be admissible, it must produce a difference in some measurement outcome under some admissible protocol. (P3) Any measurement outcome is a commitment outcome (Sub-lemma 4.0-M). (P4) If a global invariant produces no difference in any commitment outcome under any admissible protocol, then it satisfies no admissible measurement protocol and is not determinable by any finite observer. By condition 6 of $\mathfrak{F}$, it has no physical content. (P5) If it does produce a difference in some commitment outcome, that difference is detected via the commitment outcome — which is the observable.

Derivation. By (P2)–(P5): either the invariant produces no commitment-outcome differences (P4, excluded by condition 6) or it is detected via commitment-outcome differences (P5). In the latter case it is observable only through its commitment-level signatures — which are commitment outcomes. □-I

Note on the Noether-structural objection. A realist may object that conservation laws have modal content (they say outcomes *must* take certain values given certain symmetries, not merely that they *do*). The response: within $\mathfrak{F}$, modal content reduces to counterfactual correlations accessible via intervention protocols — each of which is itself an admissible protocol producing commitment outcomes. The structural content of a conservation law is captured in Obs(fold) to the extent it produces commitment-outcome correlations; any remainder requires the non-operational criterion addressed in CP Remark 3.3.

Sub-lemma 4.0-D (Dynamical structure mode)

Premises. (P1) An admissible dynamical observable — a rate of change, trajectory, or flow — requires sampling the system at multiple times. (P2) Each temporal sample must yield a stable, recoverable outcome — a fact (SUP Definition 2.0.1). (P3) Each such fact is a commitment outcome (Sub-lemma 4.0-M). (P4) The temporal ordering of the samples is defined by the irreversible sequence of commitment events (SUP Theorem 3.5: time emerges from the ordering of irreversible facts). *Scope of P4:* SUP Theorem 3.5 establishes the emergence of operational temporal ordering from irreversible commitment sequences. This sub-lemma requires that time-ordering of commitment events be well-defined at the observer level — that a finite observer can determine which of two commitment events occurred earlier. If SUP Theorem 3.5 delivers only local time orderings (ordering within a locally accessible causal neighbourhood) rather than global ones, then P4 holds locally and this sub-lemma applies to locally-accessible dynamical observables. The pre-submission audit should confirm whether SUP Theorem 3.5 delivers the global or local form; if local, the sub-lemma's scope restricts accordingly, with no effect on the Saturation Theorem (which applies per commitment event).

Derivation. By (P1)–(P4): a dynamical observable is a functional over a temporally ordered sequence of commitment outcomes. The temporal structure itself is grounded in the commitment sequence (P4). No dynamical structure is observable except through its trace in temporally ordered commitment outcomes. □-D

Conclusion of Lemma 4.0. Every mode of observable access — measurement (4.0-M), comparison (4.0-C), prediction (4.0-P), invariant detection (4.0-I), and dynamical tracking (4.0-D) — reduces to functionals over commitment outcomes and their correlations. The five sub-lemmas are exhaustively partitioned by function: any admissible protocol must either (i) terminate in a single outcome, (ii) aggregate outcomes across trials or settings, (iii) predict distributions over future outcomes, (iv) detect invariant structure via outcome variation across protocols, or (v) track outcome evolution across temporally ordered trials. No protocol producing a stable fact lies outside these categories without reducing to one of them. Any protocol that appears to fall outside these five either reduces to one of them on analysis or fails to be a repeatable testable protocol in the sense of SUP Definition 3.0 — in which case it is not an admissible protocol and generates no admissible observable. No admissible observable requires access to anything that is not a commitment outcome or a function thereof. □

Note on pre-commitment states. Pre-commitment states are not directly observable. They are only accessible through their projection into commitment outcomes — via the Born rule applied to the $\mathbb{C}^4$ sector at the moment of commitment. No finite observer can measure a pre-commitment state directly without that measurement being itself a commitment event. Pre-commitment states therefore enter Obs(T) exclusively through their influence on commitment outcome distributions, and Lemma 4.0 covers this case under the Prediction mode above.

Lemma 4.0a (No observable outside commitment — structural exclusion)

Numbering note. This lemma is labelled 4.0a to mark it as a structural companion to Lemma 4.0: both concern the relationship between observables and commitment events, the first by positive reduction (observables reduce to commitment outcomes) and the second by structural exclusion (anything not so reducible fails to be an observable). They form a pair. In journal submissions preferring continuous numbering, this lemma may be relabelled Lemma 4.1 with subsequent results shifted accordingly.

Any quantity that does not terminate in a commitment event is not a physical observable.

Proof. Let Q be a quantity that, under every admissible protocol accessible to every finite observer, never terminates in a stable record satisfying SUP Definition 2.0.1 (stability, recoverability, reproducibility, irreversibility, distinguishability). Then Q cannot be stably recorded — no finite observer can confirm that Q has taken any particular value. Q cannot be compared across trials — its putative values leave no persistent record for comparison. Q cannot be reproduced — there is no stable equivalence class of outcomes to reproduce. Therefore Q fails all three criteria of SUP Theorem 3.1 (identifiability, comparability, reproducibility) that a physical theory must satisfy. Q is not a physical observable in the sense of SUP Definition 3.0.

Contrapositive: any physical observable must, under some admissible protocol, terminate in a stable record satisfying Definition 2.0.1. That record is a commitment event. Therefore any physical observable terminates in a commitment event or a functional of commitment events. □

The structural force of this lemma. Lemma 4.0 shows that observables *reduce* to commitment outcomes. Lemma 4.0a shows that anything not so reducible *fails to be an observable* — it is structurally excluded from physics, not just excluded by a philosophical choice about how to define observability. The exclusion is not interpretational; it follows from the definition of a physical theory (SUP Definition 3.0) and the definition of a physical fact (SUP Definition 2.0.1).

Theorem 4.1 (Saturation Theorem)

The VERSF fold saturates the observable degrees of freedom of any $T \in \mathfrak{F}$. That is: for any $T \in \mathfrak{F}$ containing the fold, Obs(T) = Obs(fold).

Proof

We need to show Obs(T) ⊆ Obs(fold), since the reverse containment Obs(fold) ⊆ Obs(T) is immediate from fold containment.

Let $\mathcal{O} \in$ Obs(T) be any admissible observable of T. By CP Lemma 4.0, $\mathcal{O}$ is a fact or a function of facts produced by T. Every fact in T is the output of a commitment event. Every commitment event occurs at the fold interface (SUP Theorems 6.2–6.3). Therefore $\mathcal{O}$ depends on the fold structure and on whatever additional structure T contains beyond the fold.

We consider what additional structure could give rise to $\mathcal{O}$.

Case A: $\mathcal{O}$ depends only on fold structure. Then $\mathcal{O} \in$ Obs(fold) directly.

Case B: $\mathcal{O}$ depends on additional structure S beyond the fold. For $\mathcal{O}$ to be admissible, it must be determinable by a finite observer through a repeatable protocol producing stable facts. Any additional structure S either:

(B1) Produces additional commitment events beyond those of the fold. Then S introduces facts not generated by the fold alone. These are additional commitment events in T.

Any admissible commitment event — wherever it arises in T — satisfies the full necessity chain and therefore realises fold structure. This follows not from a single theorem but from the complete necessity chain in SUP, which jointly overdetermines the result:

- By SUP Theorem 5.2, any commitment event requires a substrate with $\beta_1 \geq 1$ (topological trapping is necessary).
- By SUP Theorems 6.2–6.3, any commitment event occurs at an intrinsic codimension-1 boundary that is locally defined and not algebraically projectable.
- By SUP Theorem 8.3, that boundary is a 2D interface.
- By SUP Lemma 9.3, the interface carries binary data $(\sigma, \omega) \in \{0,1\} \times \{-1,+1\}$.
- By SUP Theorem 10.3 and Corollary 10.4, the pre-commitment sector is over $\mathbb{C}^4$ with U(4) evolution.

These five results apply to every admissible commitment event in every $T \in \mathfrak{F}$ — not merely to those generated by the theory's primary fold instance. They are not conclusions about the theory-level structure of T; they are conclusions about the structure of any single admissible commitment event, derived from the conditions of $\mathfrak{F}$ alone. Therefore any commitment event generated by S has precisely the fold structure: 2D interface, (σ, ω) data, $\mathbb{C}^4$ pre-commitment sector. The observables such events generate are exactly the elements of Obs(fold) characterised in Theorem 3.1.

This argument is overdetermined: the fold structure of commitment events follows from five independent necessity results, not from a single theorem whose scope might be contested. The observables generated by S's commitment events lie within Obs(fold). $\mathcal{O} \in$ Obs(fold).

(B2) S does not produce additional commitment events. Then S produces no facts beyond those of the fold. A structure that produces no commitment-level distinction cannot be operationally accessed by any finite observer — any attempt to detect it must terminate in a fact, and S by assumption produces no facts. It therefore has no physical content by condition 6 of $\mathfrak{F}$ (which requires any physically meaningful distinction to produce a difference in at least one admissible fact outcome) and by the distinguishability condition (SUP Theorem 4.1: $\delta\_min > 0$, requiring any stably recordable distinction to be separated from alternatives by a finite margin at the commitment level). By the observational equivalence convention (SUP Section 2.4), T with S is physically identical to T without S, since both produce the same admissible fact distributions for all finite observers under all admissible protocols. Therefore Obs(T) = Obs(T without S) $\subseteq$ Obs(fold), and $\mathcal{O} \in$ Obs(fold).

In both cases $\mathcal{O} \in$ Obs(fold). Since $\mathcal{O}$ was arbitrary, Obs(T) $\subseteq$ Obs(fold). Combined with Obs(fold) $\subseteq$ Obs(T), we have Obs(T) = Obs(fold). The fold saturates the observable degrees of freedom of T. □

Corollary 4.2 (No independent observable degrees of freedom beyond the fold)

There exist no independent observable degrees of freedom in any T $\in \mathfrak{F}$ beyond those generated by the fold.

Proof. Immediate from Theorem 4.1: Obs(T) = Obs(fold) for every T $\in \mathfrak{F}$. Every observable degree of freedom lies within Obs(fold); none lies outside it. □

Note. This corollary restates Theorem 4.1 in degree-of-freedom language for direct citation. The saturation result and the no-new-degrees-of-freedom result are the same claim expressed differently.

Theorem 4.3 (Correlation Closure)

All admissible correlations between observables in T $\in \mathfrak{F}$ are functions over commitment outcomes and their distributions. No correlation observable introduces independent observable degrees of freedom beyond those in Obs(fold).

Proof. A correlation is an observable quantity computed from joint or conditional distributions over multiple outcomes. For a correlation to be admissible, each contributing outcome must be accessible to a finite observer through a repeatable protocol producing stable facts (Lemma 4.0a). Each such outcome is a commitment event. The correlation is therefore a function computed over a collection of commitment outcomes — a functional of the joint commitment outcome distribution.

For a correlation to introduce a new independent observable degree of freedom, it would need to vary in a way that corresponds to no variation in commitment outcome distributions: a correlation value that changes while the full joint distribution of commitment outcomes remains identical. Any such correlation would be undetectable by any finite observer — it produces no difference in any commitment outcome distribution and is excluded by condition 6 (admissibility of description) and the observational equivalence convention (SUP Section 2.4).

Therefore: there is no correlation observable whose variation does not correspond to a variation in commitment outcome distributions. All correlations are functions over Obs(fold). No correlation introduces independent observable content beyond the fold. □

Why this matters. Correlation Closure kills three specific loopholes: (1) hidden-variable completions in which correlations encode information not present in individual outcome distributions; (2) claims that quantum correlations (entanglement, Bell-type correlations) encode structure beyond commitment outcomes; and (3) statistical reinterpretations in which a new theory agrees on all individual outcomes but differs on joint distributions. In each case, the argument is the same: any variation in correlations must appear as a variation in commitment outcome distributions, which are already in Obs(fold).

Clarification for Bell-type correlations. In Bell setups, the structural content of entanglement is supposed to reside in the relationships between joint distributions at different measurement settings — different choices of measurement bases produce different joint distributions, and the non-classical correlations concern how those distributions relate across settings. Theorem 4.3 is consistent with this: the "joint distribution of commitment outcomes" ranges over all admissible measurement choices, not a single fixed setup. Each choice of measurement bases produces its own joint commitment outcome distribution, all of which lie within Obs(fold). The correlations across different settings — including Bell-inequality-violating correlations — are functions of these per-setting joint distributions and are therefore already in Obs(fold). No measurement-setting-dependence escapes the theorem; all settings contribute commitment outcomes of fold type.

Remark 4.3 (Why this is stronger than the No-Extension argument)

The previous No-Extension Theorem (CP Theorem 5.1) shows: extra structure is either redundant or creates a distinguishable competitor. The Saturation Theorem shows: the fold already occupies all observable space, so extra structure has nowhere to go that is not already occupied. The No-Extension argument operates by exclusion — showing what extra structure cannot do. The Saturation Theorem operates by characterisation — showing that Obs(fold) already covers the entire admissible observable space. Both arguments are valid; the Saturation Theorem is the more fundamental one.

Remark 4.4 (The observable degree-of-freedom count)

Theorem 3.1 identifies the observable degrees of freedom for a single commitment event: the fold generates a four-state base (the four (σ, ω) combinations) and a $\mathbb{C}^4$ state space with U(4) evolution. The observable algebra for a single commitment event is $M_4(\mathbb{C})$ — the full algebra of

4×4 complex matrices — with observables being the Hermitian elements of $M_4(\mathbb{C})$. The U(4) evolution acts on $\mathbb{C}^4$ by conjugation on states, and the observable predictions are Born-rule probabilities of (σ, ω) outcomes under U(4)-evolved pre-commitment states.

Important scope clarification: single-event vs full-theory algebra. The identification Obs(fold) = $M_4(\mathbb{C})$ is the algebra of a *single* commitment event. The observable algebra of a complete physical theory — which involves sequences of many commitment events over time and across spatial regions — is substantially larger. For n independent commitment events, the natural algebra is $M_4(\mathbb{C})^{\otimes n}$ with $U(4)^{\otimes n}$ evolution, subject to entanglement constraints arising from the fold's commitment structure. As $n \to \infty$, this generates an inductive limit structure. The saturation result says that *every individual observable in the full theory* reduces to fold-type commitment outcomes, each of which contributes an $M_4(\mathbb{C})$ factor. The full theory's observable algebra is the tensor tower over these factors, not $M_4(\mathbb{C})$ itself. "Obs(fold) = $M_4(\mathbb{C})$" should be read as: $M_4(\mathbb{C})$ is the generating algebra whose tensor extension over all commitment events gives the complete observable algebra of the theory. This is consistent with all saturation arguments — each observable reduces to fold-type factors — but the relationship between the single-event generator and the full-theory algebra requires explicit treatment, identified as Open Task 3 below.

A precision note: the correct framing is *covariance* under U(4) evolution, not *invariance*. Observables in Obs(fold) are not required to be U(4)-invariant — they are (σ, ω)-measurable Hermitian operators on $\mathbb{C}^4$ whose statistical predictions transform covariantly as states evolve under U(4). Requiring U(4)-invariance would yield only the scalars (by Schur's lemma, since U(4) acts irreducibly on $\mathbb{C}^4$), which is far too small. The correct mathematical object is $M_4(\mathbb{C})$ acting on $\mathbb{C}^4$, with U(4) specifying the dynamics of states rather than a symmetry of observables.

Saturation says the observable algebra of any admissible theory is generated by $M_4(\mathbb{C})$ — with no additional generators outside the fold structure. The saturation theorem applies at the level of generators; the full observable algebra is the closure under composition of these generators — which for n independent commitment events is $M_4(\mathbb{C})^{\otimes n}$, as identified in the scope clarification above. An independent algebraic proof that $M_4(\mathbb{C})$ is the maximal single-event admissible observable algebra under $\mathfrak{F}$ remains an open task (Open Task 1 below).

---

5. Killing the Remaining Escape Routes

The two-branch argument of the No-Extension Theorem handles generic additional structure. Three specific structural alternatives deserve explicit treatment, because they might appear to evade the generic argument.

5.1 Hidden Sectors

The objection. A theory might contain fields, particles, or degrees of freedom that do not directly participate in commitment events but influence the universe in ways not captured by ρ(x,t). Could such a hidden sector contain observable content not in Obs(fold)?

The response. For a hidden sector to be observable, its effects must be determinable by a finite observer through a repeatable protocol. That protocol terminates in facts — commitment events producing stable records. There are exactly two cases.

If the hidden sector produces no facts of its own and does not modify the facts produced by the fold, it is invisible to any finite observer by condition 6: it has no admissible observable content.

If the hidden sector does modify observable distributions — changes commitment rates, alters correlation structures, shifts probability distributions over (σ, ω) outcomes — then it is part of the effective commitment structure of T. Any structure that modifies observable distributions is, by Lemma 4.0, observable only through its effect on commitment outcomes. Its contribution is therefore already encoded in the commitment outcome distributions and hence in Obs(fold). This kills the "influence without direct observability" loophole: hidden variables that influence probabilities without being directly observable are precisely the structures whose effects appear in commitment outcome distributions. If they affect those distributions, they are part of Obs(fold). If they do not affect those distributions, they are invisible.

A hidden sector that produces no facts is physically inert. A hidden sector that modifies observable distributions is already encoded in Obs(fold). There is no third option — including no option for hidden variables that influence statistical distributions while remaining outside Obs(fold).

5.2 Emergent Overlays

The objection. Collective variables, order parameters, or emergent fields could be defined over the fold structure and behave as independent observables not reducible to individual commitment events. Could emergent overlays extend Obs(fold)?

The response. Emergent observables are functions of underlying facts (the commitment events from which they are defined). By Lemma 4.0, any function of commitment outcomes is already in the observable algebra.

The critical point is this: no emergent observable can introduce new independent observable degrees of freedom unless it introduces new distinguishable fact outcomes. Any emergent quantity that makes a new observable distinction must produce a difference in some commitment outcome distribution — otherwise it cannot be detected by any finite observer (condition 6, admissibility of description). If it produces a difference in commitment outcome distributions, that difference is a commitment-level fact and is already in Obs(fold). If it produces no such difference, it introduces no new observable content.

This ties the emergent overlay argument directly to the distinguishability condition (SUP Theorem 4.1: $\delta_{min} > 0$) and to condition 6: any new observable degree of freedom must produce a distinction separable by at least δ_{min} at the level of committed outcomes. Emergent quantities that fail to do this are not new observable degrees of freedom — they are redescriptions of existing ones.

Emergent overlays therefore cannot extend Obs(fold). They can only rearrange existing content into coarser or more convenient descriptions. The fold is the microscopic fact structure; all emergent descriptions are functional re-expressions of it.

5.3 Dual Descriptions

The objection. Some theories admit dual descriptions — two formally distinct presentations that are equivalent in all predictions. Could a dual description of the fold generate observables that the original fold formulation does not make apparent?

The response. If two descriptions are observationally equivalent, they are physically identical by the observational equivalence convention (SUP Section 2.4). A dual description of the fold that generates the same admissible fact distributions is a different formulation of the same theory — it is VERSF presented differently, and its observable algebra is Obs(fold) by definition.

If a putative dual description generates different admissible fact distributions, it is not a dual description of the fold but a different theory — an empirical competitor to VERSF, distinguishable by experiment. It is not a "different formulation of VERSF" but a new theory that must be tested against VERSF.

Dual descriptions either preserve the observable algebra (in which case they are VERSF formulated differently) or change it (in which case they are not dual descriptions but distinct theories). Neither case extends Obs(fold).

5.4 Contextuality and the Kochen-Specker Theorem

The objection. A quantum foundations reader familiar with the Kochen-Specker theorem will ask: quantum observables do not admit simultaneous classical value assignments — their values are context-dependent (they depend on which other commuting observables are jointly measured). Does this contextuality imply that observables exist "beyond" the fold's commitment structure, in the sense that their values are not determined by any single commitment outcome?

The response. Contextuality is a feature of Obs(fold), not a challenge to it. The Kochen-Specker theorem establishes that no non-contextual hidden-variable assignment to all Hermitian operators on $\mathbb{C}^4$ simultaneously exists. This is fully consistent with the saturation result. The fold does not assign classical values to all observables simultaneously — it assigns values via the Born rule for specific measurement bases, which are themselves choices of admissible protocol. Each choice of measurement context (each choice of (σ, ω) measurement basis) produces its own commitment outcome distribution, all of which lie within Obs(fold). The context-dependence of values is the statement that different measurement choices produce different joint commitment outcome distributions — and all of those distributions are in Obs(fold).

More precisely: Kochen-Specker establishes that the function from observables to values cannot be made basis-independent. The saturation result does not require it to be — Obs(fold) carries the full contextual structure of $M_4(\mathbb{C})$, including the non-commutativity of observables and the measurement-context dependence of outcomes. Contextuality is not evidence of structure beyond

Obs(fold); it is a property of Obs(fold)'s algebraic structure. The fold generates contextuality; it does not require something outside the fold to explain it.

5.5 Worked Example: Extra Scalar Field

To make the saturation argument concrete, consider a theory T_φ that extends VERSF by adding an extra real scalar field $\varphi(x,t)$ with its own dynamics — say, a Klein-Gordon field with mass m — defined on the same substrate as $\rho(x,t)$.

Is φ observable? For φ to be an admissible observable of T_φ, finite observers must be able to determine its value through repeatable protocols producing stable facts. Any such protocol terminates in a fact — a commitment event producing a stable record. The question is whether φ produces commitment events of its own or modifies the commitment events of the fold.

Case A: φ produces no commitment events of its own and does not modify fold commitment events. Then φ is invisible to all finite observers — no protocol can access its value without producing a fact, and by assumption φ produces no facts. By condition 6, φ has no physical content. T_φ is observationally identical to VERSF. $\mathrm{Obs}(T_\varphi) = \mathrm{Obs}(\mathrm{fold})$.

Case B: φ modifies fold commitment events — for instance, by acting as a background field that shifts the probability distributions over (σ, ω) outcomes (coupling to the pre-commitment $\mathbb{C}^4$ sector, modifying the Born-rule probabilities). Then φ's physical content is entirely encoded in the modified commitment outcome distributions. Those distributions are elements of Obs(fold) — they are born-rule functionals over fold-type commitment events. φ contributes no new observable degrees of freedom beyond what the fold already generates; it only modifies the statistical structure within Obs(fold). T_φ is again observationally reducible to VERSF.

Case C: φ generates its own independent commitment events — for instance, through φ-field fluctuations that produce irreversible stable records in some detector. By Case B1 of Theorem 4.1, those commitment events have fold structure (2D interface, (σ, ω) data, $\mathbb{C}^4$ pre-commitment sector). The observable content of φ-field detections is therefore in Obs(fold). T_φ is observationally equivalent to VERSF with a modified $\rho(x,t)$ encoding both fold and φ-field commitment events.

In every case, the extra scalar field either contributes nothing new (Case A) or contributes observable content already within Obs(fold) (Cases B and C). The saturation argument applies directly: $T_\varphi \in \mathfrak{F}$, T_φ contains the fold, therefore $\mathrm{Obs}(T_\varphi) = \mathrm{Obs}(\mathrm{fold})$. The extra scalar field does not extend the observable algebra. In all cases, φ contributes no independent observable degree of freedom beyond the fold.

---

6. The Observational Reduction Theorem

The Saturation Theorem and the escape-route eliminations jointly entail the main result.

Theorem 6.1 (Observational Reduction to VERSF)

Let T ∈ 𝔉. Then T is observationally equivalent to a VERSF-realisation: T's observable algebra is Obs(fold), and T's admissible fact distributions are those of VERSF.

Proof

By SUP Theorem 11.1, T contains the fold. By Theorem 4.1, Obs(T) = Obs(fold). All admissible observables of T are determined by the fold. The admissible fact distributions of T are therefore the admissible fact distributions of the minimal fold-containing theory, which is VERSF. By the observational equivalence convention (SUP Section 2.4), T is observationally equivalent to VERSF. □

Corollary 6.2 (Physics is observationally equivalent to VERSF)

Any theory T ∈ 𝔉 — that is, any physical theory in the operational sense of SUP Definition 3.0 — is observationally equivalent to VERSF. Physics, at the level of observable content, is observationally equivalent to VERSF.

---

7. What This Adds to the Uniqueness Programme

The programme now has three papers with distinct logical contributions:

Paper	Claim	Logical character
SUP	Any T ∈ 𝔉 must contain the fold	Necessity
CP	Any fold-containing T ∈ 𝔉 is VERSF	Equivalence by case exhaustion
This paper	The fold saturates Obs(T) for any T ∈ 𝔉	Equivalence by algebraic characterisation

The third paper is not merely redundant with the second. It answers a question the second cannot: *why* fold-containment entails observational equivalence to VERSF. The answer is that the fold occupies all available observable space — there is nowhere else for admissible observables to come from. CP established the result via five-stage case analysis showing that no specific source of underdetermination escapes the fold. The saturation argument of this paper shows why that analysis succeeds: Obs(fold) is the canonical observable algebra, generated by the interface and pre-commitment components, and every admissible observable lies within it. The case-exhaustion argument and the characterisation argument are equivalent in conclusion; the characterisation argument is the more illuminating one because it identifies the mathematical object that makes saturation inevitable rather than eliminating alternatives one by one.

The combined result across all three papers is:

```
Physics      →    T ∈ 𝔉      →    T ⊇ fold      →    Obs(T) = Obs(fold)
→    T ≃ VERSF
```

```
(Def. 3.0)       (necessity)       (SUP)              (saturation)
(equivalence)
```

Every arrow is a proved theorem. The final arrow is definitional (observational equivalence) once the penultimate one is established. The programme is complete.

Connection to quantum reconstruction. Hardy (2001) reconstructs quantum mechanics from five operational axioms about probabilities over outcomes; Chiribella, D'Ariano, and Perinotti (2011) give a similar reconstruction from information-theoretic axioms. The saturation result is structurally analogous: it reconstructs the fold's observable algebra ($M_4(\mathbb{C})$ as single-event generator) from operational conditions on fact formation. The key difference is that the VERSF approach derives the observable structure from the topology and geometry of the commitment interface rather than from axiomatic constraints on probability distributions alone — the fold is the geometric object from which the algebra emerges. This situates the saturation result within the broader quantum reconstruction programme while marking its distinctive route.

---

8. Falsifiability

The Saturation Theorem fails under the following condition:

F-S1 — An admissible observable not in Obs(fold) is identified. If a quantity $\mathcal{O}$ is found that is (a) accessible to a finite observer through a repeatable protocol producing stable facts, (b) not determinable by the fold structure or its commitment events, and (c) nonetheless a stable observable of some $T \in \mathfrak{F}$, then Theorem 4.1 fails. Such an observable would require either a new type of commitment event not of fold structure (refuting SUP Theorem 8.3 or Lemma 9.3) or a physical observable not grounded in commitment events at all (refuting CP Lemma 4.0 and SUP's account of how facts arise).

F-S2 — A hidden sector with observable consequences is identified. A field or degree of freedom that (a) does not participate in commitment events, (b) is not determined by $\rho(x,t)$ or any functional of it, and (c) nonetheless produces predictions distinguishable from VERSF would directly refute Section 5.1's elimination of hidden sectors and, by implication, the Saturation Theorem.

F-S3 — A dual description generating new observables is identified. A reformulation of the fold that is formally distinct from the standard fold presentation but generates admissible fact distributions not in Obs(fold) would refute Section 5.3's treatment of dual descriptions and require either revising the observational equivalence convention or accepting that the fold admits multiple distinct observable algebras.

None of these conditions is currently supported by theoretical or experimental evidence.

---

9. Conclusion

The VERSF fold is not just necessary for physics. It is sufficient to determine all observable physics.

Once a theory contains the fold, no additional observable degree of freedom is available. The fold defines a complete canonical observable algebra Obs(fold): the topology ($\beta_1 \geq 1$) enables commitment events to exist at all; the interface (σ, ω) fixes the four-state observable base; and the $\mathbb{C}^4$ pre-commitment sector with U(4) evolution fixes the statistical structure over that base. Every admissible observable in any $T \in \mathfrak{F}$ lies within this algebra. No additional structure can introduce new observable content: it either produces fold-type commitment events (already in Obs(fold)) or produces no facts at all (redundant). Hidden sectors, emergent overlays, and dual descriptions all reduce to the same conclusion.

The result is not that VERSF is the only possible formulation of physics. It is that any theory capable of producing observable phenomena must realise the same observable structure. The fold does not constrain physics — it exhausts it. There are no additional observable degrees of freedom beyond those generated by commitment events. Every admissible correlation, every conserved quantity, every emergent collective behaviour, every hidden sector — all either reduce to fold-generated commitment outcomes or produce no admissible distinctions at all. Therefore all admissible physical theories are observationally equivalent to VERSF.

Once facts exist, physics is already VERSF — whether you realise it or not.

---

10. Open Tasks

The following items are identified as the principal strengthening tasks for the programme, in priority order. Each is a genuine gap rather than a polish item.

Open Task 1 — Algebraic maximality of $M_4(\mathbb{C})$ (independent proof). Proposition 3.3 derives maximality from distinguishability constraints, which is ultimately the same argument as Theorem 4.1. An independent proof that $M_4(\mathbb{C})$ is the maximal single-event admissible observable algebra — on purely algebraic-closure grounds, without invoking the necessity chain — would overdetermine the saturation result. The natural tool is Jordan-algebra classification or the Koecher-Vinberg theorem on formally real Jordan algebras: the admissible algebras generated by a four-state binary outcome structure with U(4) covariance form a classified family, and $M_4(\mathbb{C})$ should be maximal within it. This is the biggest unrealised strengthening available.

Open Task 2 — Complete mathematical specification of Obs(fold). Remark 4.4 identifies $M_4(\mathbb{C})$ as the single-event generating algebra but does not fully specify: state space (density matrices on $\mathbb{C}^4$, or pure states only); dynamics (U(4) or U(4) modulo global phase); measurement structure (projective measurements only, or general POVMs); and composition rules for multiple commitment events. A one-page definition box pinning all of this down — ideally as a formal appendix — would make Obs(fold) citable at a technical level.

Open Task 3 — Tensor structure for composite systems. The observable algebra of a complete physical theory over n commitment events is $M_4(\mathbb{C})^{\otimes n}$ with $U(4)^{\otimes n}$ evolution, subject to entanglement constraints. As $n \to \infty$ this generates an inductive limit structure. Whether the fold structure naturally produces $M_4(\mathbb{C})^{\otimes n}$, a Fock-like construction, or something more subtle is not addressed here and is probably a separate paper's worth of work. Physical observables over spatially separated systems require this structure to be determined.

Open Task 4 — Formal verification of the five-mode exhaustiveness claim. The five sub-lemmas of Lemma 4.0 (4.0-M through 4.0-D) have been formalised with explicit premises and derivations. The residual verification task is the exhaustiveness claim: the conclusion of Lemma 4.0 argues that any admissible protocol falls into one of the five modes or fails to be admissible. This argument is given in prose in the Conclusion of Lemma 4.0, but a fully formal proof would require showing either that the five modes partition the space of admissible protocols, or that any protocol outside the five reduces to one of them. The pre-submission audit should also verify that each sub-lemma's cited SUP results deliver what the premises claim — particularly Sub-lemma 4.0-D's use of SUP Theorem 3.5 for time-ordering (see scope note in that sub-lemma).

Open Task 5 — Programme-level audit of Case B1 necessity chain. Case B1 cites five SUP results (Theorems 5.2, 6.2–6.3, 8.3, Lemma 9.3, Theorem 10.3) as applying to every admissible commitment event in every $T \in \mathfrak{F}$. The audit required is: does each result apply in the mode needed here — to every individual commitment event — rather than to the theory-level structure? For example, SUP Theorem 5.2 probably establishes that any theory supporting irreversibility requires $\beta_1 \geq 1$; this paper needs any individual commitment event to require $\beta_1 \geq 1$ of its local substrate. Verifying each is used in the correct mode is mechanical but essential before submission.

Open Task 6 — Modal Noether response. The current response to the Noether-structural objection handles the statistical-pattern content of conservation laws. A realist will push further: conservation laws have modal content — they say outcomes *must* take certain values given certain symmetries, not merely that they *do*. That modal content is distinct from the statistical pattern. The response — that within $\mathfrak{F}$, modal content reduces to counterfactual correlations accessible via intervention protocols — is real but requires a fuller treatment. This belongs as an extension of CP Remark 3.3 rather than this paper, but should be written somewhere in the programme.

Dependencies

Used here	Source	Content
SUP Definition 3.0	SUP Section 3	Operational definition of physics
SUP Definition 2.4	SUP Section 2	Committed record density $\rho(x,t)$

Used here	Source	Content
SUP Section 2.4	SUP Section 2	Observational equivalence convention
SUP Theorem 5.2, Lemma 5.3	SUP Section 5	$\beta_1 \geq 1$ necessity
SUP Theorems 6.2–6.3	SUP Section 6	Interface as locus of commitment
SUP Theorem 8.3	SUP Section 8	2D minimal interface
SUP Lemma 9.3, Corollary 9.4	SUP Section 9	(σ, ω) binary data, four-state base
SUP Theorem 10.3, Corollary 10.4	SUP Section 10	$\mathbb{C}^4$ with U(4) derived
SUP Theorem 11.1	SUP Section 11	Fold is necessary
CP Lemma 4.0	CP Section 4	Observables reduce to facts
CP Condition 6	CP Section 3	Admissibility of description
CP Theorem 5.1	CP Section 5	No-Extension Theorem
CP Theorem 6.1	CP Section 6	Closure Theorem

References

VERSF Programme (primary dependencies)

Taylor, K. (2024–2025). *Structural Uniqueness of Physical Law from Fact Formation Constraints* [SUP]. AIDA Institute. versf-eos.com

Taylor, K. (2024–2025). *Any Physical Theory Is VERSF: Completion of the VERSF Uniqueness Programme* [CP]. AIDA Institute. versf-eos.com

Taylor, K. (2024–2025). *Exclusion of Independent Distinguishability Sources: A No-Go Theorem for Multi-Field Ontologies in the VERSF Framework*. AIDA Institute. versf-eos.com

Quantum reconstruction and foundations

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

Chiribella, G., D'Ariano, G. M., & Perinotti, P. (2011). Informational derivation of quantum theory. *Physical Review A*, 84(1), 012311.

Kochen, S., & Specker, E. P. (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17, 59–87.

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]

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

Thermodynamics of computation (condition 3 substrate argument)

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.

Bennett, C. H. (1982). The thermodynamics of computation — a review. *International Journal of Theoretical Physics*, 21(12), 905–940.

Quaternionic quantum mechanics (ℍ elimination)

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 (competitor framework)

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 D. Oriti (Ed.), *Approaches to Quantum Gravity*. Cambridge University Press.

Philosophy of physics (operationalism and theory identity)

Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth*. Routledge.

van Fraassen, B. C. (1980). *The Scientific Image*. Oxford University Press.

Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica*, 43(1–2), 99–124.

VERSF Research Programme
Keith Taylor
versf-eos.com