Internal Admissible Closure from Pre-Factual Reversibility in Fact-Producing Universes

Keith Taylor
VERSF Theoretical Physics Program

For the General Reader

Physics describes a world full of possibilities — but what we actually observe are definite outcomes. A detector clicks. A particle lands. A fact is recorded.

In earlier work, we showed that requiring the universe to produce real, irreversible facts already forces strong constraints on how it must behave. Those constraints lead naturally to interference, complex numbers, and the mathematical structure of quantum mechanics.

But one question remained:

Why must all possible contributions be able to cancel or recombine cleanly?

In this paper, we show that this follows from a simple but powerful principle:

Everything before a fact is formed must remain genuinely reversible.

If some part of a system could not be undone before a fact is produced, then irreversibility would appear too early — before any actual outcome has occurred. Possibilities are supposed to stay open until the moment a fact forms. A contribution that cannot be cancelled has, in effect, already settled something, even though no commitment has been made. That would collapse the distinction between possibility and fact.

We make this precise by distinguishing two kinds of reversibility — dynamical and algebraic — and showing that the pre-factual sector must be reversible in both senses. From this, we prove:

All internally realised contributions must admit cancellation partners.

This principle — Internal Admissible Closure (IAC) — is not derived from nothing. It follows from requiring that irreversibility belongs exclusively to facts, not to the underlying possibilities from which facts emerge.

Abstract

Previous work in the VERSF reconstruction programme derives key structural features of quantum theory from the requirement that physical reality consists of stable, irreversible facts arising from pre-factual alternatives. One remaining open condition was internal admissible closure (IAC), previously introduced as a minimal closure principle ensuring that all internally realised contributions admit cancellation partners.

In this paper we derive IAC from two physically motivated principles. The first is **Pre-Factual Algebraic Reversibility** (PAR): every non-trivial internally realised transition in the recombination domain admits an algebraic reversal, because any one-way transition within the pre-factual sector would constitute a form of irreversibility forbidden by the Separation of Levels. The second is **Compositional Completeness** (CC): every non-null pre-factual contribution can occur as an internally realised component of some non-trivial pre-factual state, ensuring that every non-null $r \in \mathcal{P}$ falls within PAR's non-trivial scope.

The proof proceeds in two stages. First, Theorem 6.1 shows that PAR excludes irreducible residuals arising from non-trivial internally realised decompositions — those with $\psi^0 \neq 0$. Second, Theorem 7.1 uses CC to extend that exclusion to all non-null pre-factual contributions, yielding full IAC. The paper is explicit that it does not derive IAC from admissibility alone: the result is

PAR + CC $\Longrightarrow$ IAC

Both PAR and CC are new structural postulates whose deeper grounding from (A1)–(A2) alone is identified as the primary open question. We also clarify the level distinction between the non-negative weight structure of $\mathcal{P}$ and the signed additive structure of the amplitude algebra **A**, in which condition (S2) ultimately holds.

Table of Contents

---

1. Introduction

The VERSF Bridge Theorem establishes that any theory capable of producing stable, irreversible facts from unresolved alternatives must exhibit non-classical composition, structural interference, division-algebraic amplitude structure, and representational invariance. Subject to imported distinguishability geometry results, complex Hilbert space is the unique structure satisfying all requirements.

One structural condition remained as the primary open question:

(IAC) Internal admissible closure — no non-null internally realised affine sub-contribution remains algebraically unreversible; equivalently, every non-null admissibility class $[r] \in \mathbf{A}$ admits an additive inverse $[-r]$ such that $[r] + [-r] = [0]$ in the amplitude algebra $\mathbf{A}$.

Earlier work (*Closing the Structural Gaps*, Taylor, hereafter CSG) reduced a broad reversible-closure condition to IAC, identified it as irreducible to admissibility (A2) alone, and gave physical motivation without a formal derivation. The present paper provides that derivation.

In this paper, *pre-factual reversibility* is not merely dynamical time-symmetry. It means dynamical reversibility together with PAR and CC as the algebraic and structural conditions needed to prevent one-way internal transitions: PAR forbids algebraically irreversible non-trivial transitions within the pre-factual domain, and CC ensures every non-null pre-factual contribution lies within the non-trivial compositional domain over which PAR operates. The title is defensible precisely because these two conditions are what "pre-factual reversibility" must mean in order to entail IAC.

Honest statement of the contribution. The paper derives IAC from two principles: PAR (Pre-Factual Algebraic Reversibility) and CC (Compositional Completeness). PAR governs non-trivial internally realised transitions — it says such transitions cannot be algebraically one-way within the pre-factual domain. CC is a new structural postulate: it asserts that every non-null pre-factual state belongs to the internal compositional domain of $\mathcal{P}$, in the sense that it can occur as a sub-contribution within some non-trivial internally realised state. Together they give full IAC: CC ensures every non-null r falls under PAR's non-trivial scope; PAR then guarantees the cancellation partner.

This is a more transparent chain than earlier approaches, which closed the gap to full IAC via the formally trivial decomposition $r = 0 + r$. CC replaces that formally valid but physically hollow move with an explicit structural completeness axiom: a contribution that could never appear inside any non-trivial structured pre-factual state would be formally in $\mathcal{P}$ but compositionally disconnected from it — idle structure that the pre-factual sector's role as the domain of fact-producing alternatives does not admit. The paper's honest scope is: IAC follows from PAR + CC; both are physically motivated; neither is yet derived from (A1)–(A2) alone; this is the primary open question.

---

2. Background and Notation

We work within the framework of CSG. The basic structures are:

- **Λ** — a set of distinguishable configurations with distinguishability relation $\delta: \Lambda \times \Lambda \to [0,1]$
- **$\mathcal{P}$** — the pre-factual sector of pre-commitment states
- **w: $\mathcal{P}$ × Λ → [0,1]** — the weight function, with $w(\psi, \lambda) \geq 0$ for all ψ, λ and ∑_λ w(ψ, λ) = 1 for all normalised $\psi \in \mathcal{P}$
- **φ: $\mathcal{P}$ → Λ** — the irreversible commitment map
- **⋆_t** — a one-parameter family of composition rules ($t \in [0,1]$) on $\mathcal{P}$, including the identity at $t = 0$
- **(A2)** — admissibility: only distinctions traceable to recordable facts carry physical meaning
- **(A2c)** — admissibility closed under composition: composites of admissible operations are admissible

The null element. The **null alternative** 0 satisfies $w(0, \lambda) = 0$ for all $\lambda \in \Lambda$, giving ∑_λ w(0, λ) = 0. It is a formally adjoined boundary element of the pre-factual domain — not a normalised state in $\mathcal{P}$, but the additive identity of the recombination algebra **A**. It represents the complete absence of pre-factual content. We write $\mathcal{P}_0 := \mathcal{P} \cup \{0\}$ for the pre-factual sector together with the null element.

Observational equivalence. States $\psi, \psi' \in \mathcal{P}_0$ are *observationally equivalent*, written $\psi \cong \psi'$, if and only if $w(\psi, \lambda) = w(\psi', \lambda)$ for all $\lambda \in \Lambda$. In particular, $\psi \cong 0$ if and only if $w(\psi, \lambda) = 0$ for all λ.

Affine weight decomposition. We write $\psi = \psi' + r$ to mean:

$$w(\psi, \lambda) = w(\psi', \lambda) + w(r, \lambda) \ \forall \ \lambda \in \Lambda$$

where $w(\psi, \lambda), w(\psi', \lambda), w(r, \lambda) \geq 0$. This decomposition is *internally realised* when $\psi, \psi', r \in \mathcal{P}_0$. Decompositions requiring negative weights are excluded: for example, if $w(\psi, \lambda_0) < w(\psi', \lambda_0)$, the formal difference takes a negative value at λ_0 and has no pre-factual interpretation.

Critical level distinction. The paper operates at two levels:

- *Weight-level ($\mathcal{P}$):* affine addition $\psi = \psi' + r$ as defined above, with all weights in [0,1]. No element of $\mathcal{P}$ has negative weights.
- *Algebraic level (**A**):* the recombination algebra **A** on admissibility classes $[\psi]$. Once (S2) is established, **A** carries signed structure that extends beyond $\mathcal{P}$ as a weight space — additive inverses $[-\psi]$ exist in **A** without corresponding to elements of $\mathcal{P}$ with negative weights.

IAC and (S2) are statements about **A**, not claims that elements of $\mathcal{P}$ have weight-level inverses in $\mathcal{P}$. In its pre-algebraic formulation, IAC is expressed via the absence of irreducible residuals; at the algebraic level this is equivalent to the existence of additive inverse classes in **A**. The proof of Theorem 6.1 works at the weight level to rule out irreducible residuals within PAR's scope; Theorem 7.1 extends this to all non-null r via CC, and the result is lifted to **A** by faithfulness. The apparent tension — a weight-level cancellation partner would require $w(r, \lambda) = 0$ — is exactly the contradiction exploited to rule out non-null irreducible residuals, not a claim that non-trivial weight-level cancellation occurs.

The recombination algebra. By Proposition 3.6b of CSG, the internally realised affine decompositions of $\mathcal{P}_0$ generate a recombination algebra **A** on admissibility classes $[\psi]$. Concretely: **A** is generated by all classes $[\psi]$ with $\psi \in \mathcal{P}_0$, together with the addition $[\psi'] + [r] = [\psi]$ induced by any internally realised decomposition $\psi = \psi' + r$. The null element of **A** is [0]. The forward transitions $\psi' \to \psi$ arising from internally realised decompositions are, by construction, within **A**'s operational scope.

Faithfulness. Under the admissibility quotient, $\eta \cong 0$ implies $\eta = 0$ as an algebraic element of **A**, since [0] is the unique class with vanishing marginals.

Pointwise cancellability (Theorem 3.7 of CSG). For any non-classical regular composition rule and any $\psi \in \mathcal{P}$, $\lambda \in \Lambda$ with $w(\psi, \lambda) > 0$, there exist $\varphi \in \mathcal{P}$ and $t \in (0,1]$ such that $w(\psi \star_t \varphi, \lambda) =$

0. This drives individual configuration weights to zero via composition but does not produce a global additive inverse in **A**. PAR fills this gap.

---

3. The Separation of Levels and Two Notions of Reversibility

The pre-factual/factual distinction is foundational to the reconstruction. It rests on a **separation of levels**:

(SL) Separation of Levels. Irreversibility is the exclusive signature of the commitment process φ. The pre-factual sector $\mathcal{P}$, together with its admissible dynamics and recombination algebra **A**, is free of irreversibility.

Facts are defined as stable, irreversible records; the pre-factual sector is the domain of reversible alternatives. (SL) makes this explicit.

Irreversibility has two distinct manifestations, both of which (SL) excludes from $\mathcal{P}$:

Dynamical irreversibility. A preferred temporal direction in the composition family $\star$_t. This is excluded by CSG, which establishes time-symmetry of the composition family.

Algebraic irreversibility. A physically real pre-factual state transition — connecting two genuinely distinct, internally realisable states — that cannot be reversed by any operation in **A**. If such a transition exists, certain physically real state changes within $\mathcal{P}$ are one-way within the algebra.

CSG addressed dynamical irreversibility. The present paper addresses algebraic irreversibility.

Why algebraic irreversibility violates (SL). Suppose $\psi = \psi^0 + r$ with $\psi, \psi^0 \in \mathcal{P}_0$, $r \in \mathcal{P}$, $r \neq 0$. Then $\psi \ncong \psi^0$ (Lemma 4.4), so $\psi \to \psi^0$ is a physically real change within the pre-factual domain. If no χ with $r + \chi \cong 0$ exists in **A**, this transition is algebraically one-way — it can occur but cannot be undone within **A**. That is precisely the irreversibility (SL) confines to the commitment map.

---

4. Irreducible Residuals

Definition 4.1 (Internally realised affine sub-contribution). A state $r \in \mathcal{P}$ is an *internally realised affine sub-contribution* of $\psi \in \mathcal{P}$ if there exists $\psi^0 \in \mathcal{P}_0$ such that $\psi = \psi^0 + r$ in the affine weight decomposition of Section 2.

Definition 4.2 (Irreducible residual). A state $r \in \mathcal{P}$ is an *irreducible residual* of ψ if:

- **(i)** $r \neq 0$
- **(ii)** r is an internally realised affine sub-contribution of ψ, for some $\psi^0 \in \mathcal{P}_0$
- **(iii)** there exists no $\chi \in \mathcal{P}_0$ such that $r + \chi \cong 0$

Remark 4.2a (Condition (iii) is automatically satisfied at the weight level). Condition (iii) is automatically satisfied at the non-negative weight level, so its significance is not local existence in $\mathcal{P}$, but the obstruction it names in the attempted lift from pre-factual weight structure to algebraic inverse structure in **A**. Concretely: condition (iii) — no $\chi \in \mathcal{P}_0$ with $r + \chi \cong 0$ — holds automatically for every non-null $r \in \mathcal{P}$ at the weight level. If $r + \chi \cong 0$, then $w(r, \lambda) + w(\chi, \lambda) = 0$ for all λ; since both terms are non-negative, this forces $w(r, \lambda) = 0$ for all λ, giving $r \cong 0$, contradicting $r \neq 0$.

This is not a triviality to be glossed over; it is the precise mechanism the proof exploits. Definition 4.2 is formulated this way because irreducibility is the natural obstruction concept for the passage from the pre-factual weight structure to the algebraic inverse structure in **A**: condition (iii) names the obstruction, while Theorem 6.1 shows that PAR eliminates it on the non-trivial domain. The significance of the irreducible residual concept is exactly that PAR forbids such objects from existing within its scope: Theorem 6.1 derives a contradiction from the assumption that an irreducible residual exists in a non-trivial decomposition by showing — via PAR — that a weight-level cancellation partner would have to exist, and then weight-level non-negativity forces $r = 0$, contradicting condition (i). The signed inverses that (S2) asserts in **A** live at the algebraic level, beyond the non-negative weight space of $\mathcal{P}$.

Remark 4.2b (Scope relation between irreducible residuals and PAR). Definition 4.2 permits $\psi^0 \in \mathcal{P}_0$, which includes the formally trivial case $\psi^0 = 0$. By contrast, PAR applies only to non-trivial internally realised decompositions with $\psi^0 \neq 0$. Thus PAR governs only a proper subset of the cases permitted by Definition 4.2: it handles irreducible residuals arising in non-trivial decompositions, but is silent about the formally trivial case $\psi^0 = 0$. Theorem 6.1 is stated accordingly, covering only PAR's scope. The remaining cases are brought under PAR's scope in Section 7 via Compositional Completeness (CC).

Lemma 4.3 (Residual detectability). If $r \neq 0$, then there exists $\lambda_0 \in \Lambda$ such that $w(r, \lambda_0) > 0$.

Proof. Suppose $w(r, \lambda) = 0$ for all $\lambda \in \Lambda$. Then $r \cong 0$. By faithfulness, $r = 0$ as an algebraic element, contradicting $r \neq 0$. Therefore there exists $\lambda_0 \in \Lambda$ such that $w(r, \lambda_0) > 0$. □

Lemma 4.4 (Induced observational asymmetry). If r is an internally realised affine sub-contribution of ψ with $\psi = \psi^0 + r$ and $r \neq 0$, then $\psi \ncong \psi^0$.

Proof. By Lemma 4.3, there exists λ_0 with $w(r, \lambda_0) > 0$. Therefore:

$$w(\psi, \lambda_0) = w(\psi^0, \lambda_0) + w(r, \lambda_0) > w(\psi^0, \lambda_0)$$

By the characterisation of observational equivalence, $\psi \ncong \psi^0$. □

Lemma 4.5 (Reality of the associated transition). If r is an irreducible residual of $\psi = \psi^0 + r$ with $\psi^0 \neq 0$, then the transition $\psi \to \psi^0$ is physically real: it connects two observationally distinct, non-null, internally realisable pre-factual states.

Proof. $\psi \in \mathcal{P}$ and $\psi^0 \in \mathcal{P}$ with $\psi^0 \neq 0$ by assumption. By Lemma 4.4, $\psi \ncong \psi^0$. Both are non-null elements of $\mathcal{P}$. □

---

5. Pre-Factual Algebraic Reversibility

5.1 Definition and Justification of PAR

Definition 5.1 (Algebraically irreversible pre-factual transition). A transition $\psi \to \psi^0$ is *algebraically irreversible* if:

- **(a)** $\psi \in \mathcal{P}$, $\psi^0 \in \mathcal{P}_0$, $\psi^0 \neq 0$
- **(b)** $\psi \ncong \psi^0$
- **(c)** $\psi = \psi^0 + r$ for some $r \in \mathcal{P}$ with $r \neq 0$
- **(d)** there exists no $\chi \in \mathcal{P}_0$ such that $r + \chi \cong 0$

The requirement $\psi^0 \neq 0$ in condition (a) restricts Definition 5.1 to *non-trivial* transitions — those connecting two genuinely structured pre-factual states, not merely the transition from some r to the null element.

Principle (PAR) Pre-Factual Algebraic Reversibility. Every non-trivial internally realised pre-factual transition in the recombination domain admits an algebraic reversal within that domain. Equivalently: for every internally realised decomposition $\psi = \psi^0 + r$ with $r \neq 0$ and $\psi^0 \neq 0$, there exists $\chi \in \mathcal{P}_0$ such that $r + \chi \cong 0$.

Physical justification of PAR. Its justification proceeds in three steps.

Step 1: The transition is physically real. By Lemma 4.4, $\psi \ncong \psi^0$ whenever $r \neq 0$. Since $\psi, \psi^0 \in \mathcal{P}$ are both genuine, non-null pre-factual states, the transition $\psi \to \psi^0$ represents a real change in pre-factual state — the removal of the contribution r from one structured pre-factual state to another.

Step 2: The forward transition is within **A***'s operational domain by construction.* By the definition of **A** in Section 2, **A** is generated by admissibility classes of states in $\mathcal{P}_0$ under affine recombination. The decomposition $\psi = \psi^0 + r$ is one of **A**'s generating operations. The forward transition $\psi^0 \to \psi$ is within **A**'s operational scope by construction.

Step 3: (SL) requires the reverse transition, but generation alone does not ensure it. A semigroup is generated by elements without necessarily containing their inverses. What forces the reverse is (SL): any one-way algebraic transition between two non-null pre-factual states is a form of pre-factual irreversibility — a change within $\mathcal{P}$ that cannot be undone within **A**. (SL) confines all irreversibility to the commitment map φ, which lies strictly outside $\mathcal{P}$ and **A**. Therefore no non-trivial transition in **A**'s generating domain can be one-way.

Remark 5.1a (PAR is restricted to non-trivial transitions). PAR is stated for non-trivial decompositions $\psi = \psi^0 + r$ with $\psi^0 \neq 0$. This restriction is intentional: these are the transitions with genuine physical content — movements between two structured pre-factual states. The extension to all non-null r, including cases where no non-trivial ψ^0 is specified, is achieved via CC (Definition 5.4 below). This separation keeps the physical argument clean: PAR handles reversibility; CC handles reach.

Remark 5.1b (Honest scope). The logical structure is: (i) the forward transition $\psi^0 \to \psi$ is in **A**'s domain by construction; (ii) if its reverse is absent, that is algebraic irreversibility within $\mathcal{P}$; (iii) (SL) forbids pre-factual algebraic irreversibility; therefore (iv) the reverse must be present. Step (i) uses the definition of **A**. Steps (ii)–(iv) use (SL) via PAR. Generation does not yield inverse-closure; PAR does. Whether (SL) and PAR can be derived from (A1)–(A2) alone is the primary open question.

5.2 Compositional Completeness (CC)

PAR, as defined, covers non-trivial decompositions $\psi = \psi^0 + r$ with $\psi^0 \neq 0$. It therefore constrains only those non-null contributions that actually occur within the non-trivial compositional domain of $\mathcal{P}$. To extend IAC to all non-null $r \in \mathcal{P}$, we need a second axiom governing whether every element of $\mathcal{P}$ participates in that domain.

Definition 5.4 (Compositional Completeness — CC). The pre-factual sector is *compositionally complete* if every non-null $r \in \mathcal{P}$ participates in the non-trivial internal compositional domain of $\mathcal{P}$. That is: for every non-null $r \in \mathcal{P}$, there exist $\psi^0, \psi \in \mathcal{P}$ with $\psi^0 \neq 0$ such that:

$$\psi = \psi^0 + r$$

is a non-trivially internally realised affine decomposition.

Why CC is needed. PAR governs non-trivial internally realised transitions. It therefore constrains only those non-null contributions that actually occur within the non-trivial compositional domain of the pre-factual sector. If some non-null $r \in \mathcal{P}$ could never occur as a sub-contribution in any such decomposition, PAR would be silent about it. Universal IAC would then fail — not because reversibility breaks down where it applies, but because part of the formally admitted pre-factual sector would lie outside the domain to which reversibility applies.

Why CC is natural. CC is a domain-coherence condition on the meaning of membership in $\mathcal{P}$. The pre-factual sector is not intended as a merely extensional list of algebraic placeholders; it is the domain of physically real pre-factual alternatives from which facts arise. A non-null element of $\mathcal{P}$ that never appears in any non-trivial internally realised decomposition would be formally admitted as pre-factual while being compositionally inert — present in the set, but absent from the structured pre-factual processes that generate the recombination algebra **A**. Such an element would belong to $\mathcal{P}$ extensionally but not structurally. CC excludes precisely this mismatch between formal membership and compositional participation.

Why CC is modest. CC is not a preparation postulate, nor does it require that every non-null r be directly isolatable or operationally accessible on its own. It requires only that every non-null r occur *somewhere* within the internally realised compositional fabric of the pre-factual sector. It is therefore a minimal completeness requirement on the structural meaning of membership in $\mathcal{P}$.

CC is not derivable from admissibility (A2) or PAR alone. It is a new structural completeness axiom for the pre-factual sector, stated as an explicit premise and identified as an open question in Section 10.

Remark 5.5 (Independence of PAR and CC). PAR and CC address different questions and neither implies the other. PAR says that a given non-trivial internally realised contribution, once embedded in a structured pre-factual transition with $\psi^0 \neq 0$, must admit algebraic reversal within **A**. It says nothing about whether every non-null $r \in \mathcal{P}$ can be so embedded. CC says that every non-null r can be so embedded, but says nothing about whether the resulting transition is reversible. The two principles have disjoint content: PAR handles what happens once a non-trivial embedding exists; CC handles whether such an embedding exists at all.

5.3 Scope Remarks

Remark 5.6 (Scope of PAR and CC). PAR applies only to non-trivial transitions arising from internally realised affine decompositions with $\psi^0 \neq 0$. CC applies to all non-null $r \in \mathcal{P}$. Neither asserts that every operation in the composition family $\star$_t has an additive inverse, nor that all pre-factual dynamics are algebraically reversible.

Remark 5.7 (Dynamical vs. algebraic reversibility). Dynamical reversibility — established in CSG — concerns the composition family $\star$_t. Algebraic reversibility — established here via PAR and CC — concerns the additive structure of **A**. These are independent: pointwise cancellability (Theorem 3.7 of CSG) drives individual configuration weights to zero via composition but does not produce a global additive inverse in **A**. PAR and CC together fill the remaining gap.

---

6. The Core Theorem

Theorem 6.1 (No irreducible residuals within PAR's scope). Under the Separation of Levels (SL) and Pre-Factual Algebraic Reversibility (PAR), no non-null element $r \in \mathcal{P}$ that arises in a non-trivial internally realised decomposition $\psi = \psi^0 + r$ with $\psi^0 \neq 0$ is an irreducible residual.

Proof. Suppose for contradiction that r is a non-null irreducible residual of ψ arising in a non-trivial internally realised decomposition $\psi = \psi^0 + r$ with $\psi^0 \neq 0$.

By condition (i) of Definition 4.2, $r \neq 0$. By Lemma 4.5, the transition $\psi \to \psi^0$ is physically real: $\psi \ncong \psi^0$, with both states non-null elements of $\mathcal{P}$. By condition (iii) of Definition 4.2, no $\chi \in \mathcal{P}_0$ satisfies $r + \chi \cong 0$.

Checking Definition 5.1: condition (a) holds since $\psi \in \mathcal{P}$ and $\psi^0 \in \mathcal{P}$ with $\psi^0 \neq 0$; condition (b) holds by Lemma 4.4; condition (c) holds by assumption; condition (d) holds by condition (iii). The transition $\psi \to \psi^0$ is therefore algebraically irreversible in the sense of Definition 5.1. Since $\psi^0 \neq 0$, this transition falls within PAR's declared scope. PAR is contradicted.

Therefore no such r exists. □

Remark 6.1a (Why Theorem 6.1 does not yet yield full IAC). Theorem 6.1 applies only to irreducible residuals arising from non-trivial internally realised decompositions — that is, decompositions with $\psi^0 \neq 0$, which is exactly PAR's scope. It does not yet exclude the possibility that a non-null r might arise only via the formally trivial decomposition $r = 0 + r$ with $\psi^0 = 0$. For such an r, Definition 4.2 is satisfied but PAR has nothing to say, because the associated transition $(r \to 0)$ has $\psi^0 = 0$ and therefore falls outside PAR's scope. The extension from Theorem 6.1's restricted conclusion to full IAC is carried out in Theorem 7.1 using Compositional Completeness (CC).

---

7. Internal Admissible Closure

The conclusions of this section are stated directly at the level of the amplitude algebra **A**. As established in Remark 4.2a, no non-null $r \in \mathcal{P}$ can have a weight-level cancellation partner in $\mathcal{P}_0$ — non-negativity alone prevents it. The weight-level argument is used in Theorem 6.1 to rule out irreducible residuals; the algebraic inverses asserted below live in **A**, not in $\mathcal{P}$ as a weight space.

Theorem 7.1 (Full Internal Admissible Closure from PAR and CC). Assume SL, PAR, and CC. Then no non-null element of $\mathcal{P}$ is an irreducible residual. Consequently, every non-null admissibility class $[r] \in \mathbf{A}$ admits an additive inverse $[-r]$ such that $[r] + [-r] = [0]$ in **A**.

Proof. Let $r \in \mathcal{P}$ be any non-null element. By CC (Definition 5.4), there exist $\psi^0, \psi \in \mathcal{P}$ with $\psi^0 \neq 0$ such that $\psi = \psi^0 + r$ is an internally realised non-trivial decomposition.

Suppose r remained an irreducible residual — that is, suppose condition (iii) of Definition 4.2 holds: there exists no $\chi \in \mathcal{P}_0$ with $r + \chi \cong 0$. Together with $r \neq 0$ and the CC embedding $\psi = \psi^0 + r$ with $\psi^0 \neq 0$, conditions (i) and (ii) of Definition 4.2 are also met. Hence r would be an irreducible residual arising in a non-trivial internally realised decomposition, contradicting Theorem 6.1. Therefore r is not an irreducible residual: condition (iii) fails, and there exists $\chi \in \mathcal{P}_0$ with $r + \chi \cong 0$.

By faithfulness, this observational nullity lifts to algebraic nullity, so $[r] + [\chi] = [0]$ in **A**. Hence $[\chi]$ is the additive inverse of $[r]$. Since r was arbitrary, every non-null class in **A** has an additive inverse. The null class $[0]$ is its own inverse. □

Remark 7.2 (Why CC replaces the trivial decomposition). In earlier versions, the step from PAR's non-trivial scope to full IAC was taken via the trivial decomposition $r = 0 + r$. That move is purely formal — it treats the null element as a legitimate ψ^0, which it is not in the physically motivated sense of PAR. CC replaces this formal step with a structural completeness premise: every non-null r is embeddable in a genuine non-trivial pre-factual decomposition, placing it squarely within PAR's scope. The proof chain is then: CC embeds r non-trivially → Definition 4.2 conditions met → Theorem 6.1 rules out the irreducible residual → the cancellation partner exists → $[-r]$ is the inverse in **A**.

Corollary 7.3 (S2 in the amplitude algebra). Under the conditions of Theorem 7.1:

$$\forall [r] \in \mathbf{A}, [r] \neq [0] \Longrightarrow \exists [\chi] \in \mathbf{A} \text{ such that } [r] + [\chi] = [0]$$

Condition (S2) holds in the amplitude algebra **A**.

Proof. Immediate from Theorem 7.1. □

Remark 7.4 (Why the conclusion lives in A, not $\mathcal{P}$). As noted in Remark 4.2a, at the weight level no non-null $r \in \mathcal{P}$ can have a cancellation partner in $\mathcal{P}_0$: non-negativity forces any solution to $r + \chi \cong 0$ to have $w(r, \lambda) = w(\chi, \lambda) = 0$ for all λ, giving $r = 0$ by faithfulness — contradicting $r \neq 0$. The weight-level impossibility is the reductio that rules out irreducible residuals in Theorem 6.1. The inverses $[-r]$ are features of the extended algebra **A**, which carries signed structure beyond the non-negative weight space of $\mathcal{P}$. This is the standard pattern in quantum reconstruction: the probability-level structure is used to establish the algebraic-level structure, not to model it directly.

---

8. Consequences for Algebraic Structure

From Theorem 7.1:

- Every non-null admissibility class in **A** admits an additive inverse: $[\psi] + [-\psi] = [0]$
- **A** carries the structure of an abelian group under addition
- **A** is globally cancellable
- Equivalently, no non-null element of $\mathcal{P}$ remains as an irreducible residual

Combined with the previously established results of the VERSF programme, the division-algebra conclusion of the Bridge Theorem follows: **A** is a finite-dimensional division algebra over $\mathbb{R}$, and by the Frobenius theorem the only candidates are $\mathbb{R}$, $\mathbb{C}$, $\mathbb{H}$.

---

9. Relation to the VERSF Reconstruction Programme

The complete derivation chain now reads:

Step	Principle	Result
Fact-production	Facts are stable, irreversible records	Pre-factual / factual distinction
Admissibility (A2)	Only fact-grounded distinctions are meaningful	Commitment-traceable observables
Operational congruence (CSG §3.4)	Admissible dynamics act on equivalence classes	MCC
MCC + Cauchy equation	Marginal compositional consistency	Interference structure
(SL) — dynamical component	Irreversibility confined to φ	Dynamical reversibility of $\mathcal{P}$
(SL) — algebraic component	One-way non-trivial transitions violate (SL)	PAR (§5.1)
CC	Every non-null r embeds non-trivially in $\mathcal{P}$	Every r falls within PAR's non-trivial scope
PAR + CC → Theorem 6.1 + Theorem 7.1	No irreducible residuals; inverses exist in **A**	IAC + (S2) in **A** (Corollary 7.3)
(S2) + (A4) + Frobenius	Division algebra over $\mathbb{R}$	$\mathbb{R}$, $\mathbb{C}$, or $\mathbb{H}$
Distinguishability geometry	Excludes $\mathbb{R}$ and $\mathbb{H}$	Complex Hilbert space

The remaining ungrounded assumptions are: (A2c), metric homogeneity (A3), commutativity (C5), finite-dimensionality, the semigroup law (v), exchange symmetry (vi), and the imported distinguishability geometry results.

---

10. Structural Status of PAR and CC

The derivation of Internal Admissible Closure (IAC) in this paper rests on two structural principles: Pre-Factual Algebraic Reversibility (PAR) and Compositional Completeness (CC). Neither is derived from admissibility (A2) alone. This section makes their status explicit: what they assume, why they are natural, and how they relate to IAC.

10.1 Why Additional Structure Is Unavoidable

The core problem is not whether cancellation partners *can* exist, but whether they are *forced*. Admissibility (A2) constrains which distinctions are physically meaningful. It does not guarantee that every internal transition is reversible, or that every formal element of the pre-factual sector participates in the internal compositional structure. This gap is precisely what IAC fills. Any derivation of IAC must therefore introduce additional structural content beyond admissibility alone. The role of this paper is to make that additional structure explicit, minimal, and physically interpretable — rather than implicit or hidden inside algebraic assumptions.

10.2 Pre-Factual Algebraic Reversibility (PAR)

Statement. No non-trivial internally realised pre-factual transition is algebraically one-way.

Why PAR is natural. PAR follows directly from the central organising principle of the framework: irreversibility belongs exclusively to facts. A non-trivial transition between two structured pre-factual states ($\psi \to \psi^0$ with $\psi^0 \neq 0$) is physically real, internal to the pre-factual domain, and independent of commitment. If such a transition were one-way within the algebra, irreversibility would already be present before any fact is formed. This collapses the distinction between possibility (reversible, pre-factual) and fact (irreversible, committed). PAR is therefore the algebraic expression of the separation between possibility and fact.

Why PAR is not trivial. PAR is strictly stronger than dynamical time-symmetry and pointwise cancellability. Those results allow weights to be redistributed or locally eliminated via composition, but do not guarantee a global algebraic inverse in **A**. PAR adds exactly the missing condition: internal transitions are not merely evolvable, but reversible within the algebra that describes them.

10.3 Compositional Completeness (CC)

Statement. Every non-null pre-factual contribution participates in the non-trivial compositional domain of $\mathcal{P}$.

Why CC is needed. PAR applies only to non-trivial decompositions $\psi = \psi^0 + r$ with $\psi^0 \neq 0$. Without CC, it is logically possible that some non-null $r \in \mathcal{P}$ exists but never appears inside any non-trivial internally realised state. Such an r would belong to $\mathcal{P}$ formally but lie outside the domain where PAR applies — PAR cannot constrain it, and the derivation of IAC fails for those elements. CC closes this gap.

Why CC is natural. CC should be read as a coherence axiom on the pre-factual sector itself. Without it, the theory admits a mismatch between formal membership and structural participation: some non-null states would count as pre-factual by stipulation, yet never occur in any internally realised non-trivial decomposition. Such elements would lie outside the compositional domain that generates **A**, and hence outside the domain governed by PAR. CC removes this pathology by requiring that every non-null state in $\mathcal{P}$ participate somewhere in the internal structure of the pre-factual sector. In this sense CC is not an arbitrary strengthening, but the minimal bridge from extensional membership in $\mathcal{P}$ to compositional relevance within $\mathcal{P}$.

Why CC is not trivial. CC is a new structural completeness axiom for the pre-factual sector. It is not implied by admissibility (A2), recombinability (R2), or the definition of $\mathcal{P}$ as a set. No specific existing condition within the current framework is known to entail CC; the route to a derivation is genuinely open and does not yet have a named candidate axiom comparable to (R2)'s role for PAR. Crucially, CC does not enlarge the physical content of the theory; it restricts the interpretation of membership in $\mathcal{P}$ so that the pre-factual sector contains no compositionally idle elements. It is therefore a consistency condition on domain meaning rather than an arbitrary bolt-on assumption.

10.4 Why PAR + CC Is the Right Level of Strength

A key concern is whether PAR and CC are simply IAC restated. They are not identical to IAC:

- PAR constrains reversibility of non-trivial transitions
- CC ensures all elements lie within that non-trivial domain
- Only together do they imply that every non-null element admits a cancellation partner

Each principle is individually consistent with IAC's failure: PAR alone leaves compositionally disconnected elements unconstrained by any reversibility requirement; CC alone says nothing about whether any transition is reversible. Only their conjunction forces IAC.

The structure is therefore:

PAR (no premature irreversibility within non-trivial transitions)
+ **CC** (every non-null element lies in the non-trivial domain)
⇒ **IAC** (universal algebraic cancellation)

This separation matters. It shows that IAC is not a monolithic assumption but the result of two conceptually distinct requirements: no premature irreversibility, and no disconnected pre-factual structure. The gap from prior work has been isolated into two precise conditions rather than collapsed into a single undifferentiated closure postulate.

10.5 Honest Logical Status

The correct logical position is:

- IAC is not derived from admissibility alone
- IAC follows from PAR + CC
- PAR and CC are physically motivated, structurally minimal for the task, but not yet derived from (A1)–(A2)

This is not a weakness — it is a clarification. The original gap has not been hand-waved away; it has been isolated into two precise structural principles. One possible route for PAR: (R2) recombinability, if strengthened to require transitive closure over physically real pre-factual transitions, may yield it; this connection is not formalised here.

10.6 Secondary Open Questions

- Whether PAR and CC can be unified under a single condition
- Whether dynamical and algebraic reversibility can be unified under a single condition
- Extension of PAR and CC beyond the single-configuration commitment framework

11. Conclusion

We have shown that internal admissible closure follows from two physically motivated structural principles — Pre-Factual Algebraic Reversibility (PAR) and Compositional Completeness (CC) — both of which are required by the foundational distinction between possibility and fact.

The argument rests on recognising that irreversibility has two faces: dynamical and algebraic. Previous work established dynamical reversibility of the pre-factual sector. The present paper addresses algebraic reversibility through a two-principle structure. PAR says that every non-trivial internally realised pre-factual transition admits an algebraic reversal — because one-way transitions between structured pre-factual states would constitute pre-factual irreversibility, violating the Separation of Levels. CC says that every non-null pre-factual contribution belongs to the internal compositional domain of $\mathcal{P}$ — it can occur as a sub-contribution within some non-trivial internally realised pre-factual state, meaning its formal membership in $\mathcal{P}$ matches its internal compositional relevance. Together: CC places every non-null r within PAR's non-trivial scope, and PAR then guarantees the cancellation partner. Full IAC follows.

The honest outcome is therefore this: internal admissible closure follows from two physically motivated requirements on the pre-factual domain. PAR expresses the algebraic side of pre-factual reversibility: no non-trivial internally realised transition within the pre-factual domain may be one-way. CC expresses structural completeness of that domain: every non-null pre-factual contribution must be embeddable in some non-trivial internally realised pre-factual state, placing it within PAR's scope. In this strengthened sense of pre-factual reversibility — dynamical, algebraic, and structurally complete — internal admissible closure follows. The remaining open question is whether PAR and CC can themselves be derived from the more primitive admissibility conditions (A1)–(A2), or whether they stand as independent structural principles of fact-producing universes.