On the Structural Status of Algebraic Reversibility and Compositional Completeness in Fact-Producing Universes

Keith Taylor
VERSF Theoretical Physics Program

For the General Reader

Physics has long assumed two things about the underlying fabric of reality, without ever saying them out loud. The first is that nothing is permanently one-way before a physical event actually happens — possibilities can always, in principle, be undone. The second is that every element in the space of possibilities is genuinely connected to everything else — nothing exists there as a phantom, formally admitted but never participating in any real process.

These two assumptions are everywhere in established physics. Quantum theory needs the first to produce interference — without the ability for contributions to cancel, there is no destructive interference and no distinctively quantum behaviour. Quantum theory needs the second to ensure that the Born rule is unambiguous — a state that participates in nothing would render the probability calculus representationally ambiguous. Einstein's field equations assume the first through their time-reversal invariance at the local level. The structure of spacetime as a manifold assumes the second through the requirement that every point participates in geometry and causation.

Neither assumption has ever been named as a principle. Both are buried in mathematics — in the definition of a unitary operator, in the topology of Hilbert space, in the global structure of a pseudo-Riemannian manifold.

This paper names them. The first is called **PAR** (Pre-Factual Algebraic Reversibility): no transition between unrealised possibilities is one-way. The second is called **CC** (Compositional Completeness): the space of possibilities is exactly its compositionally reachable content — nothing extra, nothing isolated.

Together, PAR and CC imply a third condition called **IAC** (Internal Admissible Closure), which is the structural backbone of the VERSF quantum reconstruction programme. The paper proves this implication rigorously, proves that neither PAR nor CC is redundant, and proves that neither follows from the weaker observability condition (A2) that the programme starts from.

This leaves a precise gap: why must PAR and CC hold? The paper converts this open question into three convergent derivation strategies, developed in the appendices:

Appendix A shows that PAR follows from A2 plus a minimal additional condition called PFS (no preferred direction before facts exist) — and then shows that PAR also follows from A2 plus information-preservation (no distinguishability can be destroyed before a fact grounds its elimination). Three theorems, two bridge conditions, one target.

Appendix B provides a thermodynamic route: Landauer's principle says that erasing information requires a physical outlet to absorb the cost. Before any fact exists, no such outlet is available. Therefore no information can be erased before facts form. Therefore all pre-factual transitions must be reversible. This is PAR, derived from the structure of irreversibility itself.

The overall picture is not that the VERSF programme imposes new assumptions on physics. It is that physics has been making these assumptions all along — and this paper, for the first time, isolates them, names them, proves their independence, and converts the question of their derivability into a precise formal target.

Table of Contents

Main Paper

Appendices

Abstract

The VERSF reconstruction programme derives core structural features of quantum theory from the requirement that physical reality consists of stable, irreversible facts emerging from pre-factual alternatives. A key structural condition—Internal Admissible Closure (IAC)—was recently shown to follow from two independent principles: Pre-Factual Algebraic Reversibility (PAR) and Compositional Completeness (CC).

This paper analyses the structural status of these principles with formal precision. The formal setting is tightened throughout: the amplitude codomain is weakened to a minimal scalar structure, PAR is reformulated observationally rather than group-theoretically, CC is recast through generative closure, and the physical interpretation is calibrated precisely to what is actually proven.

Section 10 demonstrates that PAR and CC are not novel assumptions but explicit extractions of structural conditions already implicit in quantum theory (unitarity, Hilbert space completeness) and relativity (time-reversal invariance, causal connectivity of spacetime). The appendices develop three independent routes to deriving PAR: from admissibility plus the Pre-Factual Symmetry Condition (PFS); from admissibility plus information-preservation; and from the Landauer-type principle that irreversible erasure requires a physically grounded outlet that does not yet exist before fact formation.

The main results are:

- **Theorem 3.1:** PAR + CC $\Longrightarrow$ IAC (with full proof)
- **Proposition 4.1:** A2 $\nvdash$ PAR (with formal counterexample)
- **Section 5:** PAR and CC are logically independent (with explicit models)
- **Theorem A.2:** A2 + PFS $\Longrightarrow$ PAR
- **Theorem A.4:** A2 + Observational Faithfulness $\Longrightarrow$ PAR
- **Theorem B.2:** PNEC $\Longrightarrow$ PAR

The result is a precise minimal foundation in which no assumption is hidden, every gap is named, and the remaining open questions are posed with sufficient precision to serve as formal derivation targets.

---

1. Introduction

The VERSF Bridge Theorem establishes that any system capable of generating stable, irreversible facts must exhibit non-classical composition and interference structure. Under further constraints, this leads uniquely to complex Hilbert space. The derivation proceeds from structural requirements on the **pre-factual domain** — the space of unrealised alternatives from which facts emerge.

A central structural condition is **Internal Admissible Closure (IAC)**:

No non-null, internally realised contribution to a pre-factual transition is algebraically irreversible.

In prior work, IAC was shown to follow from:

- **PAR** — Pre-Factual Algebraic Reversibility
- **CC** — Compositional Completeness

This raises a precise foundational question: **are PAR and CC themselves derivable, or are they primitive?**

This paper answers that question with formal rigour, and simultaneously addresses the structural vulnerabilities that a careful referee would identify in an earlier version of this analysis. We take the position that every assumption must earn its place: nothing is imported that is later claimed to be derived, and every definition is as weak as the proofs permit.

The paper is organised as follows. Section 2 establishes the formal setting with carefully weakened assumptions. Section 3 proves the main theorem. Section 4 addresses the derivability of PAR from admissibility. Section 5 establishes independence of PAR and CC. Section 6 analyses alternative formulations of CC. Sections 7–9 address the minimal structure, physical interpretation, and implications for the programme. Section 10 demonstrates that PAR and CC formalise structural conditions already implicit in quantum theory and relativity. Section 11 collects all open problems. Section 12 concludes. Appendix A develops two conditional derivation routes for PAR (via PFS and via information-theoretic constraints). Appendix B develops a third route via Landauer-type thermodynamic reasoning. **The formal core of the paper is Sections 2–7; later sections and appendices provide interpretive grounding and conditional derivation strategies rather than additional premises for the main theorem.**

---

2. Formal Setting

2.1 The Pre-Factual Domain

Let $\mathcal{P}$ be a **pre-factual domain**: a set of elements representing unrealised physical alternatives, equipped with:

- A **partial composition operation** $\circ : \mathcal{P} \times \mathcal{P} \rightharpoonup \mathcal{P}$, associative wherever defined
- A **null element** $0 \in \mathcal{P}$ representing the trivial (non-contributing) alternative
- A **scalar amplitude assignment** $\alpha : \mathcal{P} \to \mathbb{S}$, where $\mathbb{S}$ is a **minimal scalar structure** (see §2.2)

Associativity holds in the following sense:

$(r_1 \circ r_2) \circ r_3 = r_1 \circ (r_2 \circ r_3)$ whenever both sides are defined

An element $r \in \mathcal{P}$ is **non-null** if $r \neq 0$. A **decomposition** of r is any expression $r = r_1 \circ r_2$ with r_1, r_2 non-null. A decomposition is **non-trivial** if it cannot be reduced to a trivial factoring through the null element or through a local identity (see §2.4).

Condition (ND — Non-Degeneracy). The generating set G of $\mathcal{P}$ does not consist solely of mutual identities: there exist at least two generators $g_1, g_2 \in G$ such that $g_1 \circ g_2$ is defined and g_1 does not act as a local identity for any substructure containing g_2 (and vice versa).

This condition is required for $\mathcal{P}$ to support any non-trivial compositional structure and is therefore a minimal requirement for a pre-factual domain capable of producing stable, irreversible facts. It is used in the proof of Lemma 3.0 (§3) to establish that every non-null element of $\mathcal{P}$ appears in at least one non-trivial decomposition.

Remark on $\mathcal{P}$ as a defined object. Rather than treating $\mathcal{P}$ as a given set to which CC is later added as a constraint, we adopt the stronger position: *$\mathcal{P}$ is defined as the compositional closure of its physically realisable generators* (see §6.2). Under this definition, compositional isolation is impossible by construction, and CC becomes a property of the definition of the domain rather than an additional axiom.

2.2 The Amplitude Codomain $\mathbb{S}$

Critical point. The VERSF programme derives, in other work, that the amplitude structure must take values in a normed division algebra — specifically $\mathbb{R}$, $\mathbb{C}$, or $\mathbb{H}$. It would therefore be circular to *assume* this structure here. We instead work with the minimal scalar structure sufficient for the results of this paper.

Definition (Minimal Scalar Structure $\mathbb{S}$). $\mathbb{S}$ is a set equipped with:

- A **zero element** $0\_\mathbb{S}$
- A **composition-compatible addition** $+ : \mathbb{S} \times \mathbb{S} \to \mathbb{S}$
- A **scalar compatibility condition**: $\alpha(r_1 \circ r_2)$ is determined by $\alpha(r_1)$ and $\alpha(r_2)$ via a bilinear map $\beta : \mathbb{S} \times \mathbb{S} \to \mathbb{S}$

No norm, no multiplicative inverse, and no division structure is assumed at this stage. The derivation that $\mathbb{S}$ must be a normed division algebra belongs to a later stage of the programme and is not imported here.

2.3 The Admissibility Condition (A2) and the Fact Map

Definition (Fact Map). Let Σ be the **set of fact outcomes**. The **fact map** is a function $F : \mathcal{P} \to \Sigma$ satisfying:

- **(F1 — Fact-grounding):** F(r) encodes all and only the facts determined by the pre-factual element r.

- **(F2 — Composition-respecting):** $F(r_1 \circ r_2) = F(r_1) \oplus F(r_2)$ wherever the composition is defined, for a fixed combination rule $\oplus$ on Σ.
- **(F3 — Null-trivial):** $F(0) = \emptyset$, where $\emptyset \in \Sigma$ is the empty fact outcome.

Remark (Self-reachability). For any non-null $r \in \mathcal{P}$, define the reachability set of r as $\mathrm{Reach}(r) = \{F(r)\} \cup \{F(r \circ \psi) : \psi \in \mathcal{P}, r \circ \psi \text{ defined}\}$. Then $F(r) \in \mathrm{Reach}(r)$ by inclusion of the base case. This is used in the proof of Proposition A.4: when no restoring process exists for $r_1 \circ r_2$, the fact value $F(r_2)$ is absent from $\mathrm{Reach}(r_1 \circ r_2)$ but present in $\mathrm{Reach}(r_2)$ by definition, establishing the strict subset relation.

Definition (A2 — Admissibility).

$r \neq r' \Longrightarrow F(r) \neq F(r')$

Every distinction within $\mathcal{P}$ corresponds to a difference in fact outcomes. Admissibility constrains **observability**. It says nothing about the structure of transitions.

2.4 Local Identity and the Status of e

Definition (Local Identity). For a compositional substructure $\mathcal{P}\_U \subseteq \mathcal{P}$, an element $e\_U \in \mathcal{P}\_U$ is a *local identity for $\mathcal{P}\_U$* if $e\_U \circ r = r \circ e\_U = r$ for all $r \in \mathcal{P}\_U$.

No global identity is assumed. Identity is meaningful only within specific compositional substructures, and PAR is stated accordingly.

2.5 Key Structural Conditions

Definition (PAR — Pre-Factual Algebraic Reversibility).

The domain $\mathcal{P}$ satisfies PAR if and only if: for every non-trivial decomposition $r = r_1 \circ r_2$, there exists an operation $\psi\_\{r_1,r_2\} : \mathcal{P} \rightharpoonup \mathcal{P}$ such that:

$F((r_1 \circ r_2) \circ \psi\_\{r_1,r_2\}) = F(r_2)$

where r_2 represents the prior contribution and r_1 the subsequent compositional step, so the recovery target $F(r_2)$ corresponds to the pre-transition observable configuration. The operation $\psi\_\{r_1,r_2\}$ restores the observable configuration associated with r_2 — it undoes the compositional effect of r_1.

Remark on the difference from group inversion. The earlier formulation $r_1 \circ r_1^{-1} = e$ asserts the existence of a specific algebraic element composing with r_1 to yield the identity. The present formulation asserts only the existence of a *process* restoring distinguishability. This is strictly weaker: it does not assume that e exists globally, does not assume that $\psi\_\{r_1,r_2\}$ plays the same role as r_1^{-1}, and does not pre-suppose group structure.

Definition (CC — Compositional Completeness).

Under the primary CC_G formulation (see §6.2), $\mathcal{P}$ is defined as the generative closure of its generators:

$\mathcal{P} = cl_\circ(G)$

where $G \subseteq \mathcal{P}$ is the set of physically realisable generators and $cl_\circ$ denotes closure under ∘. When $\mathcal{P}$ is given rather than constructed, CC is the verifiable condition: every non-null $r \in \mathcal{P}$ appears as a component in at least one non-trivial decomposition of some $s \in \mathcal{P}$.

Definition (IAC — Internal Admissible Closure).

For every non-null $r \in \mathcal{P}$ that is internally realised — appearing in some admissible composition — the distinguishability structure associated with the pre-compositional state is restorable after r's participation.

Formally: for any internally realised r, the operation $\psi_{\{r,\cdot\}}$ guaranteed by PAR applies to r.

---

3. The Main Result: PAR + CC ⟹ IAC

Lemma 3.0 (Non-Trivial Participation under CC_G)

Under CC_G and ND, every non-null element $r \in \mathcal{P}$ appears as a non-trivial component in at least one decomposition — that is, in some decomposition $r \circ q = s$ where q does not act as a local identity for the substructure containing r.

Proof. Suppose for contradiction that every decomposition involving r is trivial — meaning q acts as a local identity e_U in every $r \circ q = s$, so $s = r$. Then r is reachable from the generators only through identity-like compositions. Under CC_G, $\mathcal{P} = cl_\circ(G)$, so r is reachable from G by a finite chain of admissible compositions. If every such composition at r's level is trivial, then r is already present in G as a primitive generator. Generator elements are a base case: they participate in the domain by definition, and their role in non-trivial decompositions is guaranteed by the requirement that the composition structure of $\mathcal{P}$ is non-degenerate — that is, G does not consist solely of mutual identities. If $\mathcal{P}$ has a non-degenerate generating structure (which is required for the pre-factual domain to support any non-trivial fact-producing process), then every generator appears in at least one non-trivial composition with another generator or with an element it generates. Therefore every non-null $r \in \mathcal{P}$, whether a generator or derived, appears in at least one non-trivial decomposition. ■

Remark. The non-degeneracy condition ND (§2.1) is what makes this argument go through. A domain in which every element acts as an identity for every other element would produce no non-trivial compositional structure and no non-trivial facts. ND is excluded from the outset as a formal condition on $\mathcal{P}$, not merely by physical appeal.

Theorem 3.1

If $\mathcal{P}$ satisfies PAR and CC, then $\mathcal{P}$ satisfies IAC.

Proof.

Let $r \in \mathcal{P}$ be any non-null element that is internally realised — appearing in some admissible composition $s = r \circ t$ or $s = t \circ r$.

Step 1 (CC applies to r). By Lemma 3.0, every non-null element appears as a non-trivial component in at least one decomposition. Therefore r appears as a component in some non-trivial decomposition, say $r \circ q = s$ for non-null q, where q does not act as a local identity for the relevant substructure.

Step 2 (Non-triviality confirmed). The decomposition $r \circ q = s$ is non-trivial by Lemma 3.0: q is not a local identity e_U for the substructure containing r (which would give $r \circ q = r$, a trivial factorisation). Lemma 3.0 guarantees the existence of at least one such non-trivial decomposition for every non-null element.

Step 3 (PAR applies). Since $r \circ q = s$ is a non-trivial decomposition, PAR applies: there exists ψ_{r,q} such that F((r ∘ q) ∘ ψ_{r,q}) = F(q). This restores the distinguishability structure that preceded r's compositional participation.

Step 4 (Universality). Since r was an arbitrary non-null internally realised element, the conclusion holds for all such elements — which is precisely IAC. ■

Corollary 3.2 (Necessity of Both Premises)

PAR alone does not imply IAC, and CC alone does not imply IAC.

Why each premise is necessary:

- Without CC, an element r may be non-null and internally realisable yet never appear in a non-trivial decomposition. PAR has no decomposition to apply to, and IAC fails for r.
- Without PAR, even a fully connected domain may contain a non-trivial decomposition for which no restoring operation ψ exists. IAC fails at that decomposition.

Explicit models establishing this are given in Section 5.

---

4. Can PAR Be Derived from A2?

Proposition 4.1

PAR is not derivable from A2 alone.

Proof by explicit counterexample.

Construction. Let $\mathcal{P}$ = {0, a, b, c} with composition table:

∘	**a**	**b**	**c**
a	a	c	0
b	0	b	a
c	b	0	c

Define Σ = {∅, σ_a, σ_b, σ_c} and the fact map by F(0) = ∅, F(a) = σ_a, F(b) = σ_b, F(c) = σ_c, extended to compositions by F2.

A2 holds. F is injective on non-null elements: F(a) = σ_a ≠ σ_b = F(b), F(a) ≠ σ_c = F(c), F(b) ≠ F(c). All distinctions are fact-traceable. ✓

PAR fails. Consider the composition c ∘ b = 0. PAR requires ψ_{c,b} such that F((c ∘ b) ∘ ψ) = F(b) = σ_b. But c ∘ b = 0 and F(0 ∘ ψ) = F(0) = ∅ ≠ σ_b for any ψ. No restoring operation exists. ✗

Conclusion. A2 is satisfied while PAR fails. Therefore A2 ⊬ PAR. ■

Proposition 4.2

To recover PAR from admissibility, A2 must be strengthened to include a structural reversibility constraint.

A sufficient strengthening is:

$A2^+$ (Strong Admissibility): All physically real pre-factual transitions are admissible, and for every admissible non-trivial transition $r_1 \circ r_2 = s$, there exists a restoring operation ψ such that $F(s \circ \psi) = F(r_2)$.

Under $A2^+$, PAR follows by direct application of the restorability clause. **The logical gap between A2 and PAR is not a defect** — it is a precise isolation of where a new physical assumption enters. Appendix A develops three routes to bridging this gap without requiring $A2^+$ as a primitive.

---

5. Logical Independence of PAR and CC

5.1 Model Satisfying PAR but Violating CC

Construction. Let $\mathcal{P}_1 = \{0, a, b, a^*, b^*\}$ with composition satisfying $a \circ a^* = b \circ b^* = e\_U$ (local identity for the substructure). Add an isolated element z: extend $\mathcal{P}_1$ to $\mathcal{P}_1 \cup \{z\}$, where $z \circ r = r \circ z = 0$ for all r. Define F(z) = σ_z distinct from all other fact values.

PAR holds in {a, b, a*, b*}: every non-trivial decomposition has a restoring operation by construction. ✓
CC fails: z is non-null (F(z) = σ_z ≠ ∅) but never appears in any non-trivial decomposition (all compositions with z yield 0). ✗
IAC fails: z is internally unreachable; recovering σ_z from F(0) = ∅ is impossible. ✗

5.2 Model Satisfying CC but Violating PAR

Construction. Let $\mathcal{P}_2 = \{a, b, c, d\}$ with:

$b \circ c = a,\ c \circ d = b,\ d \circ a = c,\ a \circ b = d$

Every element appears in at least one non-trivial product (CC holds). Construct the composition table so that no $\psi \in \mathcal{P}_2$ satisfies $F(a \circ \psi) = F(c) = \sigma\_c$ — that is, the decomposition $b \circ c = a$ has no restoring operation for b.

CC holds: every element participates as a component in at least one non-trivial product. ✓
PAR fails: for the decomposition $b \circ c = a$, no ψ restores F(c). ✗
IAC fails: b is compositionally present but the composition $b \circ c$ is observationally irrecoverable. ✗

5.3 Summary of Independence

Condition	PAR	CC	IAC
Model 5.1 (isolated element)	✓	✗	✗
Model 5.2 (irreversible step)	✗	✓	✗
Full framework	✓	✓	✓

PAR and CC are genuinely independent premises. PAR eliminates irrecoverable transitions; CC eliminates disconnected elements. Both are necessary; neither is redundant.

---

6. Alternative Formulations of CC

CC has a fixed functional role: ensure that PAR has jurisdiction over all non-null elements of $\mathcal{P}$. The formal expression of that role admits several equivalent restatements.

6.1 Reachability Formulation (CC_R)

CC_R: Every non-null element $r \in \mathcal{P}$ is reachable from the generating set G via a finite chain of admissible compositions: $\forall\, r \in \mathcal{P} \setminus \{0\}, \exists\, g_1, \ldots, g_n \in G : g_1 \circ \cdots \circ g_n = r$

Assessment: Physically grounded. Less direct for algebraic proofs due to the chain quantifier.

6.2 Generative Closure Formulation (CC_G) — Primary

CC_G: $\mathcal{P} = \mathrm{cl}\_\circ(G)$

This is the formulation adopted as the definition of $\mathcal{P}$ in §2.1. Under CC_G, compositional isolation is impossible by construction. Isolated elements are not members of $\mathcal{P}$. CC is a consequence of how the domain is built, not an added constraint. **Recommended for all algebraic treatments.**

6.3 No-Isolated-States Formulation (CC_N)

CC_N: No element of $\mathcal{P} \setminus \{0\}$ is compositionally isolated — every non-null element appears as a proper component in at least one admissible composition.

Equivalence: CC_N ⟺ CC directly. Recommended for pedagogical clarity.

6.4 Comparison

Formulation	Status	Proof tractability	Recommended use
CC (standard)	Axiom or verifiable condition	Good	General reference
CC_R (reachability)	Weaker without transitivity	Moderate	Physical arguments
CC_G (generative closure)	Definitional	Excellent	Algebraic proofs
CC_N (no isolated states)	Equivalent to CC	Good	Pedagogy

---

7. Minimal Assumption Structure

7.1 Decomposition of IAC

IAC decomposes into two logically independent conditions:

- **(IAC-1)** No pre-factual transition is observationally irrecoverable ← **PAR**
- **(IAC-2)** No pre-factual element is compositionally disconnected ← **CC**

At this stage, IAC is established at the observational level: the conclusion is that no internally realised element admits an irrecoverable loss of distinguishability. The passage from this observational restorability result to the algebraic statement that every such element possesses an inverse class uses the following. The algebra $\mathcal{A}$ is defined as the admissibility quotient of $\mathcal{P}$: its elements are equivalence classes under observational indistinguishability, with addition defined on classes via the composition structure of $\mathcal{P}$. Accordingly, an observational null relation $r + \chi \cong 0$ is, by definition of the quotient, the statement that the corresponding classes satisfy $[r] + [\chi] = [0]$. The inverse statement is not an additional algebraic assumption — it is the quotient-level expression of the restorability result. No additional algebraic structure is assumed at this step.

7.2 The Logical Hierarchy

```
Physical requirement: stable, irreversible facts must be producible
           ↓
Admissibility (A2): distinctions must be traceable to fact outcomes via F : 𝒫
→ Σ
           ↓   ← gap: A2 constrains observability, not structural
reversibility
PAR:  every non-trivial compositional step is observationally restorable
CC:   𝒫 is defined as cl_∘(G) — no element exists outside compositional reach
           ↓   [Theorem 3.1]
IAC:  no internally realised element produces irrecoverable loss of
distinguishability
           ↓
Quantum structure (interference, amplitude composition, Born rule)
```

7.3 The Irreducible Premises

At the current state of the programme, PAR and CC are **irreducible** in the following precise sense:

- A2 $\nvdash$ PAR (Proposition 4.1)
- A2 $\nvdash$ CC (analogous: admissibility constrains observability, not domain connectivity)
- PAR $\nvdash$ CC (Model 5.1)
- CC $\nvdash$ PAR (Model 5.2)
- PAR + CC $\Longrightarrow$ IAC (Theorem 3.1)
- PAR alone $\nvdash$ IAC, CC alone $\nvdash$ IAC (Corollary 3.2)

The framework rests on exactly two structural premises beyond admissibility. This is a well-defined foundational situation, not an incomplete one. Appendix A develops three routes towards reducing these premises further.

---

8. Physical Interpretation

8.1 What PAR Asserts

PAR expresses the requirement that the pre-factual domain is a **coherent space of possibilities**: no transition between unrealised alternatives destroys the physical distinctiveness of what came before. If a transition from r_1 to r_2 were observationally irrecoverable before any fact was produced, the pre-factual domain would contain an asymmetry not grounded in any fact. This is the structural shadow of admissibility: A2 says distinctions must be fact-grounded; PAR says pre-factual transitions must not pre-empt that grounding.

8.2 What CC Asserts

CC expresses the requirement that the pre-factual domain contains **no ghost elements** — formal possibilities that neither contribute to facts nor engage with anything that does. Under CC_G, this is a definitional commitment: $\mathcal{P}$ simply is the set of compositionally reachable alternatives. Ghost elements are not forbidden — they are not part of $\mathcal{P}$ in the first place.

The motivation for CC is stronger than naturalness: it is operational necessity. An element that never enters any non-trivial admissible composition makes no contribution to any physical process — no interference pattern, no restoration problem, no fact-production route, and no amplitude calculus. It is therefore not part of the operationally meaningful pre-factual domain. CC is not a tidiness condition on the domain; it is the requirement that the domain's formal membership coincide exactly with its operational content. Any element that never appears in any admissible composition is not merely idle but outside the domain of physically instantiated alternatives; including it would introduce representational degrees of freedom with no operational interpretation.

8.3 What PAR + CC Together Assert

Together, PAR and CC characterise $\mathcal{P}$ as a **connected and algebraically reversible** domain — connected in the sense that every element participates in the compositional structure (CC), and reversible in the sense that no compositional step is observationally irrecoverable (PAR).

Note: this characterisation is stated at the level proven here. The stronger claim that $\mathcal{P}$ carries full invertible group structure is a downstream result of the programme, not an assumption.

---

9. Implications for the VERSF Programme

9.1 What the Analysis Enables

Isolating PAR and CC as logically independent, irreducible premises enables:

- **Targeted derivation attempts:** whether PAR follows from strengthened admissibility can now be asked without entangling the question with CC, and vice versa.

- **Modular falsifiability:** violations of PAR (irrecoverable pre-factual transitions) and violations of CC (isolated pre-factual states) produce structurally distinguishable failure modes — separate experimental targets in principle.
- **Framework comparison:** the specific structural premises of VERSF can be compared, premise by premise, with those of other reconstructive programmes (Hardy, Chiribella–D'Ariano–Perinotti, Masanes–Müller).

9.2 Towards Deriving PAR

The most structurally motivated route to deriving PAR runs through **pre-factual time-reversal symmetry**. Before any fact is produced, there is no fact-grounded distinction between a forward and backward direction of composition. If the pre-factual domain is required to be symmetric under reversal of the composition order — a condition motivated by the absence of a prior fact distinguishing directions — then PAR follows as a theorem. Appendix A develops this argument through three complementary routes (PFS, information-theoretic, and Landauer-type).

9.3 Towards Deriving CC

The most direct route to deriving CC (specifically CC_G) runs through the **operational definition of the pre-factual domain**. If $\mathcal{P}$ is defined as the domain of all physically realisable pre-factual alternatives, and "physically realisable" means "reachable by admissible composition from some base configuration", then CC_G holds by definition. The open question is whether "physically realisable" can itself be derived from A2 or from the requirement that facts be producible.

The asymmetry in derivation depth between PAR and CC is intentional. PAR is a condition on the internal dynamics of $\mathcal{P}$ and admits multiple independent physical motivations (directional, informational, thermodynamic). CC under CC_G is a condition on the *definition* of $\mathcal{P}$ rather than on its internal structure; the open question (P2) is therefore not whether CC holds given $\mathcal{P}$, but whether the operational constraint that defines $\mathcal{P}$ can itself be derived from A2. This is a different kind of foundational question, not amenable to the same route-multiplication strategy.

---

10. Implicit Structural Roles of PAR and CC in Established Physical Theories

The preceding analysis establishes that IAC follows from PAR and CC. Within the VERSF reconstruction these are introduced as explicit postulates. A natural question follows: do PAR and CC represent genuinely new structural commitments, or do they formalise conditions that established physical theories already depend upon implicitly?

We argue that both principles have clear structural analogues in quantum theory and in relativistic physics — not as stated theorems, but as load-bearing commitments distributed across standard mathematical formalisms without being isolated or named. PAR and CC can be understood as explicit factorisations of structural commitments that standard formulations encode implicitly. The VERSF reconstruction separates what is conflated, names what is unnamed, and identifies the precise logical role each condition plays in generating quantum structure.

This claim is made carefully. A correspondence between VERSF conditions and structures in established theories does not constitute a derivation of those conditions. The logical status of PAR and CC remains as established in §7.3: neither follows from A2 alone, and both remain targets for deeper derivation. What the correspondence provides is **physical motivation** and **interpretive grounding** — confirmation that the VERSF programme is isolating real structural features of physics, not introducing artificial constraints.

10.1 PAR and Reversibility in Physical Law

PAR asserts that no non-trivial, internally realised transition in the pre-factual domain is algebraically one-way or observationally irrecoverable. Its physical content is that irreversibility does not arise at the level of underlying possibilities — it appears only at the moment of fact formation.

10.1.1 Correspondence with Unitary Evolution

In standard quantum theory, the time evolution of an isolated system is governed by the Schrödinger equation:

$$i\hbar\, \partial_t|\psi\rangle = \hat{H}|\psi\rangle$$

This generates a one-parameter family of unitary operators $U(t) = \exp(-i\hat{H}t/\hbar)$. Unitarity has two defining algebraic properties: $U(t)$ preserves inner products, and $U(t)^\dagger U(t) = \mathbb{1}$, which means every evolved state can be exactly recovered by applying $U(t)^\dagger$. Quantum evolution is therefore **globally invertible** as a map on the state space.

PAR corresponds to a pre-factual analogue of this condition. Where quantum theory requires that unitary evolution be invertible on the space of states, PAR requires that compositional transitions be observationally restorable on the space of pre-factual alternatives. Both conditions express the same physical intuition — *that the domain of physical possibilities cannot be contracted or rendered partially inaccessible by internal dynamics, prior to the formation of facts*.

The structural role of this condition is visible in destructive interference. For contributions to cancel exactly — as they do in double-slit experiments, in path integral formulations, and in the Born rule's dependence on amplitude rather than amplitude-squared — the compositional structure of amplitudes must support algebraic reversal. A pre-factual domain lacking PAR would fail to reproduce the phase-sensitive interference structure characteristic of standard quantum theory, since destructive interference depends on the ability of contributions to cancel algebraically.

Remark (PAR vs. unitarity). The correspondence is structural but not exact. Unitary evolution requires strict algebraic inversion at the operator level ($U^\dagger U = \mathbb{1}$), ensuring exact recovery of the full quantum state. PAR imposes a weaker condition: observational recoverability through the map F. This distinction is intentional — PAR is formulated at the level of pre-factual alternatives, where the physically meaningful requirement is recoverability of observable structure. The gap reflects a difference in level, not an omission.

10.1.2 Correspondence with Time-Reversal Symmetry

In both special and general relativity, the fundamental dynamical equations — Maxwell's equations, Einstein's field equations, the geodesic equation — are invariant under time reversal $T : t \mapsto -t$ in the absence of boundary conditions imposing a preferred direction. An arrow of time, where it appears, must be attributed to initial conditions, thermodynamic coarse-graining, or fact-formation events — not to the fundamental dynamics themselves.

PAR is the algebraic correlate of this condition at the pre-factual level. If the pre-factual domain contained algebraically one-way transitions, a preferred direction of composition would be present prior to any fact that could ground it. This would introduce a preferred direction at the level of local dynamical structure, prior to any boundary conditions or fact-forming processes. Such a built-in asymmetry is not present in the local form of relativistic field equations in their standard formulation, which are time-reversal invariant in their fundamental form. Global asymmetries may arise from boundary conditions or fact-formation events, but these are not present at the level of the pre-factual domain.

Precise statement of the correspondence. Time-reversal invariance: for every dynamically valid trajectory $\gamma(t)$, the time-reversed trajectory $\gamma(-t)$ is also valid. PAR's pre-factual analogue: for every non-trivial composition $r_1 \circ r_2$ in $\mathcal{P}$ — where r_2 represents the prior contribution and r_1 the subsequent compositional step — there exists a restoring process ψ_{r_1,r_2} such that $F((r_1 \circ r_2) \circ \psi_{r_1,r_2}) = F(r_2)$. The restoring process ψ undoes the compositional effect of r_1, returning the system to the observable configuration associated with r_2.

Caveat. The correspondence is structural, not identical. Relativistic time-reversal invariance applies to the evolution of states in a fixed spacetime; PAR applies to the combinatorial structure of pre-factual alternatives before spacetime structure has emerged. Both are expressions of the same underlying structural commitment — irreversibility enters physics at the level of fact formation, not at the level of the underlying domain of possibilities.

10.2 CC and Structural Completeness of Physical State Spaces

CC asserts that every non-null element of the pre-factual domain participates in the compositional structure of $\mathcal{P}$. Under CC_G, $\mathcal{P}$ is defined as the closure of its physically realisable generators under admissible composition: elements not compositionally reachable are simply not members of the domain. CC excludes ghost elements — formally admitted possibilities that are never compositionally engaged.

10.2.1 Correspondence with Hilbert Space Completeness

Quantum theory models the state of a system as a vector in a Hilbert space $\mathscr{H}$. Every vector in $\mathscr{H}$ represents a physically realisable state: the formalism contains no "inert" vectors that are formally members of $\mathscr{H}$ but participate in no physical process. Several features of the standard formalism enforce this operationally:

- **Superposition closure.** If $|\psi_1\rangle$ and $|\psi_2\rangle$ are valid states, then $\alpha|\psi_1\rangle + \beta|\psi_2\rangle$ is a valid state for any $\alpha, \beta \in \mathbb{C}$. No state is isolated from the superposition structure.
- **Dynamical accessibility.** For any two pure states $|\psi\rangle$ and $|\varphi\rangle$ in the same superselection sector, there exists a unitary U such that $U|\psi\rangle = |\varphi\rangle$. Every state is dynamically reachable from every other.
- **Measurement participation.** Every state participates in the Born rule: for any observable $\hat{A}$, the expectation value $\langle\psi|\hat{A}|\psi\rangle$ is defined and physically meaningful.

Standard quantum theory excludes inert states not by explicit axiom, but by the structure of the Hilbert space formalism itself: the space is closed under the operations (superposition, unitary evolution, inner product) that constitute physical participation. CC_G is the pre-factual analogue: $\mathcal{P}$ is closed under admissible composition, so elements not compositionally reachable from the generators are not in the domain.

The Born rule dependence. If $\mathcal{P}$ contained isolated elements, the amplitude assignment $\alpha : \mathcal{P} \to \mathbb{S}$ would assign non-zero amplitude values to contributions that never appear in any composition. These amplitudes would be formally defined but operationally ungrounded. More precisely, the mapping from the formal pre-factual domain to the space of physically realised processes would fail to be surjective onto its operational domain, introducing a structural mismatch between representation and dynamics. The Born rule would then become ambiguous: inert contributions carry non-zero amplitudes but no operational meaning, so including or excluding them from the sum becomes a representational choice rather than a physically determined one. This violates the requirement that physical predictions be independent of representational redundancy. This is not merely a formal ambiguity: it would imply that physical predictions depend on representational choices rather than physical structure, contradicting the operational meaning of probability. CC removes this ambiguity by ensuring every element with non-zero amplitude participates in at least one composition.

10.2.2 Correspondence with Relativistic Spacetime Connectivity

In general relativity, certain physically well-behaved spacetime models — in particular, globally hyperbolic spacetimes — exhibit a CC-like structural commitment: formal membership in the spacetime manifold coincides with participation in geometry and causal structure. In such models, every point $p \in \mathscr{M}$ participates in:

- the metric geometry: $g_{\mu\nu}(p)$ contributes to curvature
- the Einstein field equations: $G_{\mu\nu}(p) = 8\pi G \cdot T_{\mu\nu}(p)$
- the causal structure: p lies within at least one causal domain

In such settings, spacetime contains no regions formally present but entirely disconnected from the geometry, dynamics, and causal structure.

Caveat. This correspondence is limited. Physically relevant spacetimes such as de Sitter space contain causal horizons that restrict accessibility without violating energy conditions or global hyperbolicity. The analogy with CC is therefore structural rather than exact: CC excludes compositionally disconnected elements outright, whereas relativistic theories permit limited causal accessibility while maintaining participation in the global geometric structure.

10.3 PAR and CC as Explicit Structural Extractions

Standard formulations of quantum theory and relativity assume reversibility at the level of underlying laws — unitary evolution, time-reversal invariance — and completeness of the physical domain — all states participate in superposition and dynamics; all spacetime points participate in geometry and causation. But these assumptions are embedded in the mathematical formalisms, distributed across multiple axioms without separation, and not identified as independent structural principles with distinct logical roles.

The VERSF reconstruction extracts these implicit conditions, names them, and separates their logical contributions. The analysis that isolates IAC into PAR and CC as its governing structural factors reveals that these two conditions — which in standard formulations are conflated in the structure of Hilbert spaces and manifolds — play genuinely independent logical roles. Each can fail independently (Section 5); each does independent work in Theorem 3.1.

Its contribution is not to add new physical content, but to isolate structural content already carried implicitly by established theories and assign its logical role explicitly. In this sense, PAR and CC should be understood not as additional assumptions imposed on the framework, but as factorisations of structural conditions that standard formulations encode implicitly within their mathematical frameworks — here extracted, named, and assigned distinct logical roles for the first time.

10.4 Asymmetry in the Correspondence

PAR's analogue in quantum theory — unitary invertibility — is well-established, well-named (unitarity), and its physical consequences are extensively studied. The correspondence between PAR and unitarity is therefore tight and well-supported.

CC's analogue — Hilbert space completeness and dynamical accessibility — is less explicitly articulated in standard treatments. It is enforced structurally by the definition of Hilbert space and by superselection sector theory, but is rarely stated as an independent physical principle. This asymmetry does not weaken the case for CC. It strengthens the case for the VERSF programme's contribution: by naming and isolating CC, the reconstruction makes explicit a condition that quantum theory enforces structurally without identifying as a principle.

10.5 Logical Status of the Correspondences

The correspondences established in this section are **structural analogies**, not derivations. The logical status of PAR and CC remains: both are sufficient for IAC (Theorem 3.1), both are physically well-motivated, and neither is yet derived from A2 alone. The derivation strategies are developed formally in Appendix A.

The identification of PAR and CC as explicit extractions of implicitly assumed structure sharpens the foundational question of the programme. The issue is no longer why quantum structure arises from abstract axioms, but why any fact-producing universe must organise its domain of possibilities into a reversible and structurally complete system. The VERSF programme reframes this not as a mathematical coincidence, but as a structural requirement tied to the very notion of fact formation. This question is now sharp enough to function not merely as an open problem, but as the central structural test of any attempt to derive quantum theory from first principles.

---

11. Open Problems

Primary targets (main paper):

P1. Derive PAR from A2 plus composition-order symmetry (formalised as PFS — see Appendix A). Characterise whether PFS is itself derivable from A2.

P2. Derive CC_G from the operational requirement that the pre-factual domain consists precisely of physically realisable alternatives. Determine whether "physically realisable" can be grounded in A2 alone.

P3. Determine whether PAR and CC can be unified into a single structural condition implying both — reducing the premise count above A2 to one.

Secondary targets (main paper):

P4. Characterise the class of structures satisfying PAR + CC. Does this class correspond to a known algebraic category (e.g., connected inverse semigroups)?

P5. Determine whether the gap between A2 and PAR has a structural counterpart in other reconstructive programmes, or is specific to the VERSF architecture.

P6. Investigate whether the weakened PAR (observational restorability) and the stronger PAR (algebraic inversion) are equivalent under additional constraints derivable within the programme.

Appendix targets:

P7. Derive PFS from A2 alone, completing the chain A2 $\Longrightarrow$ PFS $\Longrightarrow$ PAR (Appendix A, §A.6).

P8. Derive or replace Observational Faithfulness (Definition A.5) — the bridge condition linking information preservation to observational invertibility — from A2 or from PFS (Appendix A, §A.11).

P9. Derive the Pre-Factual No-Erasure Condition (PNEC) from A2, completing the Landauer route without independent thermodynamic motivation (Appendix B, §B.9).

P10. Proposition A.6 establishes that under Observational Faithfulness, no-pre-factual-information-loss implies PFS — settling one conditional direction of the equivalence question. The remaining open questions are: (a) does PFS imply Observational Faithfulness, or are they independent? (b) is PNEC formally equivalent to no-pre-factual-information-loss, or does the difference in grounding (thermodynamic vs. admissibility-based) mark a genuine structural distinction?

P11. Formalise the Landauer analogy beyond the structural level — determine whether PNEC can be given a quantitative grounding in the VERSF ontology as thermodynamic structure emerges from fact formation.

---

12. Conclusion

This paper establishes the following results with full formal rigour:

1. **Theorem 3.1:** PAR + CC $\Longrightarrow$ IAC
2. **Proposition 4.1:** A2 $\nvdash$ PAR (explicit counterexample via formal fact map $F : \mathcal{P} \to \Sigma$)
3. **Section 5:** PAR and CC are logically independent (explicit models 5.1 and 5.2)
4. **Section 6:** CC admits multiple formulations; CC_G (generative closure) is definitionally primary
5. **Section 7:** IAC decomposes into exactly two independent conditions, PAR and CC
6. **Section 10:** PAR and CC formalise structural conditions already implicit in quantum theory and relativity
7. **Theorem A.2:** A2 + PFS $\Longrightarrow$ PAR
8. **Theorem A.4:** A2 + Observational Faithfulness $\Longrightarrow$ PAR
9. **Theorem B.2:** PNEC $\Longrightarrow$ PAR

Six structural vulnerabilities present in earlier versions have been addressed: the amplitude codomain is weakened to $\mathbb{S}$; PAR is reformulated in terms of observational restorability; the A2 $\nvdash$ PAR counterexample uses a formal fact map; CC is reframed as the definition of $\mathcal{P}$ under CC_G; identity is treated as local; and the physical interpretation is calibrated precisely to what is proven.

Internal Admissible Closure is not a primitive assumption. It is the joint consequence of two independently motivated structural requirements:

No observational distinguishability is destroyed irrecoverably before facts are produced — **PAR**
The pre-factual domain is exactly its compositionally reachable content — **CC**

PAR is the unique structural condition identified by three independent derivation routes — admissibility logic (no ungrounded directional asymmetry), information-theoretic logic (no ungrounded elimination of distinction), and thermodynamic logic (no erasure without a physically grounded outlet) — each of which excludes pre-factual irreversibility for distinct physical reasons. All three converge on the same conclusion from different directions.

Whether PAR and CC are themselves derivable from deeper conditions is the **single remaining foundational question** of the programme. It is now posed with sufficient precision — eleven open problems, three theorem-level partial results, five named bridge conditions — to serve as a formal research target. In this sense, the problem of deriving quantum theory is reduced to a single structural question: why must any fact-producing universe forbid pre-factual irreversibility and exclude compositionally disconnected possibilities?

Appendix A: On the Derivation of Pre-Factual Algebraic Reversibility from Admissibility

---

A.1 Objective

The purpose of this appendix is to investigate whether PAR can be derived from A2, or whether it must be taken as an independent structural postulate.

The main results are:

Theorem A.2: A2 + PFS ⟹ PAR
Theorem A.4: A2 + Observational Faithfulness ⟹ PAR

where PFS is the Pre-Factual Symmetry Condition and Observational Faithfulness is defined in §A.11. Both are conditional derivations: PAR follows from A2 together with a minimal, physically motivated strengthening in each case. It does not follow from A2 alone.

The appendix also characterises precisely why A2 alone is insufficient (§A.2), identifies the remaining open problem (§A.6), and establishes the convergence of the two routes (§A.12).

A.2 Admissibility and the Absence of Fact-Grounded Asymmetry

Recall the admissibility condition:

(A2) A distinction between elements of $\mathcal{P}$ is physically meaningful if and only if it is traceable to a difference in fact outcomes via $F : \mathcal{P} \to \Sigma$.

What A2 does not say. A2 constrains distinguishability. It does not directly constrain *directionality* — the question of whether a transition $r_1 \circ r_2$ can proceed one way but not the other. A structure could satisfy A2 while containing algebraically irreversible transitions, provided those transitions do not introduce new observable distinctions. Proposition 4.1 exhibits exactly this: admissibility holds while a compositional step collapses to 0, making recovery impossible.

This gap motivates the following definition.

Definition A.1 (Pre-Factual Symmetry Condition — PFS).

The pre-factual domain $\mathcal{P}$ satisfies PFS if and only if:

No physically meaningful directional asymmetry exists within $\mathcal{P}$ that is not grounded in a formed fact.

Formally: for any non-trivial composition $r_1 \circ r_2$ in $\mathcal{P}$, if the assignment of a preferred compositional direction — forward or backward — would constitute a physically meaningful distinction, then that distinction must be traceable to a fact outcome in Σ. Since $\mathcal{P}$ contains no formed facts, no preferred compositional direction is admissible.

Relation to A2. PFS extends the scope of admissibility from *distinctions between elements* to *asymmetries between directions of transition*. It is strictly stronger than A2: the counterexample of Proposition 4.1 satisfies A2 while violating PFS (the irreversible transition $c \circ b = 0$ introduces a directional asymmetry — forward step possible, restoring step not — not grounded in any fact).

A.3 From PFS to Reversibility

Setup. Let $r_1 \circ r_2$ be a non-trivial composition in $\mathcal{P}$, where r_2 is the prior contribution and r_1 the subsequent step.

Assume for contradiction that no restoring process exists:

There is no $\psi \in \mathcal{P}$ such that $F((r_1 \circ r_2) \circ \psi) = F(r_2)$

The irreversibility defines a directional asymmetry. The forward direction (r_2 to $r_1 \circ r_2$) is physically available. The backward direction (recovery of $F(r_2)$) is not. This constitutes a preferred compositional direction.

The asymmetry is not fact-grounded. The composition occurs entirely within $\mathcal{P}$. No fact has been formed. The directional asymmetry cannot be traced to any element of Σ.

PFS forbids this. By PFS, no physically meaningful directional asymmetry exists in $\mathcal{P}$ that is not grounded in a formed fact. Contradiction.

Proposition A.1.

Under A2 and PFS, every non-trivial pre-factual composition is observationally restorable.

Proof. By the argument above: assume irreversibility → defines preferred direction → not fact-grounded → PFS forbids it → contradiction. The result holds universally. ■

A.4 Main Result: Admissibility Route

Theorem A.2 (Conditional Derivation of PAR — Admissibility Route).

A2 + PFS ⟹ PAR

Proof. PAR asserts that for every non-trivial composition $r_1 \circ r_2$, there exists ψ_{r_1,r_2} such that $F((r_1 \circ r_2) \circ \psi_{r_1,r_2}) = F(r_2)$. Proposition A.1 establishes precisely this under A2 and PFS. ■

Corollary A.3 (Necessity of PFS).

A2 alone does not imply PAR. Proof: Proposition 4.1 exhibits a structure satisfying A2 but violating PAR. That structure also violates PFS. So PFS is not redundant — it does logical work that A2 cannot. ■

Logical map:

```
A2 alone
   ↓  [insufficient — Proposition 4.1]
A2 + PFS
   ↓  [Theorem A.2]
PAR
   ↓  [with CC, Theorem 3.1]
IAC
```

A.5 Interpretation of the PFS Route

PAR is not an arbitrary assumption. It is the algebraic expression of a single, physically grounded requirement:

No irreversibility without facts.

Before facts exist, there is nothing to ground a preferred direction of transition. PFS extends the scope of admissibility from element-distinctions to directional-asymmetries. The physical motivation is identical: facts are the only source of physical reality in the pre-factual setting; therefore nothing physically real, whether a distinction or a directionality, can exist prior to them.

The strengthening from A2 to A2 + PFS is minimal. PFS does not introduce new physical content in the sense of positing new entities or processes. It extends the scope of an existing condition. If PFS can be derived from A2 — a question addressed in §A.6 — then PAR is not an independent postulate at all.

A.6 Remaining Gap and Derivation Strategy for PFS

Can PFS itself be derived from A2?

This would complete the chain A2 ⟹ PAR. The present state is that this derivation has not been achieved. We characterise the gap and propose a strategy.

What derivation of PFS from A2 would require. PFS states that no physically meaningful directional asymmetry exists in $\mathcal{P}$ that is not grounded in a fact. Derivation from A2 would require showing that any admissible structure must treat forward and reverse directions of every pre-factual composition as equally available.

Proposed derivation strategy. Suppose the pre-factual domain contains a compositional asymmetry: forward direction of $r_1 \circ r_2$ is available, but no restoring process exists. Can any physical agent, operating only within the pre-factual domain, detect the asymmetry before any fact is formed?

If detectable before fact formation: it constitutes a physically meaningful distinction not grounded in facts — violating A2 directly.
If undetectable before fact formation: it has no physical meaning and cannot be said to exist in any physically real sense — making it inadmissible as a structural feature.

In either case, the asymmetry is inadmissible. This argument, if formalised, would derive PFS from A2. Formalisation requires a precise account of "detectability within the pre-factual domain before fact fixation." Detection here refers to any operational distinction definable within the admissible structure of $\mathcal{P}$ prior to commitment, not to post-factual measurement outcomes. This is a well-posed technical problem and the target of P7.

If PFS is derived from A2:

```
A2  ⟹  PFS  ⟹  PAR  [with CC] ⟹  IAC  ⟹  Quantum structure
```

Reversibility of the pre-factual domain would then be a consequence of admissibility rather than an independent postulate.

A.7 Conclusion of the PFS Route

Implication	Status
A2 ⟹ PAR	✗ False — Proposition 4.1
A2 + PFS ⟹ PAR	✓ Theorem A.2
A2 ⟹ PFS	Open — strategy in §A.6
A2 + PFS + CC ⟹ IAC	✓ Theorem A.2 + Theorem 3.1

A.8 Information-Theoretic Formulation of Pre-Factual Reversibility

We now offer a complementary route framing the same conclusion in information-theoretic terms. The central observation:

Algebraic irreversibility corresponds to the destruction of distinguishable information.

If no information can be destroyed in the pre-factual domain — a condition derived from A2 below — then no transition can be irreversible, and PAR follows.

Definition A.3 (Distinguishability Measure).

Let $\mathcal{D}(r)$ denote the set of elements of $\mathcal{P}$ observationally distinguishable from r — all $r' \in \mathcal{P}$ with $F(r) \neq F(r')$. Define the *information content* of r:

$$I(r) = \log |\mathcal{D}(r)|$$

where $|\mathcal{D}(r)|$ is the cardinality of the distinguishable neighbourhood of r under F. This is a distinguishability measure, not Shannon entropy. It is defined entirely in terms of F, so every quantity it tracks is fact-traceable — consistent with A2.

Definition A.4 (Information Loss under Composition).

A composition $r_1 \circ r_2$ is *information-losing* if:

$$I(r_1 \circ r_2) < I(r_2)$$

The composed element has a smaller distinguishable neighbourhood than the input r_2.

A.9 Information Loss and Algebraic Irreversibility

Proposition A.4 (Irreversibility Implies Information Loss).

If $r_1 \circ r_2$ is algebraically irreversible — no ψ exists such that $F((r_1 \circ r_2) \circ \psi) = F(r_2)$ — then $r_1 \circ r_2$ is information-losing: $I(r_1 \circ r_2) < I(r_2)$.

Proof. Suppose no restoring process ψ exists. Then $F(r_2) \notin \text{Reach}(r_1 \circ r_2)$, where $\text{Reach}(r) = \{F(r)\} \cup \{F(r \circ \psi) : \psi \in \mathcal{P}, r \circ \psi \text{ defined}\}$. Since $F(r_2) \in \text{Reach}(r_2)$ by the self-reachability base case (see §2.3 Remark), the reachability set of $r_1 \circ r_2$ is a strict subset of that of r_2. It follows that $\mathcal{D}(r_1 \circ r_2) \subsetneq \mathcal{D}(r_2)$, and therefore $I(r_1 \circ r_2) < I(r_2)$. ■

Remark. The converse — information loss implies irreversibility — requires an additional condition (Observational Faithfulness, §A.11) and is not automatic.

A.10 No Pre-Factual Information Loss from Admissibility

Proposition A.5 (No Pre-Factual Information Loss).

Under A2, any destruction of distinguishable information in $\mathcal{P}$ must be grounded in a formed fact. Since $\mathcal{P}$ contains no formed facts, no information loss is admissible in $\mathcal{P}$.

Proof. Suppose $r_1 \circ r_2$ is information-losing: $I(r_1 \circ r_2) < I(r_2)$. Then elements previously distinguishable from r_2 have become indistinguishable from $r_1 \circ r_2$ — a distinction has been eliminated. Under A2, the elimination of a distinction is a physically meaningful event (it changes which elements can be discriminated). For it to be admissible, it must be traceable to a fact-formation event. No fact-formation event has occurred within $\mathcal{P}$. Thus any pre-factual information loss would constitute an admissible, fact-ungrounded elimination of distinction, contradicting A2. Therefore no information-losing composition can occur in $\mathcal{P}$ under A2. ■

A.11 From Information Preservation to PAR

Proposition A.5 establishes that all compositions in $\mathcal{P}$ are information-preserving under A2:

$I(r_1 \circ r_2) = I(r_2)$ for all non-trivial compositions in $\mathcal{P}$

The gap. Information preservation — $I(r_1 \circ r_2) = I(r_2)$ — asserts that the *cardinality* of the distinguishable neighbourhood is preserved. It does not automatically follow that the specific fact value $F(r_2)$ is recoverable from $r_1 \circ r_2$. Cardinality preservation is compatible with a rearrangement of which elements are distinguishable — a permutation preserving $|\mathcal{D}|$ while changing $\mathcal{D}$ itself.

To close this gap, we require:

Definition A.5 (Observational Faithfulness).

The distinguishability measure I is *observationally faithful* if:

$I(r_1 \circ r_2) = I(r_2)$ implies that the map $r_2 \mapsto r_1 \circ r_2$ is an observational isomorphism — there exists ψ such that $F((r_1 \circ r_2) \circ \psi) = F(r_2)$.

Remark. This is an additional condition on the relationship between I and F. We flag it explicitly rather than absorbing it silently. It holds when the distinguishability structure of $\mathcal{P}$ is sufficiently fine-grained that cardinality preservation implies structural recovery.

Theorem A.4 (Information-Theoretic Derivation of PAR).

Under A2, Proposition A.5 (no-pre-factual-information-loss), and Observational Faithfulness (Definition A.5), PAR holds.

Proof. By Proposition A.5, all compositions are information-preserving: $I(r_1 \circ r_2) = I(r_2)$. By Observational Faithfulness, information preservation implies observational invertibility: there exists ψ with $F((r_1 \circ r_2) \circ \psi) = F(r_2)$. Since $r_1 \circ r_2$ was arbitrary, PAR holds universally. ■

Summary of conditions:

Condition	Role
A2 (Admissibility)	Grounds all distinctions in facts
Proposition A.5	Derives no-information-loss from A2
Observational Faithfulness	Bridges information preservation to restorability
Conclusion	**PAR**

A.12 Relation Between the Two Routes

The PFS route (§A.3) proceeds via *directional asymmetry*: an irreversible transition introduces a preferred direction not grounded in facts, which is inadmissible under A2 + PFS.

The information-theoretic route proceeds via *distinguishability collapse*: an irreversible transition destroys distinguishable information, which is inadmissible under A2 (since the collapse is a fact-ungrounded elimination of a distinction).

Both routes track the same underlying phenomenon — irreversibility as the destruction of pre-factual structure — from different angles. The PFS route requires PFS; the information-theoretic route requires Observational Faithfulness. Neither bridge condition is trivially derivable from A2 alone, but both are independently motivated.

Proposition A.6 (Convergence).

If Observational Faithfulness holds, then no-pre-factual-information-loss (from A2) implies PFS.

Proof sketch. An irreversible transition introduces both a directional asymmetry and an information loss. Under Observational Faithfulness, the information loss is sufficient to establish the directional asymmetry: if $F(r_2)$ is not recoverable, the forward-only character of the transition is precisely the directional asymmetry that PFS forbids. ■

The two routes therefore unify under Observational Faithfulness. The question of which bridge condition is more fundamental — PFS or Observational Faithfulness — is itself target P10.

Appendix B: Landauer-Type Constraint and the Thermodynamic Grounding of Pre-Factual Reversibility

---

B.1 Objective

This appendix strengthens the case for PAR by connecting it to the thermodynamic logic of irreversibility. The central idea is not that Landauer's principle applies unchanged to the pre-factual domain — temperature, heat flow, and bath coupling belong to already-instantiated physical regimes. Rather, Landauer's principle provides the established physical model for a structural fact that is more general:

Genuine irreversibility requires a cost-bearing outlet.

In the VERSF framework, fact formation is the first point at which such an outlet can be physically grounded. Prior to fact formation, no irreversible record exists and no thermodynamic sink is available. Pre-factual irreversibility is therefore not merely unmotivated — it is physically ungroundable.

The main result:

Theorem B.2: PNEC ⟹ PAR

B.2 Landauer's Principle as the Standard Template

Landauer's principle states that the erasure of one bit of logically irreversible information carries a minimum thermodynamic cost:

$\Delta Q \geq k_B T \ln 2$

in the presence of a thermal environment at temperature T, where k_B is Boltzmann's constant.

The key structural content:

- Logical irreversibility is not free
- Erasure requires export of entropy into an environment
- Irreversibility becomes physical only when distinguishable information is no longer recoverable from the system plus environment

A logically irreversible map reduces distinguishability: multiple prior states are sent to the same later state. In any physical realisation this requires entropy export to an environment that absorbs the lost structure. Erasure requires a physically present sink.

B.3 The Absence of an Irreversibility Sink in the Pre-Factual Domain

In the VERSF setting, $\mathcal{P}$ is the domain of unresolved alternatives prior to any commitment. By definition:

- No fact has yet formed
- No irreversible record yet exists
- No committed macroscopic branch has been singled out
- No thermodynamically grounded sink for lost distinguishability has yet appeared

The Landauer template requires exactly such a sink. If distinguishability is genuinely destroyed — if previously distinct pre-factual elements are rendered indistinguishable — the lost information must be offloaded somewhere: a record, an environment, a heat bath.

Before fact formation, no such structure is available. The first physically grounded sink appears precisely at fact formation: the committed fact is itself the record into which the cost of irreversibility is paid.

Structural inference. If a pre-factual composition were genuinely irreversible, it would require an information-loss channel whose physical ground does not yet exist — consuming a resource that the pre-factual domain does not yet contain.

B.4 Pre-Factual Irreversibility as Premature Erasure

Consider a non-trivial pre-factual composition $r_1 \circ r_2$. Suppose no restoring process $\psi_{\{r_1,r_2\}}$ exists:

$F((r_1 \circ r_2) \circ \psi_{\{r_1,r_2\}}) = F(r_2)$ has no solution

Then some distinguishable structure present in r_2 has been erased by the composition with r_1. By the Landauer template, such erasure requires grounded entropy export. Yet no fact has formed and no physical sink exists.

The one-way transition therefore amounts to **premature erasure without a physical erasure channel**: distinguishability destruction consuming a resource the domain does not yet possess.

B.5 The Pre-Factual No-Erasure Condition

Definition B.1 (Pre-Factual No-Erasure Condition — PNEC).

No genuinely irreversible loss of distinguishable structure occurs in the pre-factual domain $\mathcal{P}$ prior to fact formation.

Equivalently: any apparent reduction in distinguishability within $\mathcal{P}$ prior to commitment is either only apparent (a coarse-graining), observationally reversible, or not grounded as a physical event within $\mathcal{P}$.

Relation to Proposition A.5. PNEC is the thermodynamic analogue of no-pre-factual-information-loss (Appendix A, §A.10). Both prohibit the destruction of distinguishable structure in $\mathcal{P}$. They differ in their grounding:

- Proposition A.5 derives no-information-loss from A2: the elimination of a distinction is inadmissible because it is not fact-grounded.
- PNEC grounds no-erasure in the Landauer template: erasure requires a physical outlet that does not yet exist in $\mathcal{P}$.

These are independently motivated routes to the same structural prohibition.

B.6 From No-Erasure to PAR

Theorem B.2 (Landauer-Type Route to PAR).

If $\mathcal{P}$ satisfies PNEC, then all non-trivial pre-factual compositions are observationally restorable:

PNEC ⟹ PAR

Proof. Let $r_1 \circ r_2$ be any non-trivial composition in $\mathcal{P}$. Suppose for contradiction that no restoring process $\psi_{\{r_1,r_2\}}$ exists. Then, by Proposition A.4, the composition is information-losing: $I(r_1 \circ r_2) < I(r_2)$. This constitutes an irreversible loss of distinguishable structure in $\mathcal{P}$ — a genuine erasure event. By PNEC, no such erasure is permitted prior to fact formation. Contradiction. Therefore $\psi_{\{r_1,r_2\}}$ must exist for every non-trivial composition, and PAR holds. ■

Remark. This proof invokes only the direction "irreversibility ⟹ information loss" (Proposition A.4). It does not require Observational Faithfulness. PNEC ⟹ PAR is therefore a cleaner route than Theorem A.4 in one respect: it requires fewer bridge conditions.

B.7 Relation to the Three Routes

Three independent conditional routes to PAR have now been established:

Route	Key condition	Bridge needed	Main result
Admissibility (§A.3)	A2 + PFS	PFS	A2 + PFS ⟹ PAR
Information-theoretic (§A.8–A.11)	A2 + Obs. Faithfulness	Observational Faithfulness	A2 + OF ⟹ PAR
Landauer (this appendix)	PNEC	—	PNEC ⟹ PAR

Each route identifies a distinct physical reason why the pre-factual domain cannot support genuine irreversibility:

- **Admissibility route:** a directional asymmetry without a fact to ground it is inadmissible.
- **Information-theoretic route:** a distinction lost without a fact to eliminate it is inadmissible.
- **Landauer route:** an erasure without a physical outlet to absorb it is structurally unsupported.

These are not three restatements of the same argument. Each invokes a different physical concept — directionality, distinguishability, thermodynamic cost — and requires a different bridge condition. Their convergence on PAR from three independent directions is the strongest available evidence that PAR is a genuine structural necessity of the pre-factual domain, not an arbitrary postulate.

Relation between PNEC and Proposition A.5:

A2 ⟹ No-pre-factual-information-loss (Proposition A.5)
Landauer template + absence of physical sink ⟹ PNEC
PNEC ≃ No-pre-factual-information-loss (same structural prohibition, different grounding)

Whether PNEC and no-pre-factual-information-loss are formally equivalent is target P10.

B.8 Relation to Established Physical Theories

Quantum theory. Unitary evolution is information-preserving: $U(t)^\dagger U(t) = \mathbb{1}$ ensures the full state-space structure is recoverable at every time. Information loss in quantum theory is associated exclusively with measurement — a fact-formation event in the VERSF framework. The Landauer structure is exact: before measurement, no erasure channel is open; at measurement, the committed outcome is precisely the physical record into which the cost of irreversibility is paid.

Black hole physics. The dominant modern position, supported by Hawking radiation analyses and holographic duality (AdS/CFT), holds that information is preserved in black hole dynamics: apparent loss is due to coarse-graining over microstates, not fundamental destruction. *Caveat: the information paradox remains an active research area; this reflects current theoretical consensus rather than established fact.*

Thermodynamics and statistical mechanics. In Boltzmann's framework, thermodynamic entropy increase reflects coarse-graining over microstates, not the destruction of microstate information. Liouville's theorem guarantees that phase space volume is preserved under Hamiltonian dynamics. Apparent information loss in thermodynamics is always a consequence of restricting attention to macroscopic variables. The VERSF no-erasure condition is the pre-factual expression of this: at the level of underlying structure, nothing is destroyed.

B.9 Caveats and Scope

Limitation 1: No numerical Landauer bound. The claim is not that $\Delta Q \geq k_B T \ln 2$ applies to the pre-factual domain. Temperature, thermal equilibration, and heat bath coupling belong to already-instantiated physical regimes. The bound is cited as a template for structural logic, not as a directly applicable formula.

Limitation 2: PNEC is not derived from A2. Unlike the information-theoretic route, the Landauer route does not derive PNEC from admissibility. PNEC is independently motivated by the Landauer template and the absence of a physical sink in $\mathcal{P}$. Whether PNEC can be derived from A2 is open target P9.

Limitation 3: The analogy is between structural roles, not ontological levels. Landauer's principle operates within thermodynamics; PNEC operates at the level of pre-factual structure from which thermodynamics eventually emerges. The correspondence is between the *structural role* that irreversibility cost plays in each setting, not between the physical mechanisms.

B.10 Conclusion

Landauer's principle establishes that erasure is not free: genuine irreversibility requires a physically grounded cost-bearing outlet. The pre-factual domain contains no such outlet. Pre-factual irreversibility would amount to premature erasure — distinguishability destruction consuming a physical resource that the domain does not yet possess.

Irreversibility requires cost. Cost requires a fact-grounded outlet. Before facts exist, no such outlet is available. Therefore before facts, no irreversibility.

Theorem B.2 formalises this as PNEC ⟹ PAR. Landauer-type reasoning provides the strongest available physical motivation for PNEC — as the only mechanism known in physics by which logical irreversibility acquires physical meaning, it provides the clearest available motivation for why no erasure channel exists before facts form — and PNEC implies PAR. The claim is not that Landauer's principle derives PAR directly; it is that the structural logic of Landauer, applied to the pre-factual setting, gives compelling independent grounds for PNEC, which in turn closes the gap to PAR. Together with the admissibility route and the information-theoretic route, this gives three independent convergent arguments for the same conclusion. PAR is the point at which admissibility logic, information-theoretic logic, and thermodynamic logic all arrive together.

Full logical map across both appendices:

```
A2
 ├─→ [Admissibility route, §A.3]
 │       A2 + PFS ⟹ PAR  [Theorem A.2]
 └─→ [Information-theoretic route, §A.8–A.11]
         A2 ⟹ No info loss  [Proposition A.5]
         No info loss + Observational Faithfulness ⟹ PAR  [Theorem A.4]
         PFS ⟺ No info loss  [under Observational Faithfulness, Proposition
A.6]

Landauer template + absence of physical sink in 𝒫
 └─→ [Landauer route, §B.3–B.6]
         PNEC ⟹ PAR  [Theorem B.2]
         PNEC ≃ No info loss  [same prohibition, different grounding]
```

```
All three routes ⟹ PAR  [conditionally]

Open targets (see §11):
  P7:  Can PFS be derived from A2?
  P8:  Can Observational Faithfulness be derived from A2 or PFS?
  P9:  Can PNEC be derived from A2?
  P10: Are PFS, Observational Faithfulness, and PNEC equivalent?
  P11: Can the Landauer analogy be formalised beyond structural analogy?
```

Appendix C: Notation Summary

Symbol	Meaning
$\mathcal{P}$	Pre-factual domain
$\circ$	Partial composition on $\mathcal{P}$ (associative where defined)
0	Null element of $\mathcal{P}$
e_U	Local identity for compositional substructure $U \subseteq \mathcal{P}$
$\mathbb{S}$	Minimal scalar structure (amplitude codomain)
α	Amplitude assignment $\alpha : \mathcal{P} \to \mathbb{S}$
Σ	Set of fact outcomes
F	Fact map $F : \mathcal{P} \to \Sigma$
ND	Non-Degeneracy condition on the generating set G of $\mathcal{P}$ (§2.1)
$\oplus$	Combination rule on Σ (from F2); operates on fact outcomes in Σ, distinct from the composition operation $\circ$ on $\mathcal{P}$
G	Generating set of $\mathcal{P}$
$cl_\circ$	Closure under $\circ$
$\psi_{r,q}$	Restoring operation for the decomposition $r \circ q$
$\mathcal{D}(r)$	Distinguishable neighbourhood of r under F
I(r)	Information content: log
A2	Admissibility condition
$A2^+$	Strengthened admissibility (includes restorability)
PAR	Pre-Factual Algebraic Reversibility
CC	Compositional Completeness
CC_R	Reachability formulation of CC
CC_G	Generative closure formulation of CC (primary)
CC_N	No-isolated-states formulation of CC
IAC	Internal Admissible Closure
PFS	Pre-Factual Symmetry Condition
PNEC	Pre-Factual No-Erasure Condition
k_B	Boltzmann's constant

Symbol	Meaning
$\nvdash$	Does not derive
$\Longrightarrow$	Implies

Appendix D: Responses to Anticipated Referee Concerns

R1. "You assume the normed division algebra that you later claim to derive."
Addressed in §2.2. The amplitude codomain is now the minimal scalar structure $\mathbb{S}$, carrying only addition and scalar compatibility. No norm, multiplicative inverse, or division structure is assumed. The derivation that $\mathbb{S}$ must be a normed division algebra belongs to a later stage of the programme and is explicitly not imported here.

R2. "PAR assumes global inversion — the very group structure you are trying to prove."
Addressed in §2.5. PAR is now formulated in terms of observational restorability: the existence of a process $\psi_{r,q}$ such that $F((r \circ q) \circ \psi_{r,q}) = F(q)$. This is strictly weaker than algebraic inversion. No global identity, no specific inverse element, and no group structure is assumed.

R3. "The counterexample for A2 $\nvdash$ PAR is ad hoc."
Addressed in §2.3 and §4. The fact map $F : \mathcal{P} \to \Sigma$ is a formal object with three explicitly stated properties (F1–F3). Admissibility is verified by checking F-injectivity on non-null elements. PAR failure is demonstrated by showing that $c \circ b = 0$ with $F(0) = \emptyset \neq \sigma_b = F(b)$, making recovery impossible.

R4. "CC is just a closure axiom."
Addressed in §2.1 and §6.2. Under CC_G, $\mathcal{P}$ is *defined* as $cl_\circ(G)$. Isolated elements are not forbidden — they are simply not members of $\mathcal{P}$ by definition. CC is a consequence of how the domain is constructed, not an added axiom.

R5. "The identity element e is undefined."
Addressed in §2.4. e_U is explicitly defined as a local identity for a compositional substructure $U \subseteq \mathcal{P}$. No global identity is assumed. PAR is stated without reference to a global e.

R6. "Calling $\mathcal{P}$ a connected, invertible algebraic system overstates what has been proven."
Addressed in §8.3. $\mathcal{P}$ is characterised as *connected and algebraically reversible under PAR + CC*. The stronger claim — that $\mathcal{P}$ carries full invertible group structure — is noted as a downstream result, not an assumption.

R7. "The PAR–unitarity analogy is imprecise."
Addressed in §10.1.1 (Remark: PAR vs. unitarity). Unitary evolution requires strict operator inversion ($U^\dagger U = \mathbb{1}$); PAR requires only observational recoverability via F. The distinction is intentional and explained: PAR is formulated at the level of pre-factual alternatives, where F-

recoverability is the physically appropriate condition. The gap reflects a difference in level, not an omission.

R8. "The Landauer analogy imports thermodynamic structure that doesn't yet exist at the pre-factual level."
Addressed in Appendix B, §B.9. The claim is explicitly structural, not quantitative. The numerical Landauer bound $\Delta Q \geq k_B T \ln 2$ is not applied. The argument uses only the structural role of the Landauer template: irreversibility requires a cost-bearing outlet, and no such outlet exists in $\mathcal{P}$ prior to fact formation. Limitation 3 of §B.9 states explicitly that the correspondence is between structural roles, not ontological levels.